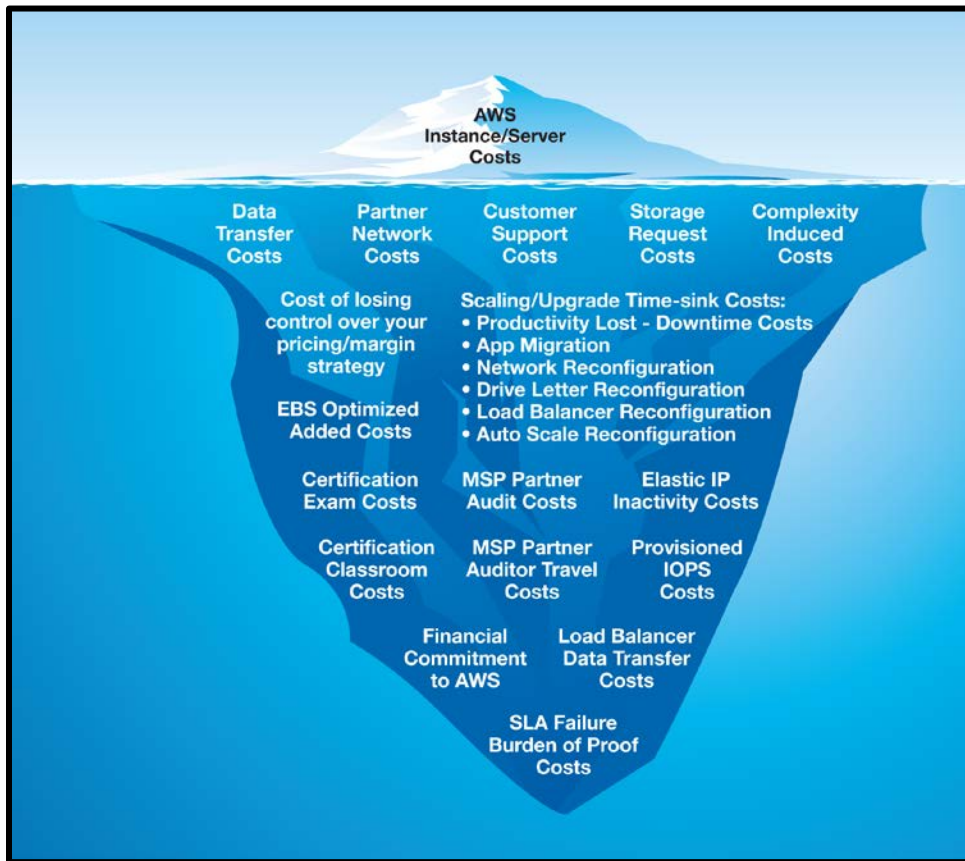


The Total Cost of Doing Business with Amazon Web Services™

Looking below the surface to discover the real costs, and threats of using AWS™



This is the final paper in a series of five providing information on the total cost of doing business (TCODB) with Amazon Web Services (AWS). The information in this paper includes all the information from the first four papers.

There are plenty of cloud vendor cost comparison websites and tools on the Internet, but they typically only compare infrastructure costs. While infrastructure cost comparisons are valuable, it is unfortunate that these websites don't dig into the hidden costs that so many large (commodity) cloud service providers (CSPs) impose on their unsuspecting service provider partners.

Finding and analyzing the hidden costs of a large CSP is a daunting task. The series of papers specifically looked at various aspects of doing business with AWS. We turned over many rocks to compile the data in these papers, but AWS is a huge, vastly complex business unit with a labyrinth of interconnected, costly, and dependent services. There's a possibility that we missed a few rocks and there are additional hidden costs and complexities you may, or may not encounter with AWS.

Before signing on the dotted line, it behooves all technology service providers, ISVs, resellers, and distributors to be aware of the potential hidden costs that add up to the TCODB with any CSP.

Many technology providers have told us that part of the challenge of doing business with AWS is that there is just too much complexity. We created a **CSP Assessment Checklist**¹ to help technology service providers know what to look for when they start, or need to realign their cloud journey with a CSP.

To comprehend the complexity highlighted in this paper, it's important to understand AWS's roots. AWS was born when Amazon engineering developed a massive infrastructure to support its burgeoning online consumer business. While it was paramount that the consumer facing front-end was user friendly, the back-end infrastructure was designed for use by Amazon developers.

Amazon realized that their infrastructure contained capacity that was unused and decided to make it available in the form of Amazon Web Services. In particular, the compute component is referred to as the Amazon Elastic Compute Cloud™, or simply EC2™.

As mentioned, the infrastructure back end was designed for Amazon developers, not end users and not partners. In short, AWS is very complex because it was designed for developers.

As a testament to this, searching for the term “developer” in the EC2 FAQ² turns up 27 occurrences including:

- *Q: What is Amazon Elastic Compute Cloud™ (Amazon EC2)?*
 - *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 computing easier for developers.***

This complexity shows up in many facets of AWS, and as is often the case with technology, complexity typically leads to **additional costs, many of which are not immediately apparent.**

As research was being performed for the series of papers, there was one consistent thought that was validated with every rock that was turned over: “this is no way to treat partners (or indirectly), their clients.”

And then there's the recent move by Amazon to enter the managed services space: "**AWS Launches Managed Services.**"³

At Egenera, our wholesale Xterity Cloud Services were designed specifically for partners to help ensure their success. With Xterity, there are no hidden costs, no hidden complexity, and no hidden agendas.

¹ CSP Assessment Checklist download URL: <http://www.egenera.com/csp-assessment-checklist/>

² <https://aws.amazon.com/ec2/faqs/>

³ <http://mspmentor.net/cloud-computing/aws-launches-managed-services>

Table of Contents

AWS Options or Complexity?	4
AWS Support Costs.....	5
AWS Amazon Partner Network Costs.....	8
AWS Purchasing Process - Complexity and Costs.....	15
AWS Service Level (SLA) and Customer Agreement Costs	23
Summary	28

AWS Options or Complexity?

One of the main problems with AWS is that the company tries to be everything to everybody, while ironically, at the same time it limits configuration flexibility and reduces or eliminates your ability to customize solutions to match your client’s needs. This leads to massive complexity and decision-making challenges.

Instead of just providing simple, scalable, and fully customizable solutions, AWS has decision matrices of “options” that require the service of expensive “AWS certified” consultants to decipher the babel. For example, when you look at the various AWS decision stacks (below) you can see how complex it is to do business with them, especially if you just want to provide reliable cloud servers to your clients, without jumping through hoops. Amazon has mastered ease-of-use when it comes to consumer purchases, but it’s a completely different story when it comes to Amazon Web Services.

AWS does have an impressive track record, and is the public cloud market leader, but better choices can be found by doing your homework. We created the various AWS advisories to help you understand the hidden costs that you should be aware of when you assess AWS. We highlight many areas where Xterity Cloud Services are by far, a better choice.

INSTANCE TYPE DECISION		PURCHASING DECISION	STORAGE DECISION	NETWORK DECISION	SCALING DECISION	SUPPORT DECISION	PARTNER NETWORK DECISION
Current Generation (50)	Previous Generation (15)	OnDemand Instances	LOCAL INSTANCE STORE	EC2-CLASSIC	ELASTIC LOAD BALANCING	BASIC	TECHNOLOGY PARTNER
General Purpose	Scheduled Reserved Instances	ELASTIC BLOCK STORAGE (EBS)	EC2-VPC	AUTO SCALING	DEVELOPER	CONSULTING PARTNER	
Memory Optimized	Standard Reserved Instances	General Purpose SSD		CloudWatch	ENTERPRISE	COMPETENCY PARTNER	
Compute Optimized	Spot Instances	Provisioned IOPS SSD		CloudWatch Logs	BUSINESS	REGISTERED	
Storage Optimized	Dedicated Host Instances	Throughput Optimized SSD		RESIZE		STANDARD	
GPU Instances	Pay By The Hour Pricing	Magnetic Volumes (EBS Standard)		MIGRATE		ADVANCED	
Micro Instances	Pay All Upfront Pricing	Cold HDD				PREMIER	
	Pay Partial Upfront Pricing	SIMPLE STORAGE SERVICES (S3)					
	Pay No Upfront Pricing	S3 STANDARD					
	Reserved Pricing	S3 STANDARD -IA Infrequent Access					
	OnDemand Pricing	GLACIER					

AWS Support Costs

To start the complexity, Amazon makes you decide between four different support levels:

- Basic / Developer / Business / Enterprise

The support levels differ in the amount of support you receive and the costs you pay, as shown in the tables on the following pages.^{4 5}

- **AWS Basic support** - Free of charge, but as you'll see in the following pages, when it comes to AWS free support, it's obvious that you get what you pay for.
- **AWS Developer support** - **Fee based**. The amount of support provided at the Developer level is a small incremental step above Basic. You get limited access to **Cloud Support Associates**. Communication is restricted to the hours between 8:00am to 6:00pm (no weekends, no holidays) and is via email.
- **AWS Business support** - **Fee based**. With this support level, you are allowed better access to **Cloud Support Engineers** but there are restrictive limits on case severity/response times, account assistance, and proactive guidance.
- **AWS Enterprise support** - **Fee based**. This AWS support level aligns closely to Egenera's comprehensive partner support, with one major difference: **Cost**. We support our partners free of charge.

Xterity support - simplicity

At Egenera, we don't have different support levels. Our partners receive 24x7x365 L2/L3 support from our engineers. And most importantly, we don't believe that customer support should be viewed as a revenue generating business unit like Amazon does. We realize that our success depends on your success. We don't believe in charging our partners for support that will help make them (and us) successful.

Support level comparisons

	AWS Basic support	AWS Developer support
Customer service & communities	24x7 access to customer service (not technical support) , documentation, whitepapers, and support forums	
Best practices	Access to four core Trusted Advisor Checks	
Technical support		Business hours (8AM-6PM), access to Cloud Support Associates via email, no holidays, no weekends
Who can open cases		One primary contact / Unlimited cases
Case Severity / response times		Normal: < 12 hours Low: < 24 hours
Architecture support		General guidance

⁴ <https://aws.amazon.com/premiumsupport/compare-plans/>

⁵ <https://aws.amazon.com/premiumsupport/pricing/>

	AWS Basic support	AWS Developer support
Launch support		
Programmatic case management		
3rd party software support		
Account assistance		
Proactive guidance		
Pricing	Included	\$49 per month

Your client relationships, and ultimately the success of your business depends on the availability of the cloud services you provide your clients. As you can see in the table above, neither service provides the level of support you're likely to require to be successful. As such, the Basic and Developer support levels can be disqualified as viable options.

Now we'll examine AWS's Business and Enterprise support levels to determine if they provide the level of support you (and your clients) expect, and deserve, from a CSP and at a partner-friendly cost.

	AWS Business support	AWS Enterprise support
Customer service & communities	24x7 access to customer service, documentation, whitepapers, and support forums	
Best practices	Access to full set of Trusted Advisor Checks	
Technical support	24x7 access to Cloud Support Engineers via email, chat & phone	24x7 access to Sr. Cloud Support Engineers via email, chat & phone
Who can open cases	Unlimited contacts / Unlimited cases (IAM supported)	
Case severity / response times	Urgent: < 1 hour, High: < 4 hours, Normal: < 12 hours, Low: < 24 hours	Critical: <15 minutes, Urgent: < 1 hour, High: < 4 hours, Normal: < 12 hours, Low: < 24 hours
Architecture support	Contextual guidance based on your use-case	Consultative review and guidance based on your applications and solutions
Launch support	Infrastructure Event Management (Available for additional fee)	Infrastructure Event Management (Included)
Programmatic case management	AWS Support API	
3rd party software support	Interoperability & configuration guidance and troubleshooting	
Account Assistance		Assigned Support Concierge
Proactive Guidance		Designated Technical Account Manager
Pricing	Greater of \$100 / month (minimum) OR 10% of monthly usage for the 1st \$0-\$10K 7% of monthly usage from \$10K-\$80K 5% of monthly usage from \$80K-\$250K 3% of monthly usage over \$250K	Greater of \$15,000 / month (minimum) OR 10% of monthly usage for the 1st \$0-\$150K 7% of monthly usage from \$150K-\$500K 5% of monthly usage from \$500K-\$1M 3% of monthly usage over \$1M

The Business and Enterprise support levels do provide enhanced support capabilities, but the problem with these offerings are the costs. The costs highlight Amazon’s practice of treating their partners as a source of revenue. Is this any way to treat a partner?

As you see below, our partner support aligns closely with the AWS Enterprise support, but with one major difference: **cost**.

	Egenera - no tiers - the same support for all our partners
Customer service & communities	24x7x365 access to customer service, documentation, whitepapers, and chat
Best practices	Access to all white papers, videos, and best practice guides
Technical support	24x7x365 L2/L3 partner support access to Sr. Cloud Support Engineers via email, chat & phone
Who can open cases	Unlimited contacts / Unlimited cases
Case severity / response times	Severity 1 < 15 minutes Severity 2 < 30 minutes Severity 3 < 2 hours Severity 4 = Next Business Day
Architecture support	Access to architects for high level review and guidance based on your use case.
Launch support	Full on-boarding support includes technical training, assigned account manager, white label marketing & sales content, sales training, announcement and press release messaging support
Programmatic case management	Unlimited access to case management system for all partners and their users
3rd party software	Infrastructure configuration guidance based on your use case
Account assistance	Assigned partner account manager
Proactive guidance	24x7x365 cloud services monitoring and management by Egenera employees
Pricing	No Charge

At Egenera, partner support is a passion; it is not a direct source of revenue. If we do our partner support job right, the results are success for both parties.

A recent customer survey (289 respondents) revealed that we’re doing our job right.

- Trouble ticket responsiveness: 97.6% responded average, good or excellent
- Engineer knowledge: 99% responded average, good or excellent
- Resolution quality: 97.2% responded average, good or excellent.

AWS Amazon Partner Network Costs

At Egenera, we feel it's outrageous to even think about charging fees for a service provider to partner with us. There's no cost for becoming an Egenera partner. This highlights the fact that our wholesale Xterity Cloud Services were designed specifically for the channel knowing that we're only successful when our partners are successful.

The AWS complexity continues when Amazon asks its service providers to choose between three different Amazon Partner Network (APN) levels:

- Technology Partner
- Consulting Partner
- Competency Partner

And just to add to the confusion, each level has sub-levels:

- **Technology Partner:**
 - Registered
 - Standard
 - Advanced
- **Consulting Partner:**
 - Registered
 - Standard
 - Advanced
 - Premier
- **Competency Partner:**
 - Managed Services Program

The various levels differ in the amount of costs you must pay to play. The one common thread in these programs is that if you want to move “up the ladder” you are required to **increase your investment in Amazon**. Just think about that, for you to generate more revenue for your business (and essentially for Amazon) you are required to pay Amazon more.

In the remainder of this section, **red text** is used to indicate additional costs your business may have to pay if you go with AWS.

APN Technology Partner – complexity & costs

*APN Technology Partners can achieve one of three performance tiers based on product status, customer engagements, **and overall business investment on AWS**. These tiers are: Registered, Standard, and Advanced. **By increasing your firm's level of AWS expertise**, you can advance through APN's tiers to receive additional benefits.⁶*

As you'll see throughout this paper, the only way your firm can increase its level of AWS expertise is by increasing your investment in AWS (e.g. training, certifications, etc.). This is a very nice business model for Amazon, but not so much for its partners.

⁶ <http://aws.amazon.com/partners/technology/>

To become an APN Technology Partner your requirements are shown in the table below⁷.

Requirements	Registered	Standard	Advanced
APN Program Fee	N/A	\$2,500/year	\$2,500/year
Products in General Availability on AWS		✓	✓
Minimum Direct, Influenced and Marketplace Monthly AWS Billings OR APN Competency (3 Month Average)			\$50,000 or APN Competency
AWS Well-Architected Review			✓
Customer Reference on AWS		2	6
AWS Support Level		Developer+	Business+
Validated Support Statement for AWS on Partner's Website			✓
AWS Partner Practice Plan			✓
Dedicated Partner Alliance Manager for AWS			✓

The benefits from these programs vary by level achieved. For additional information, refer to the AWS / Egenera partner program benefits comparison tables in the back of this section.

As shown above, the standard and advanced levels require partners to subscribe to AWS support. This “support level” requirement increases the total cost of doing business with AWS. In the first section of this paper we highlighted the various AWS support levels, and their hidden costs.

As mentioned above, increasing AWS expertise is one of the requirements for advancing up the APN partner ladder. To increase your business’ AWS expertise, you end up paying for your employees to attend AWS training classes so they can pass the requisite AWS certification exams. On one hand this is a good thing but on the other hand, this creates a Catch-22 problem. Once your employees become AWS certified they immediately become more valuable in the job market.

Retention of skilled employees is a common challenge for service providers, and employees with cloud expertise are in high demand. What happens if you’ve paid for an employee to become AWS certified so that your business can achieve a higher APN level just to see that employee (and the certification you paid for) walk out the door?

APN Consulting Partner – complexity & costs

APN Consulting Partners can achieve one of four performance tiers based on *training*, customer engagements, and *overall business investment on AWS*. These tiers are: Registered, Standard, Advanced, and Premier. By *increasing your firm's level of AWS expertise* and engagement, you can advance through APN's tiers to receive additional benefits.⁸

The partner requirements for becoming an APN Consulting Partner are shown in the table below⁶.

Requirements	Standard	Advanced	Premier
APN Program Fee	\$2,500/year	\$2,500/year	\$2,500/year
Customer References on AWS	2	6	10 within the past 12 months

⁷ <http://aws.amazon.com/partners/technology/>

⁸ <http://aws.amazon.com/partners/consulting/>

Requirements	Standard	Advanced	Premier
AWS Technical Professional Accreditations or Instructor-Led Training	2	8	20
AWS Business Professional Accreditations	2	8	20
AWS TCO & Cloud Economics Accreditations		4	8
AWS Certifications - Associate Level	2	4	20
AWS Certifications - Professional Level		2	8
Minimum Direct and Influenced AWS Billings (3-month average)	\$1K	\$50K	Must achieve minimum AWS billing requirements per region
AWS Support Level	Developer+	Business+	Business+
Advanced Level for 12+ months			✓
APN Competency			1
Audited Managed Service Provider or DevOps Consulting Competency			✓
AWS Partner Practice Plan			✓
Customer Satisfaction (CSAT) Review			✓
Dedicated Partner Alliance Manager for AWS			✓
Dedicated Project Manager for AWS			✓
Adheres to AWS Security Best Practices			✓
Demonstrates quality delivery capabilities across designated country or region			✓

Managed Service Program – complexity & costs

If you aspire to become a member of Amazon’s Managed Service Program, it’s going to take time and money. To qualify for membership in the program you must satisfy a list of costly prerequisites as shown in the table below.⁹

Program Prerequisites	
APN Membership	Advanced or Premier APN Consulting Partner
AWS Billings	\$50,000/month in AWS Revenue (direct/indirect)
Customer Engagements	≥ 6 AWS Customer References (nonpublic or public)
AWS Certifications	≥ 6 AWS Certified Staff
AWS Training	≥ 2 Trained Staff System Operations on AWS or DevOps Engineering on AWS
MSP Validation	Download the AWS Managed Service Partner Validation Checklist to prepare for operational process and capabilities review

After applying for membership and building your practice to satisfy the prerequisites, you are then subject to a partner capabilities audit. *The APN Partner’s cost for the Managed Service Program audit will be \$3,000 USD, plus auditor travel costs paid to the auditing firm.*⁷

AWS certifications – complexity & costs

⁹ <https://aws.amazon.com/partners/managed-service/>

In several of the tables shown above there are requirements for partners to obtain various types and quantities of AWS certifications. AWS certifications are listed in the table below.

AWS Certifications ¹⁰	
Associate	Professional
AWS Certified Solutions Architect – Associate	AWS Certified Solutions Architect – Professional
AWS Certified Developer – Associate	AWS Certified DevOps Engineer – Professional
AWS Certified SysOps Administrator – Associate	

Let’s look at some of the required certification exam and recommended course costs.¹¹

Certification	Practice Exam Registration Fee	Exam Registration Fee	Recommended Courses
Certified Solutions Architect - Associate	USD \$20	USD \$150	Certification Exam Readiness Workshop - USD \$225
Certified Developer - Associate	USD \$20	USD \$150	Developing on AWS - USD \$2095
Certified SysOps Administrator - Associate	USD \$20	USD \$150	AWS Technical Essentials - USD \$695 System Operation on AWS - USD \$2095
Certified Solutions Architect - Professional	USD \$40	USD \$300	Advanced Architecting on AWS - USD \$2095
Certified DevOps Engineer - Professional		USD \$300	DevOps Engineering on AWS - USD \$2095

It all adds up – Amazon partner Network complexity & costs

Listed in the table below are APN’s various partner programs and the required certifications. The USD figures shown are per individual.

Certification	Practice Exam Registration Fee	Exam Registration Fee	Recommended Courses	APN Consulting Partner <Standard>	APN Consulting Partner <Advanced>	APN Consulting Partner <Premier>
Certified Solutions Architect - Associate	USD \$20	USD \$150	Certification Exam Readiness Workshop USD \$225	(2) AWS Associate Level certifications required	(4) AWS Associate Level certifications required	(20) AWS Associate Level certifications required
Certified Developer - Associate	USD \$20	USD \$150	Developing on AWS USD \$2095			
Certified SysOps Administrator - Associate	USD \$20	USD \$150	AWS Technical Essentials USD \$695 System Operation on AWS USD \$2095			
Certified Solutions Architect - Professional	USD \$40	USD \$300	Advanced Architecting on AWS USD \$2095		(2) AWS Professional Level certifications required	(8) AWS Professional Level certifications required
Certified DevOps Engineer - Professional		USD \$300	DevOps Engineering on AWS USD \$2095			

¹⁰ <http://aws.amazon.com/partners/training/>

¹¹ <https://aws.amazon.com/certification/>

The figures shown in the table above reflect only the course and exam costs, and not all the additional costs and requirements shown in the previous tables. While it might be possible to receive discounts for some of the programs shown above, the point here is, why should you have to pay Amazon even one penny to become a more effective services provider of AWS offerings? After you consider these extra costs imposed by Amazon, you can begin to see why the total cost of doing business with Amazon usually results in higher costs and smaller margins for the service provider.

Egenera Partner Program – simplicity & no costs

At Egenera, we don't have different partner program levels. Qualifying for our partner program is as simple as filling out an **online form**¹² with your contact information. We'll get in touch with you shortly thereafter. After a brief conversation to ensure we're a match for your requirements we'll send you our Master Services Agreement (MSA). Once you review the MSA, approve it, and send it back, you're on your way to delivering high margin cloud services without the hidden costs or sales commitments like the large CSPs impose. Most of our partners usually return a signed copy of the MSA shortly after they receive it, and soon thereafter they're typically performing a proof of concept (PoC).



APN Consulting Partner Benefits comparison with Egenera Partner Program

As you review the tables below, keep in mind that Egenera does not view its partner program as a source of revenue. We view our partner program as a cost to our business that we recoup when our partners are successful. The strange thing about the APN is that the loftier your goals, the more you'll have to **invest in Amazon**. We feel this is no way to treat partners.

*APN Consulting Partners can achieve one of four performance tiers based on **training**, customer engagements, and overall **business investment on AWS**. These tiers are: Registered, Standard, Advanced, and Premier. By **increasing your firm's level of AWS expertise** and engagement, you can advance through APN's tiers to receive additional benefits.¹³*

	Registered	Standard	Advanced	Premier	Egenera
APN Partner Portal for Partner-only content and enablement	✓	✓	✓	✓	✓ Egenera partner portal
Access to APN Webcast for technical and program videos	✓	✓	✓	✓	✓ Egenera partner portal
Online Business & Technical Training & Accreditations	✓	✓	✓	✓	✓ Technical and sales training provided – no cost
20% Discount on AWS Instructor-Led Training	✓	✓	✓	✓	✓ No discount required – our partners receive training from Egenera instructors – no cost
AWS Solutions Training for Partners	✓	✓	✓	✓	✓ We provide Xterity Cloud Services training for all partners – no cost
Use of APN Logo		✓	✓	✓	✓ Our partners have access to any material they need to successful – no cost
Company Profile in the AWS Partner Directory		✓	✓	Featured	✓ We're happy to list partners on the Egenera website – no cost
Access to self-service and agency marketing platform – APN Marketing Central		✓	✓	✓	✓ White label & Egenera marketing support – no cost

¹² Partner with us online form URL: <http://www.egenera.com/partner-with-us/>

¹³ <http://aws.amazon.com/partners/consulting/>

	Registered	Standard	Advanced	Premier	Egenera
Eligible for free AWS usage via the Innovation Sandbox		✓	✓	✓	✓ Access to Xterity cloud infrastructure for POC – no cost
Eligible for free AWS usage for qualified Proof-of-Concepts		✓	✓	✓	✓ Access to Xterity cloud infrastructure for POC – no cost
Eligible for AWS Instructor-Led Training Vouchers		✓	✓	✓	✓ No vouchers required – our partners receive training from Egenera Instructors – no cost
Eligible for the AWS Channel Reseller Program		✓	✓	✓	✓*see note
Eligible for the AWS Managed Service Program		✓	✓	✓	✓*see note
Eligible for the APN Competency Program		✓	✓	✓	✓*see note
Eligible for the AWS Test Drive		✓	✓	✓	✓*see note
Eligible for Market Development Funding			✓	✓	✓*see note
Eligible for Customer-Facing Joint Webinars with AWS				✓	✓*see note
Featured on APN Premier, APN Program or Solution Webpages				✓	✓ We're happy to list our partners on our website – no cost
Eligible for prioritized marketing activities				✓	✓ We impose no eligibility requirements. We provide go to market assistance for all our partners – no cost
Eligible for AWS-Written Case Study				✓	✓ We write case studies and work with our partners to ensure technical accuracy and effective messaging – no cost
AWS Professional Service Boot Camp Training				✓	✓ We provide our partners with on-boarding training to ensure they can rapidly and efficiently begin selling cloud services to their clients – no cost
Eligible for up to 2 weeks free AWS Professional Services for Qualified Enterprise Opportunities				✓	✓ We do have billable professional services to help partners with unique situations that warrant special attention
Named AWS Partner Account Manager				✓	✓ Every partner has an assigned account manager – no cost
Named AWS Solution Architect				✓	✓ We provide our partners architecture guidance based on their specific use cases.

*Note: From marketing and sales programs to technical product requirements we work hand-in-hand with our partners to ensure they have whatever it takes for them to be successful.

APN Technology Partner Benefits - comparison with Egenera Partner Program

APN Technology Partners can achieve one of three performance tiers based on product status, customer engagements, *and overall business investment on AWS*. These tiers are: Registered, Standard, and Advanced. By *increasing your firm's level of AWS expertise*, you can advance through APN's tiers to receive additional benefits.¹⁴

	Registered	Standard	Advanced	Egenera
APN Partner Portal for Partner-only content and enablement	✓	✓	✓	✓ Egenera partner portal
Access to APN Webcast for technical and program videos	✓	✓	✓	✓ Egenera partner portal
Online Business & Technical Training & Accreditations	✓	✓	✓	✓ Technical and sales training provided - no cost
20% Discount on AWS Instructor-Led Training	✓	✓	✓	✓ No discount required - our partners receive training from Egenera instructors - no cost
Eligible for the AWS SaaS Program	✓	✓	✓	✓ We work with our partners to assist them in delivering SaaS solutions to their clients on our cloud infrastructure. Our services eliminate the infrastructure "heavy lifting," freeing our partners to focus on delivering high value, high margin software solutions.
Use of APN Logo		✓	✓	✓ Our partners have access to any material they need to successful - no cost
Company Profile in the AWS Partner Directory		✓	✓	✓ We're happy to list partners on the Egenera website - no cost
Access to self-service and agency marketing platform - APN Marketing Central		✓	✓	✓ White label & Egenera marketing support - no cost
Eligible for free AWS usage via the Innovation Sandbox		✓	✓	✓ Access to Xterity cloud infrastructure for POC - no cost
Eligible for free AWS usage for qualified Proof-of-Concepts		✓	✓	✓ Access to Xterity cloud infrastructure for POC - no cost
Eligible for Free AWS usage for Free Trial Campaigns		✓	✓	✓ Access to Xterity cloud infrastructure for POC - no cost
Eligible for AWS Instructor-Led Training Vouchers		✓	✓	✓ No vouchers required - our partners receive training from Egenera instructors - no cost
Eligible for the APN Competency Program		✓	✓	✓*see note
Eligible for the AWS Test Drive		✓	✓	✓*see note
Eligible for joint AWS Customer-Facing Webinars			✓	✓*see note
Eligible for the Market Development Funding			✓	✓*see note
Eligible for AWS-Written Case Study			✓	✓ We write case studies and work with our partners to ensure technical accuracy and effective messaging - at no cost

*Note: From marketing and sales programs to technical product requirements we work closely with our partners to ensure they have whatever it takes for them to be successful.

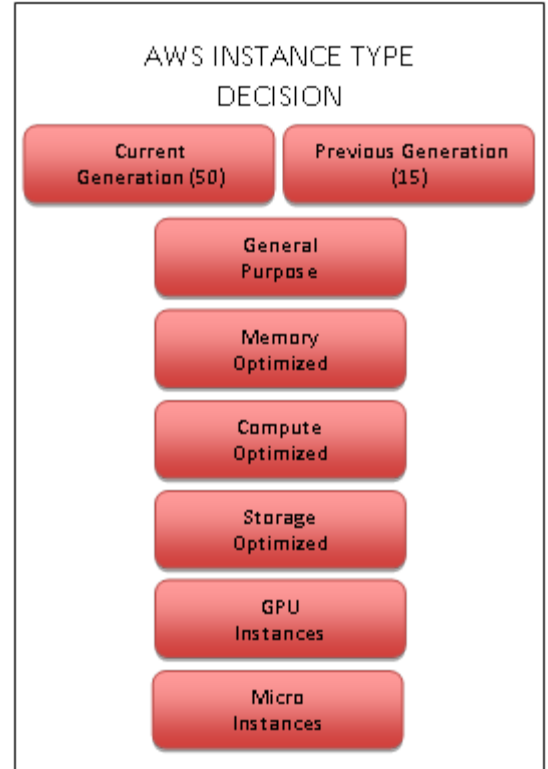
¹⁴ <http://aws.amazon.com/partners/technology/>

AWS Purchasing Process - Complexity and Costs

We built our cloud platform from the ground up with the mandates of self-service simplicity, utilizing enterprise-class hardware, and we host Xterity in world-class datacenters, and back it up with 24x7x365 L2/L3 partner support.

AWS instance (server) type decision - AWS complexity and costs

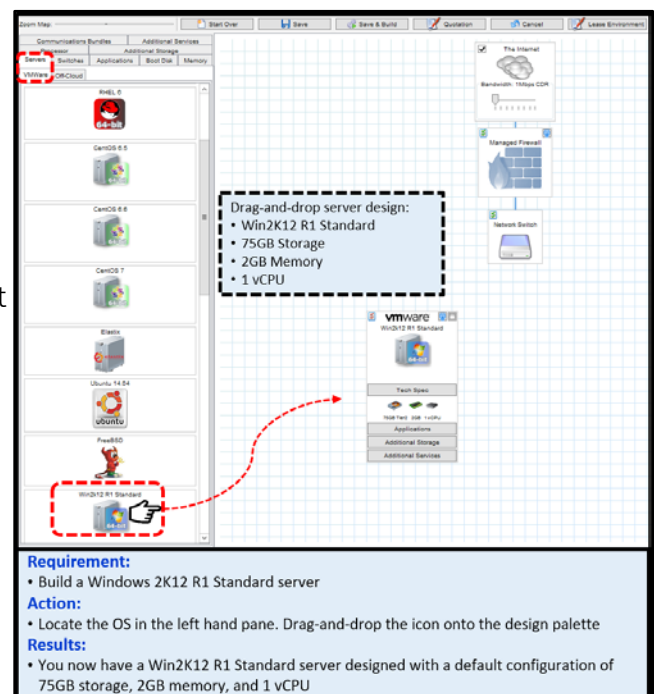
To begin with, Amazon makes you decide between several different instance types. AWS refers to virtual servers as “instances.” According to the AWS website¹⁵ there are 50+/- current generation instances. To complicate matters, AWS still offers 15 instances¹⁶ referred to as “previous generation.” AWS instances are categorized (configured) according to how AWS feels they are sized for various use cases. All the AWS instances are preconfigured (static) and you can’t easily modify the configurations. You must buy what AWS offers without any simple method of customizing the instance to meet your client’s actual requirements. This could result in performance issues due to incorrect sizing, and can put you in an unenviable position because your clients will most likely be dissatisfied with the service; which in their eyes, is being provided by you. AWS then makes you determine which preconfigured instance aligns closest to your client’s workloads, and the he issue with preconfigured configurations is that server workloads are variable.



In all fairness, AWS does provide processes for addressing variable workloads, but as you’ll see later in this paper, the processes are complex, costly and result in downtime. Server downtime might be acceptable for developers working in dev/test environments, but downtime is not acceptable for today’s 24x7x365 business world where server uptime is critical.

Server types - Xterity simplicity

The efficient (and partner friendly) way to address changeable workloads is to stay away from predefining static configurations by allowing partners to use an intuitive cloud management interface that provides the ability to adjust vCPU, memory, storage and networking specifications via drag-and-drop, or slide-bar simplicity. This method removes all complexity and facilitates “right sizing” simplicity. The Xterity platform was designed specifically with partner ease-of-use in mind.



¹⁵ <https://aws.amazon.com/ec2/instance-types/>

¹⁶ <https://aws.amazon.com/ec2/previous-generation/>

Server purchasing processes – AWS complexity and costs

Cloud computing is supposed to be synonymous with simplicity, with utility pricing where cloud services consumers pay for what they use. This seems straightforward. However, if you look at the AWS pricing maze¹⁷ you'll quickly determine that AWS pricing processes weren't designed with simplicity in mind. If you want to just launch an AWS instance (server) for a client you need to understand what Amazon means by:

- OnDemand instances
- Standard Reserved instances
- Scheduled Reserved instances
- Spot instances
- Dedicated instances
- Dedicated Host instances

Then, you need to determine which pricing policy (below) is available for the instance type you selected (above):

- Pay By The Hour pricing
- Pay All Upfront pricing
- Pay Partial Upfront pricing
- Pay No Upfront pricing
- Reserved pricing
- OnDemand pricing

Be aware that instance types, pricing policies, and even prices vary by region. In addition, each instance type has its own list of caveats that need to be highlighted. Here are just a few:

- **OnDemand Instances:** *“during periods of very high demand, it is possible that you might not be able to launch specific On-Demand instance types in specific Availability Zones”¹⁸*
- **Standard Reserved Instances:** Requires a **1-3 year commitment**. *“Determine the number of instances you want to run and the term length (1 or 3 years).”¹⁹ “When you purchase a Reserved Instance, you are billed for every hour during the entire Reserved Instance term that you select, regardless of whether the instance is running or not.”²⁰*
- **Scheduled Reserved Instances:** *“Scheduled Reserved Instances have a 1 year term commitment.”²¹ “Amazon EC2 must ensure that the EC2 instances have terminated by the end of the current scheduled time period so that the capacity is available for any other Scheduled Instances it is reserved for. Therefore, Amazon EC2 terminates the EC2 instances three minutes before the end of the current scheduled time period. You can't stop or reboot Scheduled Instances, but you can terminate them manually as needed. If you terminate a Scheduled Instance before its current scheduled time period ends, you can launch it again after a few minutes. Otherwise, you must wait until the next scheduled time period.”²²*

AWS PURCHASING DECISION

OnDemand Instances

Scheduled Reserved Instances

Standard Reserved Instances

Spot Instances

Dedicated Host Instances

Pay By The Hour Pricing

Pay All Upfront Pricing

Pay Partial Upfront Pricing

Pay No Upfront Pricing

Reserved Pricing

OnDemand Pricing

¹⁷ <https://aws.amazon.com/ec2/pricing/>

¹⁸ <https://aws.amazon.com/ec2/purchasing-options/>

¹⁹ <https://aws.amazon.com/ec2/purchasing-options/reserved-instances/getting-started/>

²⁰ <https://aws.amazon.com/ec2/pricing/>

²¹ <https://aws.amazon.com/ec2/purchasing-options/reserved-instances/>

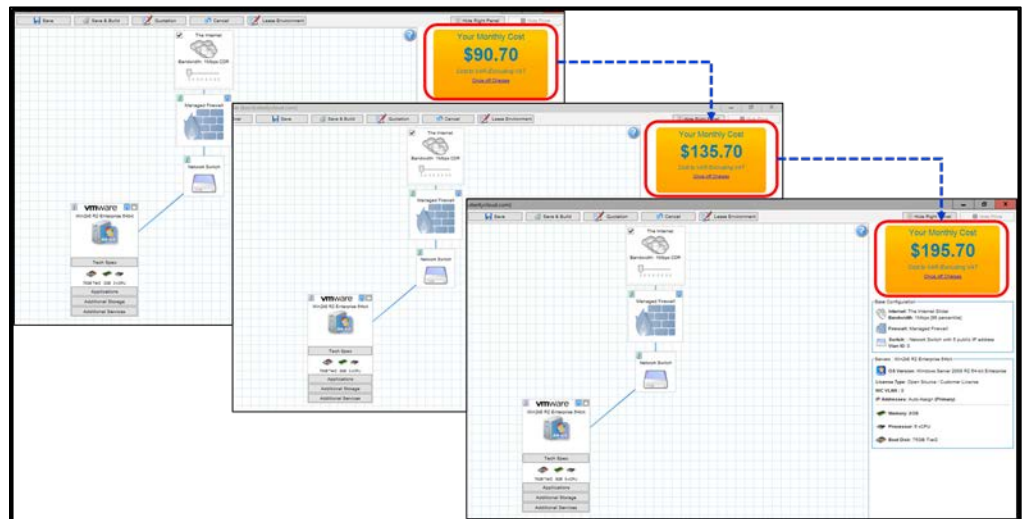
²² <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-scheduled-instances.html>

- **Spot instances:** You specify the maximum hourly price that you're willing to pay for an instance. AWS controls the "going rate" spot instance prices. *"If the Spot Price moves higher than a customer's maximum price, the customer's instance will be shut down by Amazon EC2".*²³ *"While we attempt to cleanly terminate your instances, your application should be prepared to deal with sudden shutdowns."*²⁴
- **Dedicated Instances and Dedicated Hosts:** *"An additional fee (\$2 per hour) is charged once per hour in which at least one Dedicated Instance of any type is running in a region."*²⁵

One additional "purchasing/cost challenge" you will encounter by using a large CSP is that **their price lists are readily accessible on the internet**. As a matter of fact, Amazon highlights the fact that they've dropped AWS prices 51 times.²⁶ Think about that. AWS end users (your clients) will be well aware of AWS' publicized cost reductions and they'll expect you to pass the savings on to them. The point being is that if you go with a large CSP that publishes its price lists and touts their cost reductions, **your pricing and margin strategies are dictated by the CSP**. Your ability to manage your margins is gone.

Server purchasing process - Xterity simplicity

The right way to price cloud services is to deliver on the promise of cloud by offering simple, straightforward pricing with **no hidden costs**. In addition to upfront configuration and pricing mechanisms, the cloud platform should provide comprehensive business back-end features that help partners with financial processes via consistent (wholesale), non-public costs, with an interface that provides real-time cost, pricing and margins analysis. The back-end support should enable partners to streamline their billing processes via custom quotes and invoicing.



Below are three Xterity Cloud Designer screen shots that show how the monthly cost information updates in real-time as a server's configuration is dynamically modified via drag-and-drop simplicity. The costs shown do not represent actual costs. The figures shown are for illustrative purposes only.

²³ <https://aws.amazon.com/ec2/faqs/#spot-instances>

²⁴ [ec2-downloads.s3.amazonaws.com/Intro-to-Spot-Instances.pdf](https://aws.amazon.com/ec2-downloads.s3.amazonaws.com/Intro-to-Spot-Instances.pdf)

²⁵ <https://aws.amazon.com/ec2/purchasing-options/dedicated-instances/>

²⁶ <https://aws.amazon.com/blogs/aws/happy-new-year-ec2-price-reduction-c4-m4-and-r3-instances/>

Storage purchasing process – AWS complexity and costs

With AWS storage, the decision process once again is complex. You must decide between local instance (server) storage, and networked storage. Networked storage is in the form of Elastic Block Storage™ (EBS), Simple Storage Services™ (S3), or Glacier. Glacier is long term storage with retrieval time in the hours so it will not be discussed here.

If you decide to use local instance storage, there are a couple of caveats:

- Local instance storage is alive only while the instance is alive.
- If the instance is shut down, the data stored in local instance storage is gone.

S3 storage pricing is complicated by the fact that there are two flavors:

- S3 Standard
- S3 Standard-IA (infrequent access).

It seems that picking EBS is the simple, obvious choice. Once again, the waters are muddied with decisions. EBS storage is available in various media types including solid state disks (SSD), hard disk drives (HDD), and magnetic volumes.

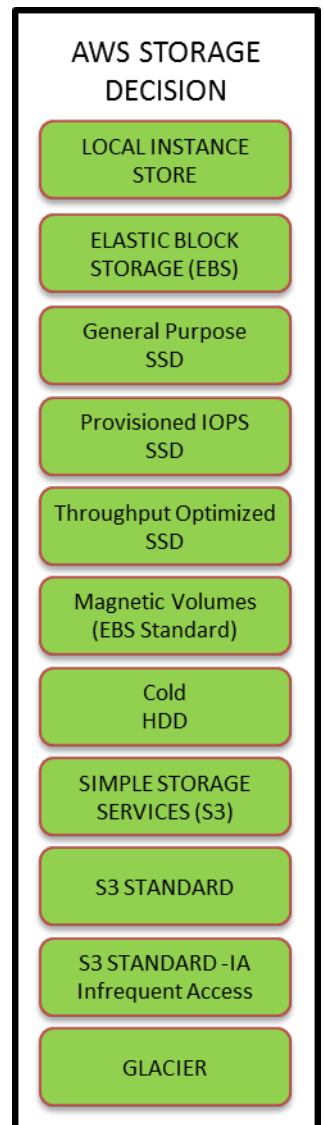
SSD is the preferred storage media and the simplest choice, but now you must decide between:

- General Purpose SSD
- Provisioned IOPS SSD
- Throughput Optimized SSD

And, if you want the highest performing networked storage possible you must select an instance (server) that is pre-configured as “EBS-Optimized.”

One of the biggest “gotcha” moments with AWS is when the invoice arrives and you’re faced with many **hidden storage related costs that are in addition to the actual provisioned storage:**

- **EBS Provisioned IOPs (SSD) volume:** *“You will be charged for the IOPS provisioned on a volume even when the volume is detached from an instance.”²⁷*
- **EBS Snapshots:** *“Copying EBS snapshots is charged for the data transferred across regions. After the snapshot is copied, standard EBS snapshot charges apply for storage in the destination region.”²⁸*
- **S3 Storage costs:** With S3 storage you can be charged for storage (per GB), request pricing (per PUT, COPY, POST, LIST, GET request), and data transfer out pricing.²⁹
- **Elastic Load Balancing (ELB) data costs:** Earlier it was mentioned that AWS has some methods of addressing changeable workloads. Horizontal scaling can be achieved via ELB. Data transferred “in” to and “out” of ELB is priced equivalent to Amazon EC2. In other words, Amazon



²⁷ <https://aws.amazon.com/ebs/pricing/>

²⁸ <https://aws.amazon.com/ebs/pricing/>

²⁹ <https://aws.amazon.com/s3/pricing/>

charges for the *“data transferred through the Elastic Load Balancer”*³⁰ in addition to the other data related charges.

And, the storage costs and complexities don’t stop there:

- Storage costs vary by region
- There is no charge for transferring data into EC2 unless you are using a public or Elastic IP address, or the source is in another availability zone (datacenter) or peered VPC in the same AWS region
- There is no charge for transferring data out from EC2 unless you are using a public or Elastic IP address or the destination is in another availability zone (datacenter) or peered VPC in the same AWS region
- **Data transferred from EC2 out to the Internet is priced on a volume transferred amount.**

But the news isn’t all bad. The 1st GB/month is free!

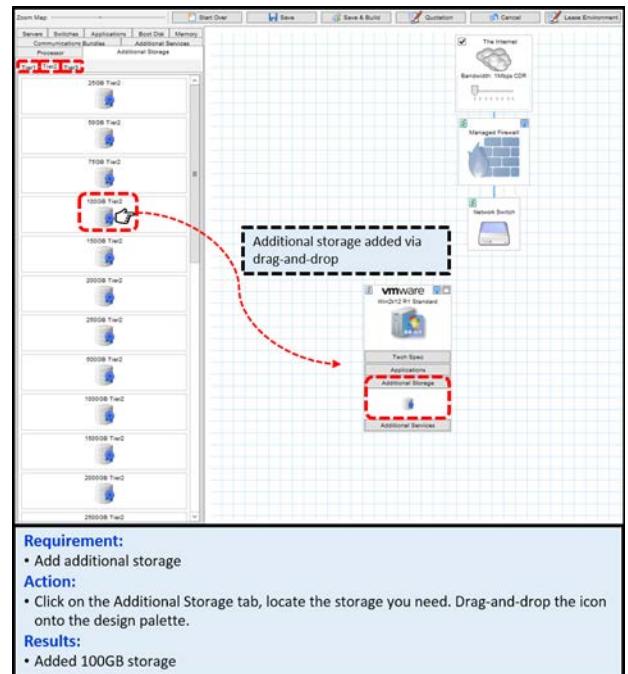
Storage purchasing - Xterity simplicity

Xterity’s interface lets you add storage via simple drag-and-drop. You only need to determine two things:

- Required storage capacity - (dynamically scalable)
- The storage performance you need - Tier 1, 2 or 3

We don’t impose data (volume) transfer fees.

There has been much press related to **Amazon’s data transfer pricing policies**. The reactions are unfavorable. AWS customers feel that the data transfer out costs lock them into AWS with a “you can check out any time you like, but you can never afford to leave” costly death grip.



³⁰ <https://aws.amazon.com/elasticloadbalancing/pricing/>

Scalability processes - AWS complexity

As previously mentioned, AWS instances are preconfigured and static in that you cannot change an instance configuration without jumping through hoops. With AWS, there are multiple (complex) ways to manually increase or decrease the size of an instance. AWS provides scaling processes for:

- Vertical scaling (scale up)
- Horizontal scaling (scale out)
- Scale up/down

But each of these processes come with **costs and complexity**.

Vertical Scaling - AWS complexity

Amazon equates “resizing” to changing an instance type: *“If the root device for your instance is an EBS volume, you can change the size of the instance simply by changing its instance type, which is known as resizing it. If the root device for your instance is an instance store volume, you must migrate your application to a new instance with the instance type that you want.”*³¹

But, the instance type you want to change to has to be compatible with the configuration of the original instance. If the instance types aren’t compatible you have to migrate your application.

Another caveat is that *“When you resize an instance, the resized instance usually has the same number of instance store volumes that you specified when you launched the original instance. If you want to add instance store volumes, you must migrate your application to a completely new instance with the instance type and instance store volumes that you want.”*³²

How complex can it be to resize an EBS-backed instance since just the instance type is changing?

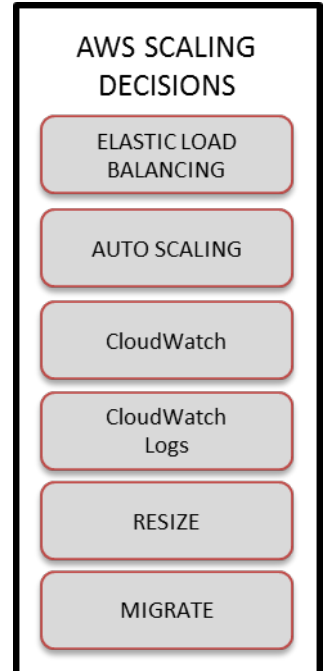
Per Amazon, *when you stop a running instance, the following happens:*³³

The instance performs a normal shutdown and stops running; its status changes to stopping and then stopped.

- Any EBS volumes remain attached to the instance and their data persists
- Any data stored in the RAM of the host computer or the instance store volumes of the host computer is gone
- The instance is migrated to a new underlying host computer when it’s started.

That seems straightforward. But, there could be possible networking issues of varying complexity depending on which AWS networking technology your AWS account uses.

- *EC2-Classic: We release the public and private IP addresses for the instance when you stop the instance, and assign new ones when you restart it.*
- *EC2-VPC: The instance retains its private IP addresses when stopped and restarted. We release the public IP address and assign a new one when you restart it.*



³¹ <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>

³² <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html>

³³ http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Stop_Start.html

- *EC2-Classic: We disassociate any Elastic IP address that's associated with the instance. You're charged for Elastic IP addresses that aren't associated with an instance. When you restart the instance, you must associate the Elastic IP address with the instance; we don't do this automatically.*
- *EC2-VPC: The instance retains its associated Elastic IP addresses. You're charged for any Elastic IP addresses associated with a stopped instance.*
- *When you stop and start a Windows instance, the EC2Config service performs tasks on the instance such as changing the drive letters for any attached Amazon EBS volumes.*
- *If you've registered the instance with a load balancer, it's likely that the load balancer won't be able to route traffic to your instance after you've stopped and restarted it. You must de-register the instance from the load balancer after stopping the instance, and then re-register after starting the instance.*
- *If your instance is in an Auto Scaling group, the Auto Scaling service marks the stopped instance as unhealthy, and may terminate it and launch a replacement instance.*
- *When you stop a ClassicLink instance, it's unlinked from the VPC to which it was linked. You must link the instance to the VPC again after restarting it.*

Auto-scaling is a no cost service; however, it requires the use of Cloud Watch™ (a fee-based service).³⁴

This level of complexity and disruption might be acceptable in a dev/test environment but this is not what you want in a production environment. As mentioned earlier, AWS was designed for developers, not service providers.

How about changing from one Instance store-backed instance to a different instance type? Should be simple, right?

When you want to move your application from one instance store-backed instance to an instance store-backed instance with a different instance type, *you must migrate it by creating an image from your instance, and then launching a new instance from this image with the instance type that you need. To ensure that your users can continue to use the applications that you're hosting on your instance uninterrupted, you must take any Elastic IP address that you've associated with your original instance and associate it with the new instance. Then you can terminate the original instance.*³⁵

- *To migrate an instance store-backed instance*
 1. *[EC2-Classic] If the instance you are migrating has an associated Elastic IP address, record the Elastic IP address now so that you can associate it with the new instance later.*
 2. *Back up any data on your instance store volumes that you need to keep to persistent storage. To migrate data on your EBS volumes that you need to keep, take a snapshot of the volumes. Create an AMI from your instance store-backed instance by satisfying the prerequisites and following the procedures in Creating an Instance Store-Backed Linux AMI. When you are finished creating an AMI from your instance, return to this procedure.*
 3. *Open the Amazon EC2 console and in the navigation pane, select AMIs. From the filter lists, select Owned by me, and select the image that you created in the previous step.*

³⁴ <https://aws.amazon.com/cloudwatch/pricing/>

³⁵ <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-resize.html#migrate-instance-configuration>

Notice that AMI Name is the name that you specified when you registered the image and Source is your Amazon S3 bucket.

- Note
 - If you do not see the AMI that you created in the previous step, make sure that you have selected the region in which you created your AMI.
- 4. Choose Launch. When you specify options for the instance, be sure to select the new instance type that you want. If the instance type that you want can't be selected, then it is not compatible with configuration of the AMI that you created (for example, because of virtualization type). You can also specify any EBS volumes that you detached from the original instance.
 - Note that *it can take a few minutes for the instance to enter the running state.*
- 5. [EC2-Classic] If the instance that you started with had an associated Elastic IP address, you must associate it with the new instance as follows:
 - a In the navigation pane, choose Elastic IPs.
 - b Select the Elastic IP address that you recorded at the beginning of this procedure.
 - c Choose Actions and then choose Associate Address.
 - d From Instance, select the new instance, and then choose Associate.

Does it have to be this complex?

How about scale-out by using Elastic Load Balancing (ELB)?

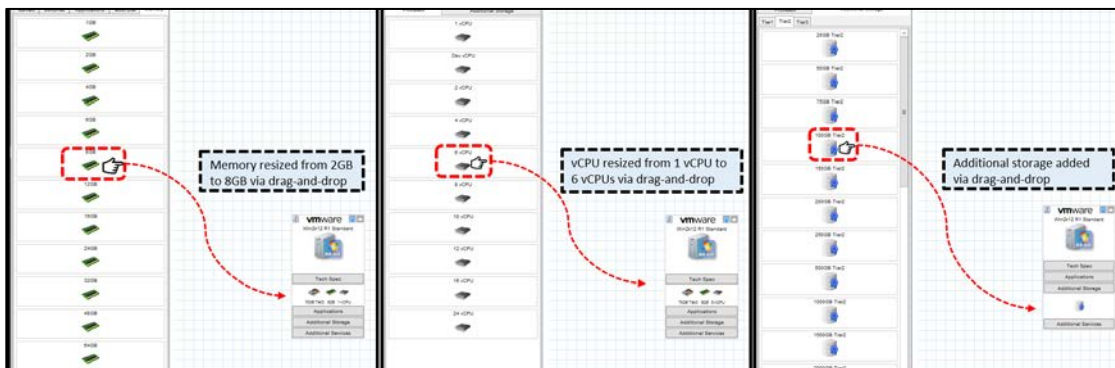
Keep in mind that when you use ELB **you are charged for the data that goes through the load balancer in addition to the other "regular" storage costs.**

If all you want to do is scale a server for your client, why should you be forced to use another product and incur additional costs? And, what if your client's application doesn't horizontally scale?

SCALING DECISION - Xterity simplicity

If you made it through all the caveats and issues mentioned in the previous pages, you're probably thinking there should be a better way to perform such a simple task as scaling a server -- after all isn't this a standard function of server virtualization?

Xterity provides a simple, intuitive interface that allows you to scale servers with a process that's as simple as drag-and-drop.



AWS Service Level (SLA) and Customer Agreement Costs

In this paper, you'll see how the fine print (legalese) within the Amazon Web Services Customer Agreement³⁶ and (multiple) Service Level Agreements were written in a manner that dilute any chance of you being rewarded compensation for AWS outages. And because your clients view you as their cloud server provider, and while Amazon gets off scot-free, your client relationships, and reputation are damaged by AWS outages.

As has been highlighted in the previous sections, Amazon doesn't do anything simple. Complexity reigns supreme with AWS and this is true with AWS' customer and service level agreements. There isn't just a single SLA to review. Amazon has multiple SLAs and a "Customer Agreement" that need to be scrutinized.

The main concern with an SLA or customer agreement from any CSP is that you need to do your homework to get a sense of whom the agreements were written to protect. This will provide insight as to what kind of partnership you're likely to have with the CSP. And remember, your CSP is an extension of your business, your client relationships, and your reputation. Your business success is in the hands of your CSP. Do you want to partner with a CSP that crafted their SLA and customer agreement documents to protect themselves?

Customer agreement – written to protect whom?

Looking at the AWS Customer Support Agreement³⁷, we found several statements that may be alarming to partners:

- **2.2 Changes To the APIs:**
 - *We may change, discontinue or deprecate any APIs for the Services from time to time but will use commercially reasonable efforts to continue supporting the previous version of any API changed, discontinued, or deprecated for 12 months after the change, discontinuation, or deprecation (**except if doing so** (a) would pose a security or intellectual property issue, (b) **is economically or technically burdensome**, or (c) is needed to comply with the law or requests of governmental entities).*
- **4.2 Other Security and Backup.**
 - ***You are responsible for properly configuring and using the Service Offerings and taking your own steps to maintain appropriate security, protection and backup of Your Content**, which may include the use of encryption technology to protect Your Content from unauthorized access and routine archiving Your Content. AWS log-in credentials and private keys generated by the Services are for your internal use only and you may not sell, transfer or sublicense them to any other entity or person, except that you may disclose your private key to your agents and subcontractors performing work on your behalf.*
- **7.2 Termination.**
 - (a) *Termination for Convenience. You may terminate this Agreement for any reason by: (i) providing us notice and (ii) closing your account for all Services for which we provide an account closing mechanism. **We may terminate this Agreement for any reason** by providing you 30 days advance notice.*
 - (b) *Termination for Cause.*

³⁶ <https://aws.amazon.com/agreement/>

³⁷ <https://aws.amazon.com/agreement/>

(i) *By Either Party.* Either party may terminate this Agreement for cause upon 30 days advance notice to the other party if there is any material default or breach of this Agreement by the other party, unless the defaulting party has cured the material default or breach within the 30 day notice period.

(ii) *By Us.* We may also terminate this Agreement immediately upon notice to you (A) for cause, if any act or omission by you or any End User results in a suspension described in Section 6.1, (B) if our relationship with a third party partner who provides software or other technology we use to provide the Service Offerings expires, terminates or requires us to change the way we provide the software or other technology as part of the Services, (c) **if we believe providing the Services could create a substantial economic or technical burden or material security risk for us**, (D) in order to comply with the law or requests of governmental entities, or (E) if we determine use of the Service Offerings by you or any End Users or our provision of any of the Services to you or any End Users has become impractical or unfeasible for any legal or regulatory reason.

- **8.6 Suggestions.** *If you provide any Suggestions to us or our affiliates, we will own all right, title, and interest in and to the Suggestions, even if you have designated the Suggestions as confidential. We and our affiliates will be entitled to use the Suggestions without restriction. You hereby irrevocably assign to us all right, title, and interest in and to the Suggestions and agree to provide us any assistance we may require to document, perfect, and maintain our rights in the Suggestions.*
- **13.3 Independent Contractors; Non-Exclusive Rights.** *We and you are independent contractors, and neither party, nor any of their respective affiliates, is an agent of the other for any purpose or has the authority to bind the other. Both parties reserve the right (a) to develop or have developed for it products, services, concepts, systems, or techniques that are similar to or compete with the products, services, concepts, systems, or techniques developed or contemplated by the other party and (b) to assist third party developers or systems integrators who may offer products or services which compete with the other party’s products or services.*

Now, lets examine some statements found in the Amazon Elastic Compute Cloud (EC2) service level agreement.³⁸

Service Commitment

AWS will use commercially reasonable efforts to make Amazon EC2 and Amazon EBS each available with a Monthly Uptime Percentage (defined below) of at least 99.95%, in each case during any monthly billing cycle (the “Service Commitment”). In the event Amazon EC2 or Amazon EBS does not meet the Service Commitment, you will be eligible to receive a Service Credit as described below.

Monthly Uptime Percentage	Service Credit Percentage
Less than 99.95% but equal to or greater than 99.0%	10%
Less than 99.0%	30%

Amazon advertises a 99.95% availability (**compared to Egenera’s 99.99%**), but when you look at their verbiage they confuse matters by saying “but equal to or greater than 99.0%.”

³⁸ <https://aws.amazon.com/ec2/sla/>

*“We will apply any Service Credits **only against future Amazon EC2 or Amazon EBS payments** otherwise due from you. At our discretion, we may issue the Service Credit to the credit card you used to pay for the billing cycle in which the Unavailability occurred. Service Credits will not entitle you to any refund or other payment from AWS.”*

Amazon plays the “region card”

It is very important to keep in mind that AWS refers to multiple datacenters within close proximity as a “region”. The individual datacenters (located within the regions) are referred to as “availability zones”. Amazon currently operates 13 regions, and within those regions there are 35 discrete datacenters.

- For example: the AWS Northern Virginia region consists of 5 availability zones (datacenters)

The key point here is that **“Availability Zones consist of one or more discrete data centers ... housed in separate facilities.”**³⁹

Here are a couple of FAQs that highlight how Amazon plays the “region” card when it comes to SLA violations:⁴⁰

Service Level Agreement (SLA)

Q. What does your Amazon EC2 Service Level Agreement guarantee?

*Our SLA guarantees a Monthly Uptime Percentage of at least 99.95% for Amazon EC2 and Amazon EBS **within a Region**.*

Q. How do I know if I qualify for a SLA Service Credit?

*“You are eligible for a SLA credit for either Amazon EC2 or Amazon EBS (whichever was Unavailable, or both if both were Unavailable) **if the Region that you are operating in** has an Monthly Uptime Percentage of less than 99.95% during any monthly billing cycle. For full details on all of the terms and conditions of the SLA, as well as details on how to submit a claim.”*

So Amazon uses “region” availability metrics for its SLA benchmark. This is worthless to you since your client’s instances are probably hosted in a separate availability zone (datacenter).

Now, lets take it a step further and see how Amazon’s EC2 SLA defines “unavailable.”

Definitions⁴¹

- *“Monthly Uptime Percentage” is calculated by subtracting from 100% the percentage of minutes during the month in which Amazon EC2 or Amazon EBS, as applicable, was in the state of **“Region Unavailable.”** Monthly Uptime Percentage measurements exclude downtime resulting directly or indirectly from any Amazon EC2 SLA Exclusion (defined below).*
- *“Region Unavailable” and “Region Unavailability” mean that **more than one Availability Zone in which you are running an instance, within the same Region, is “Unavailable” to you.***
- *“Unavailable” and “Unavailability” mean:*
 - *For Amazon EC2, when **all of your running instances** have no external connectivity.*
 - *For Amazon EBS, when **all of your attached volumes** perform zero read write IO, with **pending IO in the queue.***

³⁹ <https://aws.amazon.com/about-aws/global-infrastructure/>

⁴⁰ <https://aws.amazon.com/ec2/faqs/>

⁴¹ <https://aws.amazon.com/ec2/sla/>

Most large CSPs have similarly complex SLAs that were written in the same “tone”. The point here is that although a CSP such as AWS advertises a 99.95% SLA, you need to read the fine print to understand what it entails. As shown above, AWS’ definition of unavailable is probably different than what AWS’ partners (and their clients) define as unavailable.

Simply put, Amazon’s definitions of “Unavailable” and “Unavailability” can put you in an unenviable position with your clients. Your clients view you as the provider of their cloud server, not Amazon. Think about this as you review the scenarios shown below.

- Scenario 1:** A client informs you that one of their instances (servers) appears to be down because they have no access to it. You perform problem determination steps and arrive at the same conclusion. But, according to the EC2 SLA there are no “unavailability” issues because the SLA states that “all of your running instances” must have no external connectivity. How do you explain this to your client?
- Scenario 2:** A client informs you that something is wrong with their instance (server) because transactions don’t appear to be updating their records (stored on an EBS volume). You perform problem determination on your EBS volumes and conclude that one of the volumes is having problems. But, per the EC2 SLA, an EBS volume is unavailable only when all your attached volumes perform zero read write IO, with pending IO in the queue. How do you explain to your client that per the AWS SLA, everything is fine?

Keep in mind that AWS instances typically use EBS or S3 for storage purposes. Amazon does not offer an SLA for EBS storage. However, as we saw in the S3 SLA (above), EBS storage volumes are not considered unavailable unless all your attached volumes are experiencing issues.

S3 storage does have a separate SLA. Keep in mind that if a server’s S3 or EBS storage isn’t available then for all intent and purposes your client’s instance isn’t available. Right?

Looking at the S3 SLA⁴² we see completely different uptime percentages:

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 99.0% but less than 99.9%	10%
Less than 99.0%	25%

AWS continues with its complexity theme, by having different uptime percentages for Amazon S3 Standard – Infrequent Access (Standard-IA) storage:

Monthly Uptime Percentage	Service Credit Percentage
Equal to or greater than 98.0% but less than 99.0%	10%
Less than 98.0%	25%

Credit Request and Payment Procedures

Even if you experience an outage that complies with the various AWS requirements, just how easy does Amazon makes it for you to receive compensation for such outages?

- The process starts with you submitting a claim to AWS by opening a case in the AWS Support Center. But, if we look back at the AWS Support Costs section, we see that the basic support

⁴² <http://aws.amazon.com/s3/sla/>

level does not allow you to open a case. Therefore, to even satisfy step 1 in the credit request process you must subscribe to a (fee based) support level that allows you to open a case.

- o At Egenera, every partner receives the same level of support, we don't charge our partners for support, and anyone can open a case.
- Your credit request must include the dates and times of each incident

And last, but not least:

- You must include “your request logs that document the errors and corroborate your claimed outage.”

Once again, when it comes to the tone that the AWS customer agreement and SLAs were written in, whose interest is the top priority?

Summary

This paper was written with the goal of equipping you with the information you need to make a wise, prudent, and informed decision on who you should partner with for your journey into the cloud. This paper specifically looked deep into the inner costs of using Amazon's Web Services. As is the case with most large CSPs we uncovered a lot of hidden costs.

Many of our partners came to us after experiencing the TCODB shock with a large CSP. They found us by searching for a wholesale cloud services provider that met some critical criteria:

- Experience with cloud technologies
- 24x7x365 partner support @ no additional charge
- Cost efficiency across all aspects of cloud services including infrastructure, support, and partner programs
- Reliable, scalable enterprise-class hardware hosted in world-class, secure datacenters
- Intuitive, full-service partner cloud services portal
- An SLA that was written to protect the partner, ensuring a high level of cloud services availability

Selecting the right cloud partner enables you to evolve your business model to include higher value, higher margin services on top of the reliable "heavy lifting" cloud infrastructure services provided by your CSP. With higher value services comes stronger client loyalty (stickiness) which in the end, gives them pause when it comes to looking elsewhere.



Although trying to determine the final TCODB with any large CSP is difficult (it's akin to trying to grab a handful of Jell-O®) the main goal of this, and the previous papers is to help you understand how important it is to look beyond the infrastructure (e.g. server instance) costs.

You need to look under the surface to determine:

- What the real costs are for the 24x7x365 support you need from your CSP
- What the real costs and complexities are for modifying a configuration to meet your client's changeable workload requirements
- What the real costs are for participating in the CSPs partner program
- If there are data transfer costs for moving data in and out of the cloud
- If there are time-sink (time is money) costs imposed on you by the CSPs complexity

The cloud is here, and now's the time to start or realign your journey. Just be sure you're working with a cloud services partner, not just a cloud services provider. A cloud services partner aligns their success with yours.

About Egenera

Egenera was founded in 2000 to deliver “tomorrow’s datacenter – today.” Since then we’ve been providing highly reliable, advanced server virtualization/converged infrastructure solutions. Building on this experience, we developed Xterity Cloud Services specifically for the channel. Xterity Cloud Services provides an intuitive, drag-and-drop cloud management solution that simplifies all cloud workflow processes including customized design, deployment, scalability, management, pricing, margin analysis, and billing.

Without the inherent complexities found in other cloud services, Xterity partners are liberated to focus on providing higher value, higher margin services to their clients. Xterity Cloud Services combine the security and reliability of Tier 3 datacenters with enterprise-class hardware and software to deliver wholesale managed private and public cloud IaaS services including bare metal, disaster recovery, backup, and migration. 24x7x365 partner support and assigned account management is included at no additional cost to help ensure success for our partners and their clients.

To learn more about Egenera’s Xterity Cloud Services please contact us at:

USA - Headquarters

Egenera, Inc.
80 Central Street
Boxborough, MA 01719 USA
phone: 978.206.6300
fax: 978.206.6436
web: www.egenera.com
email: info@egenera.com

Asia/Pacific

Egenera, K.K.
Kojimachi Crystal City
West Bldg. 4-8-1 Kojimachi
Chiyoda-ku, Tokyo 102-0083 Japan
Phone: +81-3-6261-6301
web: <http://japan.egenera.com>
email: info-jp@egenera.com

EMEA - Dublin

Egenera Ltd.
4033 Citywest Avenue,
Citywest Business Park,
Dublin 24
Phone: +353 1 485 3473
email: sales@xteritycloud.ie
email: info@xteritycloud.ie

EMEA - London

Egenera Ltd.
21 St. Thomas Street
Bristol, BS1 6JS
United Kingdom
phone: +44 (0)203 808 5563
email: emea@egenera.com

EMEA - Germany

Egenera GmbH
Dornhofstrasse 34
63263 Neu Isenburg, Germany
phone: +49 6102 812230
email: emea@egenera.com