

Estimating the Cost of a GIS in the Amazon™ Cloud

An Esri® White Paper
August 2012



Copyright © 2012 Esri
All rights reserved.
Printed in the United States of America.

The information contained in this document is the exclusive property of Esri. This work is protected under United States copyright law and other international copyright treaties and conventions. No part of this work may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or by any information storage or retrieval system, except as expressly permitted in writing by Esri. All requests should be sent to Attention: Contracts and Legal Services Manager, Esri, 380 New York Street, Redlands, CA 92373-8100 USA.

The information contained in this document is subject to change without notice.

Esri, the Esri globe logo, ArcGIS, esri.com, and @esri.com are trademarks, service marks, or registered marks of Esri in the United States, the European Community, or certain other jurisdictions. Other companies and products or services mentioned herein may be trademarks, service marks, or registered marks of their respective mark owners.

Estimating the Cost of a GIS in the Amazon Cloud

An Esri White Paper

Contents	Page
Introduction.....	1
What Is ArcGIS Server on Amazon Web Services?.....	1
What Is Cloud Builder?	1
Choosing Services Offered on Amazon.....	1
Amazon Elastic Block Store	1
Elastic IP Addresses.....	2
Elastic Load Balancing	2
Auto Scaling.....	2
Amazon CloudWatch.....	2
Calculating AWS Data Center Costs	2
Hourly Rates and Data Transfer	2
The Cost of Underutilized Resources	3
The Cost of Time	3
Time to Set Up.....	3
Time to Maintain.....	3
Choosing an Instance	4
An Example Implementation	5
Using ArcGIS Server AMI and AWS Management Console.....	5
Example Deployment Components	5
ArcGIS Server AMI.....	5
Custom AMI	5
Elastic Load Balancing	6
Amazon CloudWatch plus Auto Scaling	6
Update File Geodatabase	6
Using ArcGIS Server AMI and Cloud Builder.....	7
How Costs Are Incurred When Using Cloud Builder	7
Estimating Costs	9
Instance Costs	9

Contents	Page
Storage and Data Transfer Costs.....	11
Elastic Block Store Costs.....	12
Elastic IP Address Costs	13
Elastic Load Balancing Costs	13
CloudWatch Costs	13
Amazon Calculators.....	14
Amazon EC2 Bill Sample.....	14
Summary	15

Estimating the Cost of a GIS in the Amazon Cloud

Introduction

Geographic information system (GIS) and information technology (IT) managers see the cloud as a compelling argument for leveraging off-premises hardware resources and reducing local infrastructure costs. But how does one go about calculating the cost of using cloud resources in lieu of using on-site infrastructure?

This paper will help the reader understand how costs are calculated for cloud services—using Amazon™ Elastic Compute Cloud (EC2)—and how to estimate the costs of using Amazon cloud services with an ArcGIS® for Server cloud implementation: ArcGIS for Server on Amazon Web Services (AWS®).

What Is ArcGIS Server on Amazon Web Services?

ArcGIS for Server is the core server GIS software made by Esri, used for creating, managing, and distributing GIS web services, applications, and data. ArcGIS for Server is typically deployed on premises within the organization's service-oriented architecture (SOA) or off premises in cloud computing environments or data centers.

Esri offers ArcGIS for Server on AWS as a preconfigured solution for organizations wishing to deploy their GIS servers on established cloud infrastructures rather than use their own local infrastructures. ArcGIS Server on AWS provides three Amazon Machine Images (AMIs): ArcGIS for Server and an enterprise geodatabase configured on PostgreSQL, both running in Ubuntu; ArcGIS for Server configured on Microsoft SQL Server Express; and an enterprise geodatabase configured on Microsoft SQL Server Standard, the latter two running in Windows 2008 R2.

An AMI represents a machine image or the state of a machine once it's completely configured for a certain need. In the case of ArcGIS for Server on AWS, the ArcGIS Server AMI is used to launch instances of ArcGIS Server. Each instance is configured to run on Windows® 2008 R2 and includes a 100-gigabyte (GB) drive for GIS data. However, additional configuration options are available. These options are covered in more detail later.

What Is Cloud Builder?

ArcGIS Server Cloud Builder on AWS helps users build and maintain an ArcGIS for Server site on AWS. It is a lightweight desktop application downloadable by anyone who has purchased an ArcGIS for Server license.

When a site is built in AWS, various components must be created, connected, and maintained. In previous releases of ArcGIS for Server, these tasks were performed manually in the AWS Management Console, a web application developed by Amazon. At ArcGIS 10.1, Cloud Builder hides this complexity and simplifies the process of assembling and administering sites.

Choosing Services Offered on Amazon

Amazon offers optional services that can be associated with an instance or group of instances. The pricing for these services is based on a pay-as-you-go pricing model.

Amazon Elastic Block Store

Elastic Block Store (EBS) is the storage that can be attached directly to an instance to become a local drive to the machine. The EBS can be detached and managed separately as needed. When an instance is launched from the AMI, it automatically includes

attached storage that is ready to use for any custom data. For example, the ArcGIS for Server AMI and the enterprise geodatabase AMI each include 100 GB of EBS. The EBS has a life cycle that is separate from the instance so that if the instance is terminated, the EBS volume continues to exist. This allows users to retain and manage their data separately from the instance. A typical use of EBS is to create a snapshot or backup copy of project data. For more about Elastic Block Store, visit aws.amazon.com/ebs/.

Elastic IP Addresses

Elastic IP addresses provide an IP address that never changes inside the cloud. Typically, when a new instance is launched, it is assigned a temporary external IP address. However, each time the instance is terminated or replaced, this external IP address will change. In some cases, it is necessary to assign an IP address that remains static. The elastic IP address is associated with an account and not a particular instance. It can be maintained until the user chooses to explicitly release it. For example, an elastic IP address can be assigned to a user's AWS account, then mapped to a particular ArcGIS Server instance when ArcGIS for Server on AWS is used. The IP address can be remapped to a replacement ArcGIS Server instance.

Elastic Load Balancing

Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances. It enables users to achieve even greater fault tolerance in their applications and provides the amount of load balancing capacity needed in response to incoming application traffic. Elastic Load Balancing can be associated with multiple instances within a single availability zone or across multiple zones to maintain application performance or to compensate for instances that stop working properly. For example, Elastic Load Balancing can be associated with a set of ArcGIS Server instances that support high-traffic or mission-critical applications. For more information about Elastic Load Balancing, visit aws.amazon.com/elasticloadbalancing/.

Auto Scaling

Auto Scaling ensures that the number of Amazon EC2 instances being used scales up (increases) during demand spikes to maintain performance, then automatically scales down during traffic lulls to minimize costs. Auto Scaling is enabled by Amazon CloudWatch and available at no additional charge beyond Amazon CloudWatch fees. For more information about Auto Scaling, visit aws.amazon.com/autoscaling/.

Amazon CloudWatch

Amazon CloudWatch provides visibility into AWS cloud resource utilization, operational performance, and overall demand patterns including metrics such as CPU utilization, disk reads and writes, and network traffic. For example, users can monitor use of their ArcGIS Server instances in real time via the AWS Management Console, including EBS volumes and Elastic Load Balancing. For more information about Amazon CloudWatch, visit aws.amazon.com/cloudwatch/.

Calculating AWS Data Center Costs

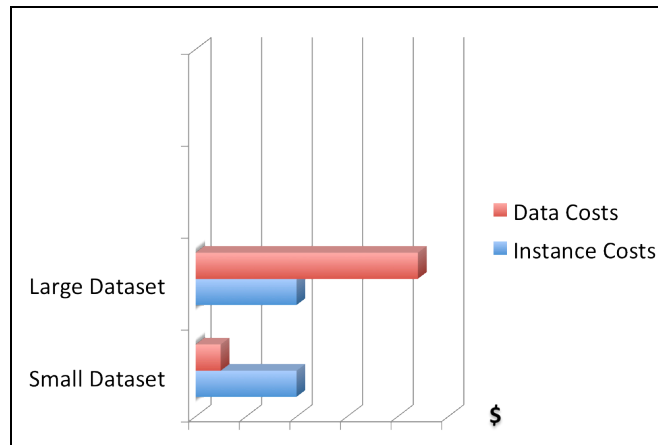
First, the general concepts that contribute to the cost of cloud computing must be taken into consideration.

Hourly Rates and Data Transfer

The costs of running an instance on hosted infrastructure depend primarily on two factors—the hourly rate of the instance type and the amount of data transferred to and from the instance. In general, if the data collection is small (e.g., no larger than 10 GB) and the expected network traffic to/from the instance is low (e.g., only 2 to 10 users are accessing the services a few times a day), the majority of the cost will be generated by the hourly rate for the instance type.

J-9979

However, as data size and network traffic to and from the instance increase, the relationship between hourly rates and data transfer rates changes. With large datasets and high traffic, the majority of costs are generated by the data transfer and storage rate, not the hourly rate of the instance.



The Cost of Underutilized Resources

On AWS, users can realize cost savings by ensuring only the resources needed are used and they are used only when needed. For example, users can manage cloud resources by planning to use a smaller machine during times of low traffic, then add more or larger machines when traffic reaches a certain threshold. Starting machines and letting them run even when they're not in use or when they're not fully utilized will not allow an organization to fully realize cost savings in the cloud.

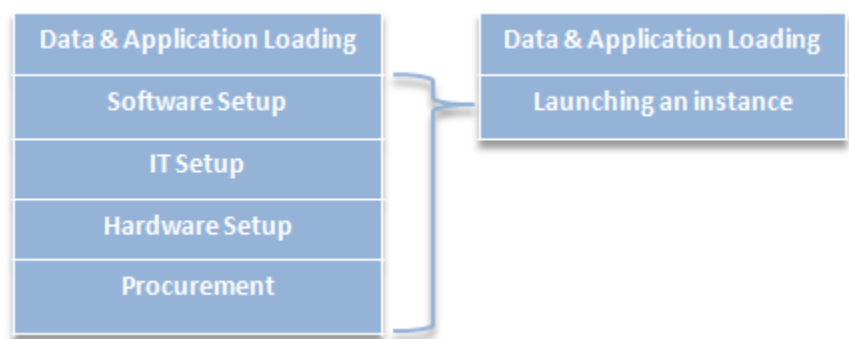
The Cost of Time

When estimating cloud computing costs, time savings is an often overlooked cost factor. Typical time-saving variables include time to set up and maintain cloud resources.

Time to Set Up

When implementing on-premises hardware, considerable time is spent getting a machine ready to use. These steps include procurement of hardware and software, installation and configuration of the operating system and software, and network configuration.

With the Amazon cloud, these steps are combined into one step—launching a specific instance type on Amazon EC2. The user is allocated a virtual machine that is ready to use for data and application setup.



Time to Maintain

Setup time can be further reduced by creating an AMI, which contains applications, libraries, data, and associated configuration settings. The AMI can be used to launch a

new instance. If at any time the instance fails, it can be terminated, and a new instance can be launched with no issues. This minimizes the time spent maintaining a machine after its initial launch.

Also, when new versions of software, updated applications, or data and maintenance patches are required, new AMIs can be created. The time used to create the new AMIs will not cause downtime for those who use the instances, because the AMI creation can be done simultaneously. Once the new AMI is available, it can be used immediately.

Choosing an Instance

Beyond the general concepts of data transfer, underutilized resources, and time, Amazon cloud consumers must consider the types of instances they are using as well as the other (optional) AWS that they may add to their solutions. Both instances and services are priced based on hourly use.

Amazon EC2 offers a number of predefined instance types from which to choose. Each instance type has a known configuration that defines the processing and memory capacity. Instance types are grouped into *families* such as Standard Instances, Micro Instances, High-Memory Instances, and High-CPU Instances.

For example, within the Standard Instances family, a user can choose between small, large, and extra-large instances. The configuration of the large instance type has 7.5 GB of memory, four EC2 compute units (two virtual cores with two EC2 compute units each), 850 GB of local instance storage, and a 64-bit platform.

Note: The EC2 compute unit is generally equivalent to a 1.0–1.2 GHz 2007 Opteron or 2007 Xeon processor.

It's also worth noting that the size of the hard drive can be easily altered, as EBS drives can be attached and removed as required. Users will only pay for the gigabytes they use. To see a complete list of each of the instance types and their configurations, visit aws.amazon.com/ec2/#instance.

Determining an instance type that provides the right capacity for applications that run in the cloud is similar to choosing the capacity of virtualized machines on premises.

On Amazon EC2, each instance type has an hourly cost that is dependent on the operating system and the software license associated with that instance. The ArcGIS Server AMI is preconfigured to run only with 64-bit instance types, and it is based on Windows 2008. It is assumed that Amazon cloud consumers will provide their own ArcGIS for Server licenses to enable use of the Esri® preconfigured AMIs. Consequently, the hourly cost for running an instance on Amazon's infrastructure is based solely on the instance type and the license fee for using Windows.

For a complete list of the hourly cost per each instance type, visit aws.amazon.com/ec2/#pricing.

An Example Implementation

In this example, ArcGIS for Server on AWS is used to support a web mapping application used by government officials, the public, and the media. Ideally, the application should be accessible 24 hours a day, seven days a week, and handle heavy traffic loads intermittently without degradation in performance. The IT and GIS staff want to be able to monitor cloud resource use and update the application and geodatabase on a quarterly basis. The organization has licensed ArcGIS for Server.


The following sections discuss two different implementations of the ArcGIS Server AMI. The first approach uses AWS Management Console. The second approach uses ArcGIS Server Cloud Builder, which leverages the new capabilities of ArcGIS 10.1 for Server.

Using ArcGIS Server AMI and AWS Management Console

The ArcGIS Server AMI is used to launch an instance that will publish and maintain three separate map services. The data used by the services is stored in a file geodatabase about 80 GB in size, and the web application is built on ArcGIS API for Flex™. The data and the application are also stored in the cloud as part of the instance.

Example Deployment Components

- ArcGIS Server AMI
- ArcGIS Server instance running
 - A file geodatabase
 - Three published services
 - The web application
- Custom ArcGIS Server AMI
- Amazon Elastic Load Balancer
- Amazon CloudWatch plus Auto Scaling



ArcGIS Server AMI

The ArcGIS Server instance is configured to include the application, the services, and the file geodatabase. The instance of ArcGIS Server is started using the Standard Large Instance type. This is the configuration from which additional production instances will be started. Users have the option to choose between using Linux or Windows when they start the instance.

Custom AMI

Once the ArcGIS Server instance is configured, it should be backed up. A custom AMI that is the exact copy of the newly reconfigured ArcGIS Server instance is created to preserve the configuration settings. This custom AMI can be used to restore/launch new instances if the production instances fail.

Note: Creating a custom AMI copies any EBS volumes that may have been attached. Be aware that this will affect costs when the custom AMI is deployed.

Elastic Load Balancing

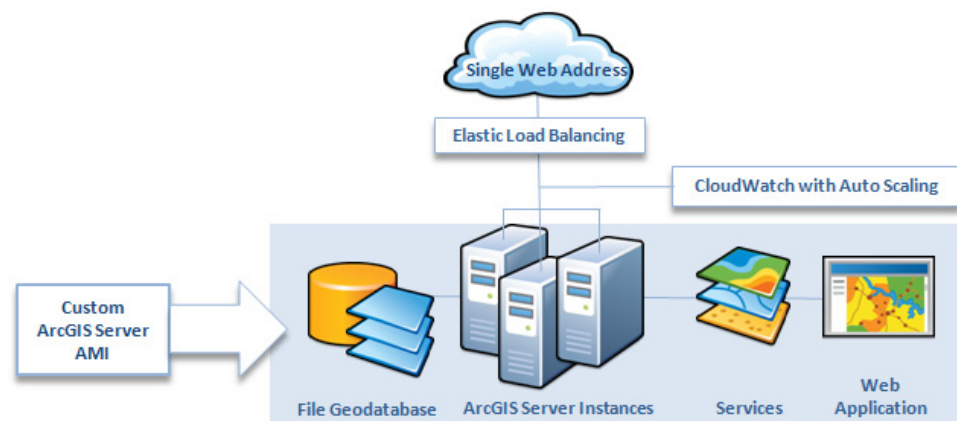
In this type of scenario, it is usually necessary for the solution to automatically launch new instances of ArcGIS Server as they are needed, reduce instances when they are not needed, and have Elastic Load Balancing (ELB) automatically unite the multiple instances under a common web address. In this way, the solution can harness the power of all the instances under the ELB to respond to requests. The production instance is added to the ELB.

Note: Adding the ELB to a configuration will increase the costs when the solution is deployed.

Amazon CloudWatch plus Auto Scaling

Amazon CloudWatch can also be configured for the instances; it produces statistics and graphs describing the utilization of the instance resources. To add and remove instances automatically, the Amazon Auto Scaling API can be employed. The Auto Scaling service can be set up to start and terminate instances according to specified criteria (e.g., when the processor average utilization exceeds 70% capacity for 10 minutes).

Note: Amazon CloudWatch will increase the costs of the solution. Auto Scaling is enabled by Amazon CloudWatch and carries no additional fees.



Update File Geodatabase

As data changes, the file geodatabase in the ArcGIS Server custom AMI may need to be updated to reflect those changes in the services and application. A new custom AMI that contains all the components, including the updated geodatabase, can be created. The new instances launched from the new AMI would have to be added to the ELB to replace the old instances.

Alternatively, EBS can be used to do data updates by directly updating the files on the EBS drive of a running instance. It's worth noting that in this case, every time a new instance is launched, the files on the EBS drives would need to be updated. Since the AMI will not include a backup of this update, snapshots can be used to create a backup copy of just the EBS drive.

Which option to use will depend on how often the data update is needed and the quantity of the data. If the data is updated often, the user might opt for the latter approach to minimize the amount of work and number of AMI versions.

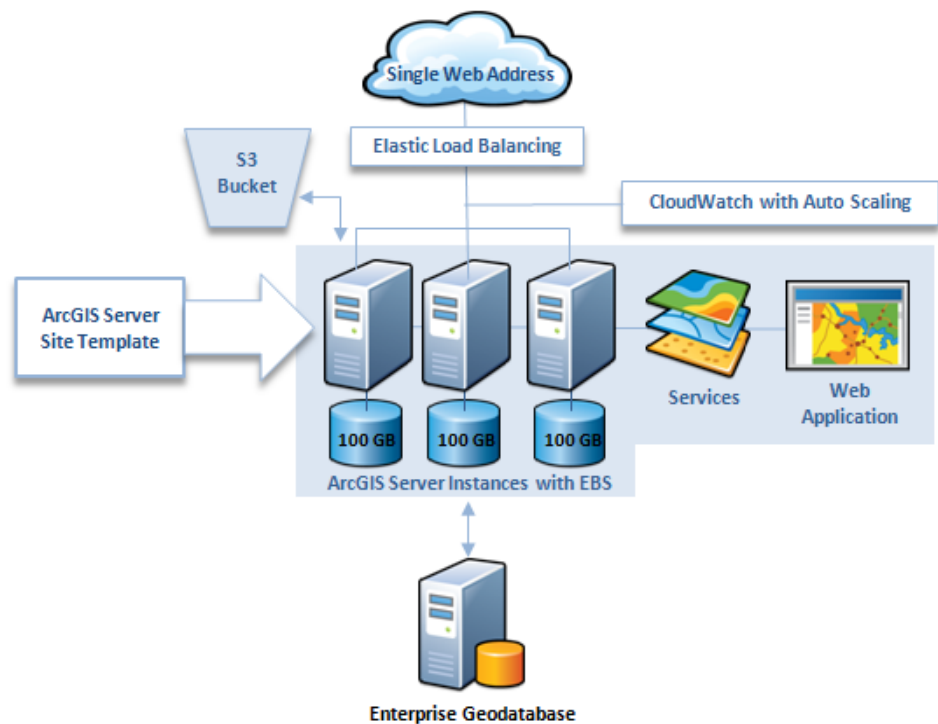
If the data is updated regularly and frequently, it's best to utilize a database. The database can then run as an additional instance to the ArcGIS Server instance. Users who want to achieve high availability for the database can run two database instances that mirror each other.

Using ArcGIS Server AMI and Cloud Builder

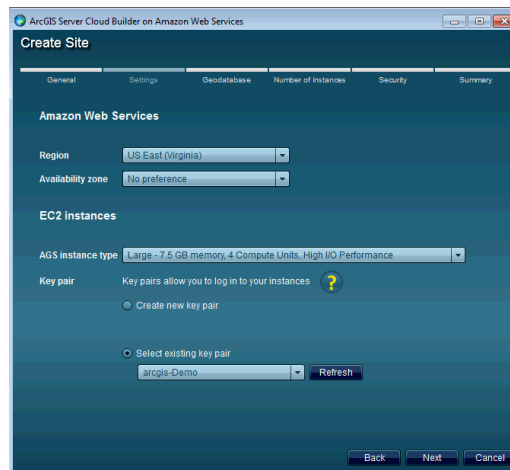
How Costs Are Incurred When Using Cloud Builder

Cloud Builder works as a wizard. Initially, users are prompted to configure it with their Amazon Web Services account. This is where the costs will be incurred when they deploy ArcGIS for Server. As Cloud Builder creates new sites and site templates, it will save information in its own configuration files. These configuration files (which are a few kilobytes in size) are located on S3 of the account Cloud Builder is configured to use.

Cloud Builder should be used not only for building an ArcGIS Server site but also for maintaining and backing it up. Cloud Builder can also create an ArcGIS Server site template that includes a customized version of ArcGIS for Server with the services and web application. The template incurs a storage fee. Charges will be levied from both Amazon EC2 and Amazon Simple Storage Service (S3) that pertain to the template. Behind the scenes, Cloud Builder is creating two different AMIs for each site template. It's also updating S3 with configuration file update to help re-create the site through Cloud Builder.



Users can choose to run an instance type of their choice using Cloud Builder; they can also specify the Amazon EC2 regions where those instances will run. Starting an instance in a specific region using Cloud Builder doesn't incur any extra cost, other than the costs charged by Amazon for that instance in that region. The cost in this case would be similar to starting an instance from the ArcGIS Server AMI using the AWS Management Console. The same cost applies when starting or stopping a site, which corresponds to stopping and starting the instances running that site.



However, it's important to note that Cloud Builder will enable detailed monitoring for Amazon EC2 instances using Amazon CloudWatch for each of those instances started, and the account will be charged for that cost. The cost for the detailed monitoring is \$3.50 per instance per month, provided at one-minute frequency. For the most updated prices, refer to aws.amazon.com/ec2/. CloudWatch is used by Cloud Builder to leverage Auto Scaling. There is no extra cost for using Auto Scaling; the cost of the instances that will be started or stopped by Auto Scaling is what is charged to a user account.



Also, it's worth noting that Cloud Builder will always create an Elastic Load Balancer when a site is created. An ELB is the main entry point for the ArcGIS Server site in the Amazon cloud. The ELB not only distributes load but also acts as a proxy to forward incoming requests on port 80 to port 6080 (which is used by ArcGIS for Server).

When starting a site, users can choose one of two operating systems: Linux or Windows. Note that each of these operating systems has a different price when used in AWS. Users can choose to start an enterprise geodatabase based on PostgreSQL or Microsoft SQL Server Express, depending on their OS choice. No extra license cost is incurred for either database. Additionally, users who choose to run Windows OS with an enterprise geodatabase can run Microsoft SQL Server Standard.

Note: When using this type of instance, the hourly rate includes the Microsoft SQL Server Standard license, which makes the hourly rate run higher than a Windows instance without it. For more information about the cost of running Microsoft SQL Server Standard with each instance size, review the information at this link: aws.amazon.com/marketplace/pp/B007NML5NE/ref=mkt_ste_ec2.



Users can also [create backups of their ArcGIS Server site](#) using Cloud Builder. ArcGIS Server Cloud Builder on Amazon Web Services allows users to create a [backup](#) that preserves the current state of their site. If an enterprise geodatabase instance is included in the site, it also participates in the backup. Restoring a backup is an easy way to return the site to its original state if it becomes corrupted or if it is being used for development and prototyping work and the user wants to start something over. The costs incurred for backup are the same as the costs of creating a site template.

Additionally, users can use Cloud Builder to delete a site. In that case, Cloud Builder will delete all the components that it has started.

It's important to note that the EBS drives are not deleted by default. Users should be sure to delete an EBS drive when no longer needed so it will not continue incurring costs.

Other features that Cloud Builder enables, like authorizing ArcGIS for Server or uploading a certificate for secure access, pose no extra costs to users when leveraged.

Estimating Costs

Instance Costs

For either of the example implementations, a cost per instance can be estimated. The hourly rate of the Standard Large Instance type is \$0.46 per hour. Assuming 720 hours per month, the estimated cost per month for one instance would be \$331.20. If the traffic to the site necessitated that the Auto Scaling feature maintain two instances at all times, then the cost per month would effectively double (i.e., \$662.40).

Region: US East (Virginia)		
	Linux/UNIX Usage	Windows Usage
Standard On-Demand Instances		
Small (Default)	\$0.080 per Hour	\$0.115 per Hour
Medium	\$0.160 per Hour	\$0.230 per Hour
Large	\$0.320 per Hour	\$0.460 per Hour
Extra Large	\$0.640 per Hour	\$0.920 per Hour
Micro On-Demand Instances		
Micro	\$0.020 per Hour	\$0.030 per Hour
High-Memory On-Demand Instances		
Extra Large	\$0.450 per Hour	\$0.570 per Hour
Double Extra Large	\$0.900 per Hour	\$1.140 per Hour
Quadruple Extra Large	\$1.800 per Hour	\$2.280 per Hour
High-CPU On-Demand Instances		
Medium	\$0.165 per Hour	\$0.285 per Hour
Extra Large	\$0.660 per Hour	\$1.140 per Hour

Please check for updated prices at aws.amazon.com/ec2/.

In some cases, additional instances are needed only during peak times. For example, a spike in traffic may occur due to an overwhelmingly popular public event, with the increase in traffic expected to last throughout that day, requiring an additional instance during that time. These situations can be scheduled and may help reduce costs.

If access to the application and/or services does not require high availability, it may not be necessary to plan for multiple instances. Additionally, an instance can be shut down during certain low-activity hours (e.g., at night or on the weekends and holidays).

Keep in mind that during normal operations, additional instances may need to be started for reasons other than web traffic or increased load on the server. For example, troubleshooting, testing, updates, or creating a new AMI may require adding more instances, at least for a temporary period. Planning some margin cost for such activities would be considered a best practice.

Amazon EC2 Reserved Instances are another way to save costs. If an instance will be used all the time for a full year, Reserved Instances can be requested and prepaid at a discounted rate. For example, given a Standard Large Reserved Linux Instance running full-time with heavy utilization for one year (8,760 hours) at \$0.064 per hour, the cost would be \$560.64, plus the \$780.00 one-time fee, for a total of \$1,340.64. This is a savings of \$6,608.16 from the On-Demand Instance equivalent scenario.

J-9979

Linux Windows				
Heavy Utilization Reserved Instances				
Region:	US East (Virginia)			
	1 yr Term		3 yr Term	
	Upfront	Hourly	Upfront	Hourly
Standard Reserved Instances				
Small (Default)	\$195	\$0.016 per Hour	\$300	\$0.013 per Hour
Medium	\$390	\$0.032 per Hour	\$600	\$0.026 per Hour
Large	\$780	\$0.064 per Hour	\$1200	\$0.052 per Hour
Extra Large	\$1560	\$0.128 per Hour	\$2400	\$0.104 per Hour
Micro Reserved Instances				
Micro	\$62	\$0.005 per Hour	\$100	\$0.005 per Hour
High-Memory Reserved Instances				
Extra Large	\$1030	\$0.088 per Hour	\$1550	\$0.07 per Hour
Double Extra Large	\$2060	\$0.176 per Hour	\$3100	\$0.14 per Hour
Quadruple Extra Large	\$4120	\$0.352 per Hour	\$6200	\$0.28 per Hour
High-CPU Reserved Instances				
Medium	\$500	\$0.04 per Hour	\$775	\$0.035 per Hour
Extra Large	\$2000	\$0.16 per Hour	\$3100	\$0.14 per Hour

Reserved Instances are also offered for medium and light utilization for cases when users expect to use the instance but not heavily all the time.

Please check for updated prices at aws.amazon.com/ec2/.

Storage and Data Transfer Costs

Input and output costs are measured by size in gigabytes. Input and output rates will go down with higher usage. This cost is one of the most difficult to estimate, since it will depend on each individual service and the way it's configured with the content. It will also depend on how many users make requests to the instances at a certain time. If the data size will be a terabyte or more (more gigabytes will still cost more, just less per gigabyte), users should do some testing and benchmarks to better estimate the costs for data transfer.

Region:	US East (Virginia)
Pricing	
Data Transfer IN	
All data transfer in	\$0.000 per GB
Data Transfer OUT	
First 1 GB / month	\$0.000 per GB
Up to 10 TB / month	\$0.120 per GB
Next 40 TB / month	\$0.090 per GB
Next 100 TB / month	\$0.070 per GB
Next 350 TB / month	\$0.050 per GB
Next 524 TB / month	Contact Us
Next 4 PB / month	Contact Us
Greater than 5 PB / month	Contact Us
<p>There is no Data Transfer charge between Amazon EC2 and other Amazon Web Services within the same region (i.e. between Amazon EC2 US West and Amazon S3 in US West). Data transferred between Amazon EC2 instances located in different Availability Zones in the same Region will be charged Regional Data Transfer. Data transferred between AWS services in different regions will be charged as Internet Data Transfer on both sides of the transfer.</p> <p>Usage for other Amazon Web Services is billed separately from Amazon EC2.</p>	
Availability Zone Data Transfer	
<ul style="list-style-type: none"> \$0.00 per GB – all data transferred between instances in the same Availability Zone using private IP addresses. 	
Regional Data Transfer	
<ul style="list-style-type: none"> \$0.01 per GB – all data transferred between instances in different Availability Zones in the same region. 	

Prices vary and change regularly. Please check current pricing at aws.amazon.com/ec2/.

Elastic Block Store Costs

EBS storage costs are calculated using the amount of input/output and the size of the storage in gigabytes. However, in this case, the input/output is calculated based on the number of requests, not their size. Since this is measured by the number of disk reads, it's usually good to be quite generous with the number of requests estimated.

Another good way to understand the impact of the costs is to use the site on a limited basis and track input/output use costs by examining the monthly bill. This can serve as a benchmark and help project what the bill might be with increased usage.

Most users try to limit the size of the EBS to just what is needed for supporting the solution. An EBS volume can be attached to only one instance. When a new instance is started, a new EBS volume is started as well.

Region:	US East (Virginia)
Amazon EBS Volumes	
<ul style="list-style-type: none"> \$0.10 per GB-month of provisioned storage \$0.10 per 1 million I/O requests 	
Amazon EBS Snapshots to Amazon S3	
<ul style="list-style-type: none"> \$0.125 per GB-month of data stored 	

Prices vary and change regularly. Please check current pricing at aws.amazon.com/ec2/.

J-9979

Snapshots are used as backups of the EBS drive. There's a different rate for storing snapshots. A snapshot can be used to create the second EBS volume rather than having to create both volumes from scratch.

Note: Users should delete their EBS drives to stop charges. Deleting the instance doesn't delete the EBS drive. Users should take special care to ensure that EBS drives are deleted when there is no need for them anymore.

Elastic IP Address Costs

Using an elastic IP address service does not include any extra costs unless it is not attached to an instance. The best practice is to keep an elastic IP address attached to an instance and in use.

No cost for Elastic IP addresses while in use

- \$0.01 per non-attached Elastic IP address per complete hour
- \$0.00 per Elastic IP address remap – first 100 remaps / month
- \$0.10 per Elastic IP address remap – additional remap / month over 100

Prices vary and change regularly. Please check updated prices at aws.amazon.com/ec2/.

Elastic Load Balancing Costs

Elastic Load Balancing costs are charged based on two components—the number of hours the ELB is running and the amount of data (per GB) that was processed through the ELB as input and output. In this case, it's safe to assume that the input and output of the ELB is equal to all the collective external input and output of the instances registered with it.

US – N. Virginia	US – N. California	EU – Ireland	APAC – Singapore
\$0.025 per Elastic Load Balancer-hour (or partial hour) \$0.008 per GB of data processed by an Elastic Load Balancer			

Prices vary and change regularly. Please check updated prices at aws.amazon.com/ec2/.

CloudWatch Costs

CloudWatch costs \$3.50 per instance per month regardless of the instance size when using detailed monitoring at a one-minute frequency. Basic monitoring at a five-minute frequency is free.

Region: <input type="text" value="US East (Virginia)"/>
Detailed Monitoring for Amazon EC2 Instances <ul style="list-style-type: none"> \$3.50 per instance per month, provided at 1-minute frequency
Basic Monitoring for Amazon EC2 instances <ul style="list-style-type: none"> \$0.00 (free of charge) per instance per month, provided at 5-minute frequency
Monitoring for Custom Metrics <ul style="list-style-type: none"> \$0.50 per metric per month

Prices vary and change regularly. Please check updated prices at aws.amazon.com/ec2/.

Amazon Calculators

Amazon provides a calculator to help users determine their potential monthly costs when using services; it can be found at calculator.s3.amazonaws.com/calc5.html.

FREE USAGE TIER: New Customers get free usage tier for first 12 months

Services Estimate of your Monthly Bill (\$ 895.47)

Choose region: US-East (Northern Virginia) & Inbound Data Transfer is Free and C

Amazon EC2

Amazon S3

Amazon RDS

Amazon DynamoDB

Amazon SimpleDB

Amazon SQS

Amazon SES

Amazon SNS

Amazon SWF

Amazon Route 53

Amazon CloudFront

Amazon ElastiCache

Amazon CloudWatch

Amazon VPC

Amazon Elastic MapReduce

AWS Import Export

AWS Premium Support

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in [the cloud](#). It is designed to make web-scale

Compute: Amazon EC2 On-Demand Instances:

Instances	Description	Operating System	Instance Type	Usage	Detailed Monitoring
1		Windows	Large	720 Hours/Month	<input checked="" type="checkbox"/>
1		Linux	Large	30 Hours/Month	<input checked="" type="checkbox"/>

Compute: Amazon EC2 Reserved Instances:

Storage: Amazon EBS Volumes:

Volumes	Description	Provisioned Storage	Average IOPS in volume	Snapshot Storage
2		100 GB-month	1000	20 GB-month of Storage

Elastic IP:

Number of Elastic IPs: 2

Elastic IP Non-attached Time: 10
Hours/Month

Number of Elastic IP Remaps: 20
Per Month

Amazon EC2 Data Transfer:

Data Transfer In: 30
GB/Month

Data Transfer Out: 60
GB/Month

Regional Data Transfer: 0
GB/Month

Public IP/Elastic IP Data Transfer: 60
GB/Month

Elastic Load Balancing:

Number of Elastic LBs: 0

Total Data Processed by all ELBs: 60
GB/Month

Another calculator is also available as a spreadsheet and can help compare costs between different options for infrastructure. It's available at [awsmedia.s3.amazonaws.com/Amazon EC2 Cost Comparison Calculator_022210.xls](http://awsmedia.s3.amazonaws.com/Amazon+EC2+Cost+Comparison+Calculator_022210.xls).

The spreadsheet contains many options to compare between using Amazon, a colocation facility, and traditional on-premises data center costs. Users can fine-tune the settings within the spreadsheet for each of those options depending on their work to model their annual total cost of ownership (TCO).

Amazon EC2 Bill Sample

The bill for using Amazon EC2 shows the costs of the instances based on use plus the cost of the EBS drives attached to those instances, as well as other AWS services used (e.g., Elastic Load Balancing, CloudWatch).

J-9979

Amazon Elastic Compute Cloud View/Edit Service			
US East (Northern Virginia) Region			
Amazon EC2 running Windows			
\$0.48 per Large Windows Instance (m1.large) instance-hour (or partial hour)	2,098 Hrs		1,007.04
Amazon EC2 EBS			
\$0.10 per GB-month of provisioned storage	1,047,867 GB-Mo		104.79
\$0.10 per 1 million I/O requests	11,394,333 IOs		1.14
\$0.15 per GB-Month of snapshot data stored	8,767 GB-Mo		1.32
\$0.01 per 10,000 gets (when loading a snapshot)	103,623 Requests		0.10
\$0.01 per 1,000 puts (when saving a snapshot)	2,488 Requests		0.02
Elastic Load Balancing			
\$0.008 per GB Data Processed by the LoadBalancer	4,048 GB		0.03
\$0.025 per LoadBalancer-hour (or partial hour)	1,057 Hrs		26.43
Amazon CloudWatch			
\$0.015 per monitored instance-hour (or partial hour)	1,171 Hrs		17.57
Download Usage Report »			1,158.43

Overall costs are represented as a combination of Amazon service reports.

Billing Statement: June 1, 2010

The billing cycle for this report is May 1 - May 31, 2010.

Expand All Services Collapse All Services		Printer Friendly Version
		Totals
Amazon Elastic Compute Cloud View/Edit Service		
Download Usage Report »		1,158.43
Amazon Simple Storage Service View/Edit Service		
Download Usage Report »		0.00
Amazon Virtual Private Cloud View/Edit Service		
Download Usage Report »		0.00
AWS Data Transfer (excluding Amazon CloudFront) View/Edit Service		
		0.63
Taxes		0.00
Total Charges due on June 1, 2010		\$1,159.06

Summary

Planning for the types of loads, usage, and features, along with the costs of each cloud service, provides a way to estimate the potential costs of cloud services.

For ArcGIS for Server on AWS, these principles apply:

- Begin with a Standard Large Instance type as a baseline for estimating instance costs.
- For solutions requiring high availability, plan for multiple instances and costs associated with Elastic Load Balancing and Amazon CloudWatch.

- Understand the extra components used by Cloud Builder and how they will affect costs.
- While an elastic IP address will not necessarily add to the cost of the solution (as long as it is attached to an instance), it makes instance replacement seamless and easy if an instance goes down or needs to be replaced. Without the elastic IP address, troubleshooting and setup costs could impact the project.
- Creating a custom AMI that reflects the solution setup can save time and money when it becomes necessary to re-create or update the setup.

It is good practice to compare estimates to the actual bill to see how accurate those estimates are. Continuous monitoring of the system and the bills is one of the best tools to keep costs under control.

For details on how to set up and use the ArcGIS Server AMIs or using Cloud Builder, please refer to the [ArcGIS for Server on AWS documentation](#).



Esri inspires and enables people to positively impact their future through a deeper, geographic understanding of the changing world around them.

Governments, industry leaders, academics, and nongovernmental organizations trust us to connect them with the analytic knowledge they need to make the critical decisions that shape the planet. For more than 40 years, Esri has cultivated collaborative relationships with partners who share our commitment to solving earth's most pressing challenges with geographic expertise and rational resolve. Today, we believe that geography is at the heart of a more resilient and sustainable future. Creating responsible products and solutions drives our passion for improving quality of life everywhere.



Contact Esri

380 New York Street
Redlands, California 92373-8100 USA

1 800 447 9778
T 909 793 2853
F 909 793 5953
info@esri.com
esri.com

Offices worldwide
esri.com/locations