



User Guide: Amazon EC2 Cost Comparison Calculator

Revised: November 8, 2011

See <http://aws.amazon.com/economics> for most current version.

Amazon Web Services (AWS) offers companies of all sizes an elastic, reliable, flexible, low-cost infrastructure web services platform in the cloud. Many companies have already launched applications in the cloud while others are currently evaluating the costs and benefits of moving some or all of their IT infrastructure to the cloud. To help financial decision makers quantify the direct economic benefits of cloud computing compared to traditional IT infrastructure alternatives, AWS has published the *Amazon EC2 Cost Comparison Calculator*, a basic Microsoft Excel-based, cost-comparison tool. While the Amazon EC2 Cost Comparison Calculator is designed to be self-explanatory, this user guide provides a detailed explanation of the calculator's inputs, assumptions, and calculations. Current versions of both this user guide and the Amazon EC2 Cost Comparison Calculator are available for free download at <http://aws.amazon.com/economics>.



Amazon EC2 Cost Comparison Calculator

The Amazon EC2 Cost Comparison Calculator (“Calculator”) uses industry data, AWS customer research, and user-defined inputs to quantify and compare the annual fully-burdened cost of owning, operating, and maintaining IT infrastructure versus the pay-for-use costs of using the Amazon Elastic Compute Cloud (“Amazon EC2”). The Calculator compares only the *direct costs* of IT infrastructure and ignores the many *indirect economic benefits* of cloud computing, including high availability, reliability, scalability, flexibility, faster time-to-market, and other cloud-oriented benefits. Decision makers are strongly encouraged to conduct a separate analysis to quantify and prioritize the economic value of these features.

In order to keep the analysis relatively simple and easy to understand, the Calculator quantifies only the most significant costs of owning and operating IT infrastructure and ignores many ancillary costs like architecture and engineering fees, cabling, real estate, IT build-out, facilities management, data center security, the cost of invested capital, and debt service. It is likely that owning and operating IT infrastructure will incur additional costs beyond those defined by the Calculator and explained in this user guide.

The Calculator is designed to be an important *first step* in helping a company quantify the direct economic benefits (or costs) of cloud computing. However, it should not replace a company’s careful analysis and consideration of its own IT infrastructure, architecture, business strategy, and cost structure. It is expected that most companies will want to customize the Calculator with their specific parameters and unique considerations. While all of the calculator’s built-in default assumptions are Amazon’s best estimates of average industry IT infrastructure costs, each assumption can (and should) be changed to represent the realities of each company’s use case.

1.0 Annual Cost of Amazon EC2

The Amazon Elastic Compute Cloud (“Amazon EC2”) is a web service that provides resizable compute capacity in the cloud on a pay-per-use basis. Amazon EC2 presents a true virtual computing environment, allowing companies to use web service interfaces to create machine images with a variety of operating systems, load them with custom application environments, manage network access permissions, and run images on as many or few systems (“Amazon EC2 instances”) as desired. Amazon EC2 instances can be added or removed instantly based on the variable demands of applications at any point in time.

The Calculator estimates all of the major costs of using Amazon EC2. These costs include: 1) instance hour costs and 2) data transfer costs. To simplify comparison with non-cloud alternatives, supplementary Amazon EC2 features like Elastic IPs, Amazon Elastic Block Storage (Amazon EBS), or Amazon Virtual Private Cloud (Amazon VPC) are excluded from this analysis.

1.1 Amazon EC2 Instance Cost

Amazon EC2 instances are billed according to actual hourly usage, rounded up to whole hours. The price per instance hour varies by Amazon EC2 Instance type, and geographic Region. Multiple Amazon EC2 Instance pricing options are available, including the option to: a) pay On-Demand Instance hourly usage rates, b) pay a one-time fee to reserve instances (for 1 or 3 year terms) and pay reduced hourly Reserved Instance hourly usage rates, or c) bid on unused Amazon EC2 capacity and qualify for hourly usage rates frequently lower than the On-Demand rate (during periods of lower demand).

Determining the number of instances needed and the percentage of time that each instance will be used is the first step in estimating the annual Amazon EC2 instance cost. On the “Main” tab of the Amazon EC2 Cost Comparison Calculator, users define the number of “baseline” and “peak” instances required to satisfy total application demand. Baseline



instances are assumed to be used for a large percentage of the year and represent the minimum base level instance demand. Peak instances are assumed to be used for a smaller percentage of the year and represent the *additional* instances needed to satisfy peak instance demand. In reality, for most applications the demand for instance hours constantly fluctuates between minimal demand (baseline instances only) and peak demand (requiring additional peak instances) due to normal variation, seasonal fluctuations, or unexpected spikes in demand.¹ Estimating the average annual usage of baseline and peak instances allows the total number Amazon EC2 instance hours to be calculated. The annual usage (hours) of baseline and peak Amazon EC2 instances is calculated by:

Hours of Baseline Instance Usage = No. of Baseline Instances * Hours per Year * Average Annual Usage of Baseline Instances

Hours of Peak Instance Usage = No. of Peak Instances * Hours per Year * Average Annual Usage of Peak Instances

Input	Description	Default Value	Data Source
No. of Baseline Instances	The minimum number of instances needed by the user's application(s) categorized by Amazon EC2 instance type (Standard Small, Standard Large, Standard Extra Large, High-Memory Extra Large, High-Memory Double Extra Large, High-Memory Quadruple Extra Large, High-CPU Medium, High-CPU Extra Large) Note: Calculator currently <i>excludes</i> Micro and Cluster GPU instances.	300 Standard Small Amazon EC2 Instances	User-defined.
Average Annual Usage of Baseline Instances	The average percentage of time (each year) that baseline instances are utilized	75%	User-defined.
No. of Peak Instances	The <i>additional</i> number of on-demand instances (beyond the defined number of baseline instances) needed to satisfy peak demand, categorized by Amazon EC2 instance type (see above)	700 Standard Small Amazon EC2 Instances	User-defined.
Average Annual Usage of Peak Instances	The average percentage of time (each year) that peak instances are utilized	10%	User-defined.
Hours per Year	The number of hours in a year	8,736	24 hours x 7 days x 52 weeks per year

After determining the annual usage of baseline and peak instances, the annual cost of Amazon EC2 instances (using On-Demand prices) is calculated by:

Annual Instance Cost (On-Demand Pricing) = (Hours of Baseline Instance Usage + Hours of Peak Instance Usage) * On-Demand Instance Price

¹ Berkeley researchers find that real world estimates of server utilization in traditional data centers range from 5% to 20% and that for many services, the peak workload exceeds the average by factors of 2 to 10. For more information, see "Above the Clouds: A Berkeley View of Cloud Computing," Technical Report No. UCB/EECS-2009-28, 2/10/2009, <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html>.

Amazon EC2 On-Demand Instance Prices

(in \$USD, price per hour)	US East (N. Virginia) Region US West (Oregon) Region		US West (N. California) Region		EU (Ireland) Region Asia Pacific (Singapore) Region		Asia Pacific (Tokyo) Region	
	Linux/Unix	Windows	Linux/Unix	Windows	Linux/Unix	Windows	Linux/Unix	Windows
Standard Instances								
Small (Default)	\$0.085	\$0.12	\$0.095	\$0.13	\$0.095	\$0.12	\$0.10	\$0.12
Large	\$0.34	\$0.48	\$0.38	\$0.52	\$0.38	\$0.48	\$0.40	\$0.48
Extra Large	\$0.68	\$0.96	\$0.76	\$1.04	\$0.76	\$0.96	\$0.80	\$0.96
High-Memory Instances								
Extra Large	\$0.50	\$0.62	\$0.57	\$0.69	\$0.57	\$0.62	\$0.60	\$0.62
Double Extra Large	\$1.00	\$1.24	\$1.14	\$1.38	\$1.14	\$1.24	\$1.20	\$1.24
Quadruple Extra Large	\$2.00	\$2.48	\$2.28	\$2.76	\$2.28	\$2.48	\$2.39	\$2.48
High-CPU Instances								
Medium	\$0.17	\$0.29	\$0.19	\$0.31	\$0.19	\$0.29	\$0.20	\$0.29
Large	\$0.68	\$1.16	\$0.76	\$1.24	\$0.76	\$1.16	\$0.80	\$1.16

Note: Hours of Baseline Usage and Hours of Peak Usage are calculated in this section, above.

Source: <http://aws.amazon.com/ec2>

In addition to calculating the annual On-Demand instance cost, the Amazon EC2 Cost Comparison Calculator also calculates the annual cost of 1 year and 3 year Reserved Instances. The Amazon EC2 Cost Comparison Calculator uses reserve instance pricing for all “baseline” instance usage and on-demand pricing for all “peak” instance usage. For baseline instances (running in the U.S. East Region, 24x7x365, or 100% annual utilization), a one-year reserved instance term reduces the annual cost of Amazon EC2 by 34% and a three-year reserved instance term reduces the annual cost by nearly 49%². The annual cost of Amazon EC2 (using 1 and 3 year reserved pricing) is calculated by:

Annual Instance Cost (w/ 1 Year Reserve Pricing) = Reserved Instance One-time Fees + Reserved Instance Usage + On-Demand Instance Usage, where:

Reserved Instance One-time Fees = No. of Baseline Instances * 1 Year Reserve One-Time Fee

Reserved Instance Usage = Hours of Baseline Instance Usage * Reserved Instance Usage Price

On-Demand Instance Usage = Hours of Peak Instance Usage * On-Demand EC2 Instance Price

Annual Instance Cost (w/ 3 Year Reserve Pricing) = Reserved Instance One-time Fees + Reserved Instance Usage + On-Demand Instance Usage, where:

Reserved Instance One-time Fees = (No. of Baseline Instances * 3 Year Reserve One-Time Fee) / 3 Years

² Steady state usage (8,736 hours per year) of an Amazon EC2 standard small Linux/Unix server in the US - Northern Virginia Region costs \$742.56 per year using On-Demand pricing, \$489.58 using 1-year Reserved instance pricing, and \$378.75 per year using 3-year Reserved Instance pricing. This represents a savings of 34% for a 1-year Reserved Instance and 49% savings for a 3-year Reserved Instance.

Note: For comparison purposes, the 3 Year Reserve One-Time Fee is depreciated over 3 years.

Reserved Instance Usage = *Hours of Baseline Usage* * *Reserved Instance Usage Price*

On-Demand Instance Usage = *Hours of Peak Instance Usage* * *On-Demand EC2 Instance Price*

Amazon EC2 Reserved Instance Prices (Linux/Unix instances only)

(in \$USD)	Pricing							
	US East (N. Virginia) Region US West (Oregon) Region				US West (N. California) Region			
	One-Time Fee		Hourly Usage Fee		One-Time Fee		Hourly Usage Fee	
Standard Instances	1 Year Term	3 Year Term	Linux/Unix	Windows	1 Year Term	3 Year Term	Linux/Unix	Windows
Small (Default)	\$227.50	\$350	\$0.03	\$0.05	\$227.50	\$350	\$0.04	\$0.06
Large	\$910	\$1,400	\$0.12	\$0.20	\$910	\$1,400	\$0.16	\$0.24
Extra Large	\$1,820	\$2,800	\$0.24	\$0.40	\$1,820	\$2,800	\$0.32	\$0.48
High-Memory Instances	1 Year Term	3 Year Term	Linux/Unix	Windows	1 Year Term	3 Year Term	Linux/Unix	Windows
Extra Large	\$1,325	\$2,000	\$0.17	\$0.24	\$1,325	\$2,000	\$0.24	\$0.32
Double Extra Large	\$2,650	\$4,000	\$0.34	\$0.48	\$2,650	\$4,000	\$0.48	\$0.64
Quadruple Extra Large	\$5,300	\$8,000	\$0.68	\$0.96	\$5,300	\$8,000	\$0.96	\$1.28
High-CPU Instances	1 Year Term	3 Year Term	Linux/Unix	Windows	1 Year Term	3 Year Term	Linux/Unix	Windows
Medium	\$455	\$700	\$0.06	\$0.125	\$455	\$700	\$0.08	\$0.145
Large	\$1,820	\$2,800	\$0.24	\$0.50	\$1,820	\$2,800	\$0.32	\$0.58
	EU (Ireland) Region & Asia Pacific (Tokyo) Region				Asia Pacific (Tokyo) Region			
	One-Time Fee		Hourly Usage Fee		One-Time Fee		Hourly Usage Fee	
	Standard Instances	1 Year Term	3 Year Term	Linux/Unix	Windows	1 Year Term	3 Year Term	Linux/Unix
Small (Default)	\$227.50	\$350	\$0.04	\$0.06	\$239	\$368	\$0.045	\$0.065
Large	\$910	\$1,400	\$0.16	\$0.24	\$956	\$1,470	\$0.18	\$0.26
Extra Large	\$1,820	\$2,800	\$0.32	\$0.48	\$1,911	\$2,940	\$0.36	\$0.52
High-Memory Instances	1 Year Term	3 Year Term	Linux/Unix	Windows	Linux/Unix	Windows	Linux/Unix	Windows
Extra Large	\$1,325	\$2,000	\$0.24	\$0.32	\$1,391	\$2,100	\$0.27	\$0.34
Double Extra Large	\$2,650	\$4,000	\$0.48	\$0.64	\$2,783	\$4,200	\$0.54	\$0.68
Quadruple Extra Large	\$5,300	\$8,000	\$0.96	\$1.28	\$5,565	\$8,400	\$1.07	\$1.36
High-CPU Instances	1 Year Term	3 Year Term	Linux/Unix	Windows	Linux/Unix	Windows	Linux/Unix	Windows
Medium	\$455	\$700	\$0.08	\$0.145	\$478	\$735	\$0.09	\$0.16
Large	\$1,820	\$2,800	\$0.32	\$0.58	\$1,911	\$2,940	\$0.36	\$0.64

Source: <http://aws.amazon.com/ec2>

Given the dynamic nature of Spot Instances prices, the Calculator does not explicitly attempt to quantify the annual cost of using Spot Instances. Spot Instances allow users to specify the maximum hourly price that they are willing to pay to run a particular instance type. There is a Spot Price for each instance type in each Region, which is the price all users will pay to run a Spot Instance, which fluctuates based on supply and demand. Customers will never pay more than the maximum price they have specified. If the Spot Price moves higher than a user's maximum price, the user's instance will

be shut down by Amazon EC2, with no further charges. Other than those differences, Spot Instances perform exactly the same as On-Demand or Reserved Instances.

Users can roughly estimate the annual cost of using Spot Instances by changing the default On-Demand Instance usage prices to reflect the average predicted Spot Price (assuming the user's bid price always exceeds the Spot Price). The annual cost of Amazon EC2 instances (using On-Demand prices) is calculated by:

Annual Instance Cost (Spot Pricing) = (Hours of Baseline Instance Usage + Hours of Peak Instance Usage) * Predicted Spot Price (variable)

Users are encouraged to go to <http://aws.amazon.com/ec2/spot-instances/> to learn more about Amazon EC2 Spot Instances.

1.2 Amazon EC2 Data Transfer

Data transfer cost for Amazon EC2 is based on amount of data transferred "in" and "out" of Amazon EC2 each month, excluding traffic coming from or going to other AWS services (e.g. Amazon S3), as there is no charge for data transfer within a Region of the AWS cloud. Users of the Cost Comparison Calculator will input the average monthly data transfer "in" and "out" (in GB) of each Amazon EC2 instance. The monthly data transfer (in GB) is calculated by:

Monthly Data Transfer "In" = ((No. of Baseline Instances * Average Annual Usage of Baseline Instances) + (No. of Peak Instances * Average Annual Usage of Peak Instances)) * Average Monthly Data Transfer "In" per Instance

Monthly Data Transfer "Out" = ((No. of Baseline Instances * Average Annual Usage of Baseline Instances) + (No. of Peak Instances * Average Annual Usage of Peak Instances)) * Average Monthly Data Transfer "Out" per Instance

Input	Description	Default Value	Data Source
Average Monthly Data Transfer "In" per Instance	The monthly data (in GB) transferred "in" to each Amazon EC2 instance via the internet; categorized by Amazon EC2 instance type (Standard Small, Standard Large, Standard Extra Large, High-Memory Extra Large, High-Memory Double Extra Large, High-Memory Quadruple Extra Large, High-CPU Medium, High-CPU Extra Large)	10	User-defined.
Average Monthly Data Transfer "Out" per Instance	The monthly data (in GB) transferred "out" to each Amazon EC2 instance via the internet; categorized by Amazon EC2 instance type (Standard Small, Standard Large, Standard Extra Large, High-Memory Extra Large, High-Memory Double Extra Large, High-Memory Quadruple Extra Large, High-CPU Medium, High-CPU Extra Large)	20	User-defined.

Note: The inputs for No. of Baseline Instances, No. of Peak Instances, and Average Annual Usage of Peak Instances are described in section 1.1, above.

The annual cost of data transfer is calculated by:

Annual EC2 Data Transfer Cost = ((*Monthly Data Transfer In* * *Data Transfer Price (In)*) + (*Monthly Data Transfer Out* * *Data Transfer Price (Out)*)) * *Months per Year*

Input	Description	Default Value	Data Source
Monthly Data Transfer In	The monthly data (in GB) transferred "in" to Amazon EC2 via the internet	Calculation	Calculated above
Monthly Data Transfer Out	The monthly data (in GB) transferred "out" of Amazon EC2 via the internet	Calculation	Calculated above
Data Transfer Price (In)	The cost per GB transferred "in" to Amazon EC2 via the internet	\$0.10	Published price on AWS website as of 11/1/2010.
Data Transfer Price (Out)	The cost per GB transferred "out" of Amazon EC2 via the internet	<p><i>All US and EU Regions</i> First GB = \$0.00 Up to 10 TB = \$0.12 Next 40 TB - \$0.09 Next 100 TB - \$0.07 Over 150 TB = \$0.05</p> <p><i>Asia Pacific (Singapore) Region:</i> First GB = \$0.00 First 10 TB = \$0.19 Next 40 TB - \$0.15 Next 100 TB - \$0.13 Over 150 TB = \$0.12</p> <p><i>Asia Pacific (Tokyo) Region:</i> First GB = \$0.00 First 10 TB = \$0.201 Next 40 TB - \$0.158 Next 100 TB - \$0.137 Over 150 TB = \$0.127</p>	<p>Published price on AWS website as of 9/1/2010.</p> <p>Tiered pricing is based on amount of data transferred each month</p>
Months per Year	The number of months per year	12	Definition.

2.0 Annual Cost of Co-Location

A traditional IT infrastructure alternative to Amazon EC2 is leasing co-located data center space. Co-location vendors offers data center space, power, cooling, and connectivity to support user-owned server and network hardware.

The annual cost of co-location is based on the total cost of purchasing, operating, and supporting server equipment co-located in a third-party's data center. The Calculator helps users quantify the major costs of co-located IT infrastructure. These costs include 1) server hardware 2) network hardware, 3) hardware maintenance, 4) server operating system, 5) co-location fees, 6) remote hands support, and 7) data transfer. Default assumptions are based on customer interviews and industry research. Users can change assumptions and prices based on their own research or using current quotes from their hardware vendors and co-location providers.

2.1 Server Hardware Cost

In section 1.1 above, users defined the number of “baseline” and “peak” EC2 instances needed to meet maximum application demand. The sum of these “baseline” and “peak” instances equals the number of servers that must be purchased by the user in the co-location scenario. For comparison purposes, the Calculator requires that servers be near-equivalent in performance to Amazon EC2 instances types³.

Like other useful assets, the initial purchase (or acquisition) price of server hardware is depreciated over the asset’s expected useful life. Using straight-line depreciation, the annual server expense is the initial purchase price minus the end-of-life salvage value, divided by the expected useful life. According to discussions with partners and customers, most companies estimate a 3 year useful life for server hardware with minimal end-of-life salvage value. The Calculator assumes zero salvage value, as a simplifying assumption. Thus, the annual cost of servers is calculated by:

$$\text{Annual Cost of Server Hardware} = \text{Total Number of Servers} * \text{Cost per Server} / \text{Server Useful Life}$$

Input	Description	Default Value	Data Source
Total Number of Servers	Total number of servers, categorized by equivalent Amazon EC2 Instance Type (Standard Small, Standard Large, Standard Extra Large, High-Memory Extra Large, High-Memory Double Extra Large, High-Memory Quadruple Extra Large, High-CPU Medium, High-CPU Extra Large) This is the sum of user-defined “baseline” and “peak” servers defined in section 1.1 above.	1,000 (300 + 700) Amazon EC2-equivalent servers (Standard Small-Equivalent instance type)	User-defined. The default value(s) are arbitrary. Available server types are designed to mirror the compute capacity of available EC2 instance types.
Cost per Server	Cost of servers, categorized by Amazon EC2 Instance type (see above)	<u>Standard Compute</u> Small = \$363 Large = \$1,452 X-Large = \$2,903 <u>High-Memory</u> X-Large = \$3,237 Double X-Large = \$5,244 Quadruple X-Large = \$10,487 <u>High Compute</u> Medium = \$682 X-Large = \$2,726	Market prices based on guidance from major hardware vendors and assuming virtualization of smaller instance types. Users can edit based on their own research.
Server Useful Life	Estimated useful life (in years) of servers	3 years	Amazon customer

³ To perform a cost comparison with Amazon EC2, this methodology requires that servers be comparable to available Amazon EC2 instances. One Amazon EC2 Compute Unit provides the equivalent CPU capacity of a 1.0-1.2 GHz 2007 Opteron or 2007 Xeon processor. See <http://aws.amazon.com/ec2> for more information on Amazon EC2 instance types.

			surveys reveal average estimated useful life of servers between 3-5 years. Most use 3 years.
--	--	--	--

2.2 Network Hardware

Network hardware allows servers to connect to one another and to the Internet. Network hardware includes firewalls, routers, switches, intrusion detection systems, and other equipment. Since network architecture is very dependent on each application's use case, a detailed calculation of network hardware cost is beyond the scope of this basic analysis. Instead, the Calculator estimates the cost of network hardware as a percentage of the initial server hardware cost. Discussions with customers have estimated this value to be between 10% and 30%. Users of this calculator are encouraged to conduct a separate analysis of their network demands to confirm initial estimates used in this model.

The Calculator depreciates network hardware using straight-line depreciation over the expected useful life of the equipment. The annual cost of network hardware is calculated by:

Annual Cost of Network Hardware = $(\text{Total Number of Servers} * \text{Cost per Server} * \text{Network Hardware Cost}) / \text{Network Hardware Useful Life}$

Input	Description	Default Value	Data Source
Network Hardware Cost	The network hardware cost, expressed as a percentage of the initial purchase (or acquisition) price of server hardware	20%	AWS customer surveys reveal network cost as 10-30% of initial hardware cost
Network Hardware Useful Life	Estimated useful life (in years) of network hardware	3 years	AWS customer surveys reveal average estimated useful life of servers between 3-5 years.

Note: The inputs for *Total Number of Servers* and *Cost per Server* are described in section 2.1, above.

2.3 Hardware Maintenance Cost

Server and network hardware is commonly purchased with annual maintenance contracts to support the reliability of the hardware devices over time. Users input the annual maintenance cost as a percentage of the initial purchase cost of server and network hardware. Discussions with customers and hardware vendors have estimated this percentage to be approximately 10%. The annual cost of maintenance for server and network hardware is calculated by:

Annual Hardware Maintenance Cost = $\text{Annual Maintenance Cost of Servers} + \text{Annual Maintenance Cost of Network Hardware}$, **where:**

Annual Maintenance Cost of Servers = $\text{Total Number of Servers} * \text{Cost per Server} * \text{Hardware Maintenance Rate}$

Annual Maintenance Cost of Network Hardware = $\text{Total Number of Servers} * \text{Cost per Server} * \text{Network Hardware Cost} * \text{Hardware Maintenance Rate}$



Input	Description	Default Value	Data Source
Hardware Maintenance Rate	Annual maintenance rate for server and network hardware as a percentage of initial purchase price	10%	AWS customer surveys reveal average annual hardware maintenance cost as 10% of initial hardware costs.

Note: The inputs for *Total Number of Servers* and *Cost per Server* are described in section 2.1, above. The input for *Network Hardware Cost* is described in section 2.2, above.

2.4 Server Operating System Cost

Servers require operating systems to function and host client applications. Server operating systems can be open-source or licensed directly from vendors. The complex nature of operating system licensing schemes is beyond the scope of this basic analysis. Users estimate the initial cost of operating system software (either Linux/Unix or Windows) on a per server basis. The initial cost is then amortized over the user-defined expected useful life of servers. The annual server operating system cost is calculated by:

Annual Server Operating System Cost = $(\text{Total Number of Servers} * \text{OS Cost per Server}) / \text{Server Useful Life}$

Input	Description	Default Value	Data Source
OS Cost per Server	Operating system cost per server	\$0 (Linux/Unix) \$2,999 (Windows)	Default Windows value is based on the U.S. cost of Windows Server 2008 Data Center Edition as defined on Microsoft's Web site as of November 2009. ⁴

Note: The inputs for *Total Number of Servers* and *Server Useful Life* are described in section 2.1, above.

2.5 Co-location Fees

In most co-located data center, servers are housed in industry standard 19" wide racks with each individual rack unit (U) measuring roughly 1.75" in height. While a basic server may occupy only 1U of rack space, more advanced models may occupy 2U or more of rack space. Space in a co-located data center can be leased in varying sizes, from individual single rack units, to complete racks. Users of co-located data centers pay an initial, non-recurring set-up cost (i.e. to pay for rack bolting, installation of power whips and breakers, etc.) and an ongoing recurring monthly co-location fee for each unit of space used (to pay for power, cooling, space, etc.).

Typically, a co-location provider will allot a certain amount of power to correspond to the amount of space that is leased in the facility. To accurately calculate the co-location cost of server hardware, the Calculator estimates the physical space required by servers as well as the power allotment required by each server. If the power required by servers

⁴ Microsoft Website, <http://www.microsoft.com/windowsserver2008/en/us/pricing.aspx>, accessed 11/30/2009

exceeds the allotment of power to each unit of space, then the Calculator assumes that additional units of space are required. The amount of co-located data center space *in terms of power* is calculated by:

Number of Rack Units Needed (Power) = $ROUNDUP(\text{Total Number of Servers} * \text{Nameplate Power per Server} / \text{Power per Rack Unit})$

The amount of co-located data center space *in terms of space* is calculated by:

Number of Rack Units Needed (Space) = $\text{Total Number of Servers} * \text{Server Space}$

The Calculator then assumes that the co-located space needed is the *maximum* of the two calculations above:

Number of Rack Units Needed (Power/Space) = $MAX(\text{Number of Rack Units Needed (Power)}, \text{Number of Rack Units Needed (Space)})$

Input	Description	Default Value	Data Source
Nameplate Power per Server	The amount of power (watts) used by servers, categorized by Amazon EC2-equivalent server type	<u>Standard Compute</u> Small = 150W Large = 240W X-Large = 391W <u>High-Memory</u> X-Large = 405W Double X-Large = 515W Quad X-Large = 630W <u>High Compute</u> Medium = 289W X-Large = 309W	Nameplate ratings of Amazon EC2 equivalent servers based on guidance from major hardware vendors.
Server Space	The rack unit space occupied by Amazon EC2 equivalent servers, categorized by Amazon EC2-equivalent server type	<u>Standard Compute</u> Small = 1U Large = 2U X-Large = 2U <u>High-Memory</u> X-Large = 2U Double X-Large = 2U Quad X-Large = 2U <u>High-Compute</u> Medium = 1U X-Large = 2U	Based on guidance from major hardware vendors (comparing Amazon EC2-equivalent servers).
Power per Rack Unit	The allotment of power (watts) provided to each unit of rack space as determined by the co-location provider	119 watts per rack unit	Assumes 5,000 kW per rack divided by 42 rack units per rack. (5,000/42 = 119)
Months per Year	The number of months in a calendar year.	12	Defined.

Note: The inputs for *Total Number of Servers* is described in section 2.1, above.

The functions “ROUNDUP” and “MAX” are spreadsheet functions in Microsoft Excel. The “ROUNDUP” function rounds a number up, away from zero where ROUNDUP(number,num_digits) and is used to round the number of rack units up to the nearest whole number. The “MAX” function returns the largest value in a set of values, where MAX(number1, number 2,) and is used to determine the “max” number of rack units needed based on power *or* space.

Once the number of rack units needed is calculated, the Calculator uses user-defined co-location prices to determine the annual cost of co-location. For simplicity and comparison purposes, the initial non-recurring co-location setup fee is depreciated over the useful life of the server hardware. The annual cost of co-location is calculated by:

Annual Cost of Co-location = *Non-Recurring Set-up Cost* + *Annual Recurring Cost*, **where:**

Non-Recurring Set-up Cost = *Number of Rack Units Needed* * *Initial Set-up Fee* / *Server Useful Life*

Annual Recurring Cost = *Number of Rack Units Needed* * *On-going Co-lo Fee* * *Months per Year*

Input	Description	Default Value	Data Source
Initial Co-Lo Set-up Fee	The initial set-up fee charged by co-location providers to install equipment in each unit of space	\$80 per rack unit (U)	Based on a survey of co-location providers. Users should consult local co-location providers.
On-going Co-Lo Fee	The monthly fee charged by co-location providers per unit of space	\$80 per rack unit (U)	Based on a survey of co-location providers. Users should consult local co-location providers.
Months per Year	The number of months in a calendar year.	12	Defined.

Note: The *Number of Rack Units Needed* is based on the “max” of power or space demand, described above in this section. Also, users of the Calculator should make sure the cost per rack unit (U) corresponds with the *Power per Rack Unit* defined above.

2.6 Co-Location Remote Hands

Periodically, remote hands will be required to perform basic maintenance and troubleshooting on physical servers located in a co-location provider’s data center. Co-location providers offer remote hands services for reconfiguring servers, installing software, performing reboots, etc. The annual cost of remote hands support is calculated by:

Annual Cost of Remote Hands Support = *Total Number of Servers* * *Server Failure Rate* * *Hours per Failure* * *Remote Hands Cost per Hour*

Input	Description	Default Value	Data Source
Server Failure Rate	The annual failure rate of servers, including episodes requiring configuration, reboot, troubleshooting, etc.	10%	AWS estimate.
Hours per Failure	The average number of hours required to respond to server failures	1 hour	AWS estimate.
Remote Hands Cost per Hour	The hourly rate for remote hands support in a co-located data center	\$150	Based on a survey of co-location provider’s Web sites.

Note: The input for the *Total Number of Servers* is described in section 2.1, above.

2.7 Data Transfer

To allow applications to transmit data over the Internet, companies pay for transit capacity from co-location providers. Typically, co-location providers charge customers for “billable utilization of bandwidth” using the 95% percentile bandwidth provisioning method. Every five minutes the amount of data transferred through the provider’s switch or router is measured and recorded in a log file. At the end of the month, the samples are sorted from highest to lowest, the top 5% of data (approximately 36 hours of a 30-day billing cycle) is discarded, and the next highest measurement becomes the 'monthly billable utilization' for the month. The billable utilization (measured in Mbps) is then multiplied by a monthly price per megabit transferred to determine the monthly data transfer cost.

The Calculator simplifies the calculation of data transfer cost by only estimating the peak billable utilization. First the user-defined total monthly data transfer (in GB) is converted to the average monthly data transfer (in Mbps). Then the average monthly data transfer is multiplied by a peak-to-average ratio (i.e. 5x) to determine the monthly billable utilization. Finally, the monthly billable utilization is multiplied by the co-location provider’s price per megabit (i.e. \$75 per Mbps) to determine the monthly data transfer cost. The billable monthly bandwidth cost is calculated by:

Monthly Billable Utilization = *Total Monthly Data Transfer * Conversion to Megabits / Seconds per Month * Bandwidth Peak to Average Factor*

After calculating the monthly billable utilization, the annual data transfer cost is calculated by:

Annual Data Transfer Cost = *Billable Monthly Bandwidth * Price per Mbps per Month * Months per Year*

Input	Description	Default Value	Data Source
Total Monthly Data Transfer	Monthly data transfer (GB) per month; (sum of data transfer “in” and “out” of servers)	Calculation	Sum of Monthly Data Transfer ‘In’ and Monthly Data Transfer ‘Out’ from section 1.2. Calculation based on user-inputs.
Conversion to Megabits	Constant used to determine the number of megabits per GB	8,192	1,024 megabytes per gigabyte, and 8 bits per byte. (1,024 x 8)
Seconds per Month	Number of second in an average month	$2,635,200 = 60sec * 60min * 24hours * 30.5days/month$	Definition.
Bandwidth Peak to Average Factor	Conversion factor used to determine the peak actual data transfer per second (as a multiple of average data transfer)	5	AWS estimate. Can be edited by the user. This value is highly dependent on the variability of the user’s application(s).
Price per Mbps per Month	Price for data transfer based on peak data transfer each month.	User input, defaults are: 0-1 Mbps = \$100 per Megabit 1-5 Mbps = \$80 per Megabit 5-10 Mbps = \$75 per Megabit 10-30 Mbps = \$70 per Megabit 30-50 Mbps = \$60 per Megabit 50-100 Mbps = \$50 per Megabit 100-150 Mbps = \$40 per Megabit > 150 Mbps = \$30 per Megabit	Defined by user; defaults are based on a survey of co-lo bandwidth rates on provider’s Web sites. Users should consult with local co-lo providers for specific pricing.
Months per Year	Months in a year	12	Definition.

3.0 Annual Cost of On-Site Infrastructure

Another alternative to using Amazon EC2 is to own and operate an independent on-site corporate data center. Companies that run data centers assume all of the fully-burdened costs of IT infrastructure.

The Calculator helps users quantify the major costs of on-site IT infrastructure. These costs include 1) server hardware 2) network hardware, 3) hardware maintenance, 4) server operating system, 5) power and cooling, 6) data center construction 7) administration, and 8) data transfer. This analysis excludes cabling, security, interest, real estate, property taxes, landscaping, and other expenses associated with owning data centers. Default assumptions are based on AWS customer interviews and industry research; users are encouraged to change assumptions and prices based on their own research.

3.1 Server Hardware

Same as section 2.1, above.

3.2 Network Hardware

Same as section 2.2, above.

3.3 Hardware Maintenance

Same as section 2.3, above.

3.4 Server Operating System

Same as section 2.4, above.

3.5 Server Operating Power and Server Cooling Power

Data center servers not only consume electricity but they also convert that electricity into heat, which must be removed from the data center to avoid overheating server equipment. The power required to cool data center servers can equal or exceed the power used by the servers themselves. The Green Grid consortium estimates that most data centers have a power usage effectiveness (PUE) ratio of between 1.3 and 3.0.⁵ A PUE of 3.0 means for every 3.0 watts of electricity entering the facility, only 1.0 watts is used to support computing activities (i.e. servers) while the other 2.0 watts support other infrastructure (i.e. cooling). The Amazon EC2 Cost Comparison Calculator calculates the total power usage (power and cooling) by estimating the (de-rated) energy demand of servers multiplied by user-defined PUE ratio. The total electricity usage is then multiplied by the average cost of electricity (per kW hour) to estimate the total annual power costs. The annual cost for server operating power and cooling is calculated by:

Annual Cost of Power and Cooling = *Total Number of Servers * Nameplate Power per Server (W) * Power Conversion Factor * Server Operating Hours per Year * Average PUE * (Electricity Price per Hour (kW)/1000)*

Input	Description	Default Value	Data Source
Nameplate Power per Server (W)	The maximum amount of power (watts) used by servers, categorized by Amazon EC2-	<u>Standard Compute</u> Small = 150W	Nameplate ratings of Amazon EC2 equivalent

⁵ Whitepaper from The Green Grid, "Green Grid Data Center Power Efficiency Metrics: PUE and DCIE", 2008.

	equivalent server type	Large = 240W X-Large = 391W <u>High-Memory</u> Double X-Large = 515W Quad X-Large = 630W <u>High Compute</u> Medium = 289W X-Large = 309W	servers based on guidance from major hardware vendors.
Power Conversion Factor	The conversion factor used to convert the nameplate power rating (maximum) to average operating load	0.50	Nameplate ratings are typically de-rated by 50% to 75%, according Raritan. ⁶ Users can edit based on own estimates of server load and power utilization.
Server Operating Hours per Year	The number of hours each year the server is turned on	8,736	Annual operating hours for a typical 24x7x52 operation
Average PUE (Power Usage Effectiveness)	The "power usage effectiveness" or PUE ratio is a metric used to determine the energy efficiency of a data center. A PUE of 3.0 means for every 3.0 watts of electricity entering the facility, only 1.0 watts is used to support computing activities (i.e. servers) while the other 2.0 watts support other infrastructure (i.e. cooling).	2.5	The Green Grid consortium estimates an industry-average PUE range of 1.3 – 3.0. Users can adjust based on their own internal metrics.
Electricity Price per Hour	Price per hour per kW of electricity	\$0.09	Average commercial value for the United States in December 2008, Energy Information Administration Electricity Statistics.

Note: The input for *Total Number of Servers* is described in section 2.1, above.

Methodology only calculates power costs associated with server hardware. Electricity used to power and cool network equipment (switches and routers) is excluded from this analysis.

3.6 Data Center Construction Cost

Servers need to be protected, powered, cooled, and monitored in a data center environment. Modern data centers are built with raised floors, fire detection/pre-action sprinklers, UPS power backup systems, utility switch-gear, generators,

⁶ Whitepaper from Raritan.com, "Data Center Power Distribution and Capacity Planning," 2009.

HVAC systems, and many other industrial components. The Uptime Institute has published a simple model to quickly forecast the construction costs of modern data centers.⁷ The Uptime Institute’s model emphasizes the “engine” capacity and functionality aspects of a data centers, in addition to the space itself. The “engine” capacity refers to the power and cooling *engine* that supports the computer room space. This engine is sized in kW by Tier level and provides for the installation of the underlying and redundant mechanical/electrical infrastructure capacity (measured in kW of uninterruptible power and cooling required by the computer and cooling load), as well as the distribution equipment required to support the desired Tier of functionality. A description of the Uptime Institute’s functionality Tiers and the “kW cost” component of each Tier is defined below:

Tier I	Tier I is composed of a single path for power and cooling distribution, without redundant components, providing 99.671% availability. \$11,500 per kW of redundant UPS capacity for IT
Tier II	Tier II is composed of a single path for power and cooling distribution, with redundant components, providing 99.741% availability. \$12,500 per kW of redundant UPS capacity for IT
Tier III	Tier III is composed of multiple active power and cooling distribution paths, but only one path active, has redundant components, and is concurrently maintainable, providing 99.982% availability. \$23,000 per kW of redundant UPS capacity for IT
Tier IV	Tier IV is composed of multiple active power and cooling distribution paths, has redundant components, and is fault tolerant, providing 99.995% availability. \$25,000 per kW of redundant UPS capacity for IT

Source: *Uptime Institute*

In addition to determining the “kW component” of data center construction costs, the Uptime Institute recommends that in all cases \$300 per ft² of computer room floor be added to the “kW component.” Thus, a Tier III data center is estimated to cost \$23,000 per kW of redundant IT power to construct, plus \$300 per ft² of computer room space. (Note: While it may be useful to add an additional “empty space” cost component to support future growth, it is excluded from the Calculator for purposes of simplicity and conservatism.)

To calculate the annual cost of an on-site corporate data center, the Amazon EC2 Cost Comparison Calculator first multiplies the expected power demand of servers (in kW) by the per kW construction cost estimate. It then determines the space needed for the computer equipment and multiplies that by \$300 per ft², or another user-defined value. Finally, the total construction cost is depreciated over the average useful life of the data center.

Annual Data Center kW Cost = (Total Servers Needed * Nameplate Power per Server (W) * Power Conversion Factor / 1,000 * Data Center Cost per kW) / Data Center Useful Life

⁷ Whitepaper from the Uptime Institute, “Cost Model: Dollars per kW plus Dollars per Square Foot of Computer Floor,” 2008. Related Uptime Institute Whitepapers include: “Tier Classifications Define Infrastructure Performance,” 2001-2005. See uptimeinstitute.org/whitepapers for more information.

Annual Data Center Computer Room Cost = $\text{Rack Units Needed} / \text{Rack Units per Rack} * \text{Space per Rack} * \text{Computer Room Cost} / \text{Data Center Useful Life}$

Input	Description	Default Value	Data Source
Data Center Cost per kW	The construction cost per kW of electricity demanded to build a data center with power, cooling, and auxiliary support systems.	Default value = \$23,000, assuming Tier III data center space (99.982% site availability) Tier I: \$11,500/kW Tier II: \$12,500/kW Tier III: \$23,000/kW Tier IV: \$25,000/kW	Based on 2007 estimates published by the Uptime Institute.
Rack Units Needed	The number of rack units needed based on space or power demand, whichever is greater.	This is calculated based on user inputs, as explained in section 2.5 above.	Calculation
Rack Units per Rack	The typical number of rack units in a standard 19" rack.	42	Survey of vendor web sites.
Space per Rack	The amount of computer room space occupied by a rack included surrounding buffer space.	22 square feet	An industry standard 19" wide rack occupies approximately 7 square feet of space, but requires approximately 15 additional square feet of space for walkways, cooling equipment, etc. based on Amazon estimates.
Computer Room Cost	The cost per square ft to build a computer room in a data center.	\$300 per square ft	Based on 2007 estimates published by the Uptime Institute.
Data Center Useful Life	The expected useful life (in years) of a newly constructed data center	15	AWS estimate.

Note: The input for *Total Servers Needed* is described in section 2.1, above. The input for *Nameplate Power per Server* and *Power Conversion Factor* are from section 3.5, above.

3.7 Administration Cost

When companies own and operate servers on-site, they need to hire full-time personnel to procure, deploy, support, and manage server assets. Server administrators perform a variety of duties, including: server upgrades, restores, patches, or moves; problem/incident management; vendor management; and monitoring. The annual server administration cost is calculated by:

Administration Cost = $ROUNDUP(\text{Total Servers Needed} / \text{Average Servers per FTE Administrator} * \text{Average Annual Salary per FTE Administrator})$

Input	Description	Default Value	Data Source
Average Servers per FTE Administrator	Number of servers that one FTE server administrator can support	50	Default value is an estimate for small-scale on-site deployments. Larger scale data center operations may allow for greater efficiency.
Average Annual Salary per FTE Administrator	Average annual salary of a FTE Administrator	\$105,000	Default value based on Amazon's estimates of average total compensation for a System Admin/Engineer

Note: The input for *Total Servers Needed* is described in section 2.1, above.

The functions "ROUNDUP" is a spreadsheet functions in Microsoft Excel. The "ROUNDUP" function rounds a number up, away from zero where $ROUNDUP(\text{number}, \text{num_digits})$ and is used to round the number of rack units up to the nearest whole number.

3.8 Data Transfer

Companies operating on-site infrastructure must lease dedicated lines from telecom providers to connect their equipment to the Internet backbone. The Calculator uses a similar methodology as was used in section 2.7 above to determine the annual cost of data transfer from on-site data centers. First the user-defined total monthly data transfer (in GB) is converted to the billable monthly bandwidth (in Mbps). Then, the billable monthly bandwidth is multiplied by a price per megabit (i.e. \$25 per Mbps) to determine the monthly data transfer cost. The annual data transfer cost is calculated by:

Billable Monthly Bandwidth = $\text{Bandwidth Used} * \text{Conversion to Megabits} / \text{Seconds per Month}$

Annual Bandwidth Cost = $\text{Billable Monthly Bandwidth} * \text{Price per Mbps per Month} * \text{Months per Year}$

Input	Description	Default Value	Data Source
Bandwidth Used (GB) (per month)	Bandwidth used (GB) per month (sum of download and upload bandwidth)	User input	User
Conversion to Megabits	Constant used to determine the number of megabits per GB	8,192	1,024 megabytes per gigabyte, and 8 bits per byte. (1,024 x 8)
Seconds per Month	Number of second in an average month	2,635,200 =	Definition.



		60sec * 60min * 24hours * 30.5days/month	
<i>Price per Mbps per Month</i>	Price for bandwidth per Mbps based on total transferred	\$25.00	Defined by user; defaults are based on a survey of AWS customers in the US.
<i>Months per Year</i>	Months in a year	12	Definition.

5.0 Conclusion

The Amazon EC2 Cost Comparison Calculator is designed to help compare the annual TCO of Amazon EC2 with traditional IT infrastructure alternatives. Amazon has provided this tool in Microsoft Excel spreadsheet format so that financial decision makers can use it as a first step in identifying, quantifying, and evaluating the cost of IT infrastructure and comparing those costs with using Amazon EC2. It is expected that most customers will want to customize the tool with their specific parameters and unique considerations. The Amazon EC2 Cost Comparison Calculator should not be used as a substitute for careful analysis of specific user scenarios that may differ from a generalized model.

Revision History

The following updates were made to the original December 7, 2009 version of this document.

February 1, 2010

- Explanation of Spot Instance pricing.
- Amazon EC2 data transfer “in” price reduced to \$0.00 per GB. Note that data transfer “in” price is currently scheduled to return to \$0.10 per GB after 6/30/2010.
- Amazon EC2 data transfer “out” prices reduced permanently by \$0.02, across all usage tiers, effective 2/1/2010.

February 22, 2010

- Added additional Amazon EC2 instance type: High-Memory Extra Large Instance.
- Added Windows Reserved Instances.

April 28, 2010

- Added the Asia Pacific (Singapore) Region pricing.

September 1, 2010

- Updated pricing for High-Memory Double Extra Large and High-Memory Quadruple Extra Large instance types

November 1, 2010

- Amazon EC2 data transfer “in” price returned to \$0.10 per GB. (See revision from February 1, 2010).

March 2, 2011

- Added the Asia Pacific (Tokyo) Region pricing.

July 1, 2011

- Updated Amazon EC2 inbound and outbound data transfer pricing for US and EU Regions.

November 8, 2011

- Added the US West (Oregon) Region pricing.

Notices

© 2009-2011 Amazon.com, Inc., or its affiliates.

Original Publish Date: December 7, 2009; Last Revision Date: November 8, 2011

This user guide and accompanying workbook is provided for informational purposes only. Amazon Web Services LLC is not responsible for any damages related to the information in these documents, which are provided “as is” without warranty of any kind, whether express, implied, or statutory. Nothing in this user guide or accompanying workbook creates any warranties or representations from Amazon Web Services LLC, its affiliates, suppliers, or licensors. This user guide and accompanying workbook do not modify the applicable terms and conditions governing your use of Amazon Web Services technologies, including the Amazon Web Services website. This workbook and accompanying user guide represents Amazon Web Services' current product offerings as of the date of issue and are subject to change without notice.