

# A Business Perspective on Deploying ArcGIS for Server in the Cloud

This session will cover models and platforms for deploying ArcGIS for Server in the cloud from a business perspective. We will help you understand how to think about costs when deploying ArcGIS for Server in cloud infrastructures, private clouds, etc.

<http://video.esri.com/watch/67/a-business-perspective-on-deploying-arcgis-for-server-in-the-cloud>

---

## Video Transcription

**00:01** Thanks for attending this afternoon.

**00:02** What we're going to talk about is a business perspective on deploying into the Amazon infrastructure.

**00:10** Myself, my name is Andrew Hendrickson. I'm a solution architect within our corporate sales management division.

**00:15** I'm going to be copresenting with Marwa Mabrouk from our service implementation services division, correct?

**00:21** And so, I will pick up and start the conversation today in discussing, you know, is the Amazon cloud right for your GIS?

**00:31** And I'll talk about a lot of different things from a business angle, but we'll also...

**00:36** ...dive a little bit in and talk about architecture as well.

**00:39** So I think it'll be necessary to talk a little bit about patterns and practices as well.

**00:44** We'll talk about business cases for using the Amazon cloud.

**00:48** Within that, we'll kind of focus on elasticity and time to market...

**00:51** ...as well as risk aversion and budgetary reasons for utilizing the Amazon infrastructure.

**00:57** We'll talk about ArcGIS 10 a little bit and we'll float a little bit back and forth and talk about some technical...

**01:04** ...we'll deep dive a little bit for some of the deployment options that I'll discuss.

**01:10** And then we'll wrap up talking about some services that we have to help get you started.

**01:19** So with ArcGIS 10 we see a change.

**01:24** You know, we see a cloud-ready system that out of the box supports the ability to have...

**01:30** ...both an on-premise server as well as a hosted server.

**01:34** In addition to our desktop, mobile, and rich Internet clients, we now have the ability to have multiple, or I should say...

**01:40** ...sort of a hybrid approach to architecting a solution to satisfy your business needs.

**01:47** So what becomes important now is to start to talk about, how does GIS fit into your IT enterprise?

**01:54** What is the governance around the GIS itself? Is GIS part of it?

**02:01** Is it part of the IT, you know, enterprise plan?

**02:06** Do we consider it mission critical? Right? 'Cause we've got to start talking about SLAs and service-level agreements...

**02:12** ...when we're talking about including the GIS on infrastructure that you do not necessarily own.

**02:21** So in a sense, are you in or out of the IT governance?

**02:25** So what this slide attempts to describe, is cloud an option for your GIS?

**02:30** In the left column I have listed as tiers, and I've just sort of named them, randomly as...

**02:37** ...SLA 1 through SLA 3, meaning service-level agreement again.

**02:42** Let me define service-level agreement.

**02:43** What I'm talking about is what is the required uptime...

**02:46** ...or expectations of that system match with your business workflow needs?

**02:52** So the way I've laid this out is, in the second column it says Uptime.

**02:56** What is the expected uptime of a particular application or technical workflow?

**03:01** Is there five nines? Is there four nines? Is there three nines?

**03:03** This is what I'm really getting at in terms of uptime. So 99.999 percent uptime.

**03:09** Does anybody have that type of uptime with their GIS today? Okay, couple, right?

**03:14** Let's move down a little bit and talk about 95 percent uptime...

**03:17** ...which basically means you get about 18 days a year that you can be down.

**03:20** Does that sound more realistic for the GIS domain?

**03:24** This is something that we need to begin to consider because the type of infrastructure that you will use to support it...

**03:30** ...will either enable you or disable you from doing that.

**03:33** So let's go from left to right on SLA 1.

**03:37** SLA 1, the uptime there required, 99.99 percent. 999 percent, which gives you 5.25 minutes of downtime a year.

**03:48** Some folks are laughing, which translates into 25.9 seconds a month...which translate into about 6 seconds a week.

**03:58** So this is a...this is a sample, by the way.

**04:02** This is just how I want to start to get you thinking.

**04:04** On the right-hand side I have listed under Infrastructure, On Premise.

**04:09** Why do I have listed On Premise?

**04:10** Because there may be the inability for a cloud vendor to support that type of uptime, depending upon what you choose to use.

**04:19** But you need to be aware of this ahead of architecting a solution and including the cloud in your potential conceptual architecture.

**04:28** Number two, the second column there, SLA 2. Uptime requirement, 95 percent.

**04:33** Well that gives you about 18 days a year that you can be down, or 36 hours a month, about 8 hours a week.

**04:39** Down meaning not available to the rest of your business.

**04:42** The infrastructure I've listed here is a combination or a hybrid of on-premise and cloud...

**04:48** ...where we're able to mix and match the SLAs

**04:51** ...on the particular pieces of infrastructure to match your business needs.

**04:57** SLA 3, 90 percent required uptime. 365 days a year.

**05:02** Again, I'm recommending in the third column the same type of architectural pattern for deployment to meet those needs.

**05:10** So the takeaway from this slide is, we need to begin to think about the GIS the same way...

**05:14** ...we think about other critical mission, or mission-critical systems.

**05:19** A good way to do this is to start thinking about the technology as it falls into deployment patterns, or patterns of usage.

**05:30** So over the past few years, we've been working a lot with large clients and small clients alike...

**05:35** ...and we've noticed trends in the way folks are deploying Esri technology.

**05:39** The first pattern that we want to talk about is data management.

**05:43** So where is the data being managed?

**05:45** Sometimes this is also referred to as asset management.

**05:49** But this directly correlates with how you're collecting, organizing, and exchanging your data.

**05:53** collect once, reuse many times.

**05:56** So from a technology perspective this correlates with our geodatabase, our geodatabase information model.

**06:03** Moving to the right, Planning and Analysis.

**06:05** So here's where we're able to take information out of that place where we're storing it...

**06:09** ...and transform it into actionable information.

**06:13** We do that how? Contemporarily, we do that with our geoprocessing framework.

**06:18** There's other ways that it can be done too, but primarily we're talking about geoprocessing.

**06:23** Further to the right, Field Mobility.

**06:24** Getting information into and out of the field.

**06:28** Having a field user be able to actually alter attributes in the geodatabase...

**06:33** ...or change features and submit them back in real time.

**06:37** The fourth pattern, Operational Awareness.

**06:40** Here's where we're able to take information and make it available to the rest of the business...

**06:43** ...whether it be via a web map, or a mashup, right? But this correlates with our web API technology.

**06:51** Why do I like using these patterns when I'm talking about the cloud as well?

**06:54** Because the technology performs differently, depending upon which pattern of usage you're using.

**07:02** So it's a good way to start thinking about your business process workflows...

**07:06** ...and which ones might be better served in the cloud versus on premises.

**07:13** So one neat thing that we've seen happen over the past few years is the usage of our web APIs is increasing tremendously.

**07:21** So it's giving, relevancy to the GIS domain.

**07:24** The folks that are editing data are actually doing things within the geodatabase editing features, but what do we have?

**07:31** We have a lot of demand coming from folks who want to see those web maps.

**07:35** They want to see the GIS data taken out of the domain and mashed up with other types of map services.

**07:42** Utilizing the Amazon EC2 paradigm, we can very quickly expand out a web mapping application.

**07:50** And we can also bring it back.

**07:51** So if we've got peaks and troughs in our usage, we can expand and bring them out...so...and bring them back.

**07:57** So the concept I wanted to leave you with is let's begin to think about the patterns of usage...

**08:01** ...within your business enterprise as well, before you make a decision as to what way you want to deploy the technology.

**08:08** Is cloud actually an option for you?

**08:16** So I thought I'd throw up a little quote, definition, you know. So what is the definition of "elasticity?"

**08:21** "The tendency of a body to return to its original shape after it has been stretched or compressed."

**08:28** Is elasticity a reason, by a show of hands, for why a lot of you were curious about utilizing Amazon?

**08:35** Maybe? Okay.

**08:37** I challenge that in the future it may be, for you.

**08:40** I may see more hands in the future.

**08:46** So I'm going to dig in a little bit on elasticity.

**08:49** You can adjust for peaks and demands, troughs, for your data management tasks that you're performing today.

**08:56** So, with our potential deployment options today...

**09:01** ...we see the availability to create high availability with ArcGIS Server in the Amazon cloud...

**09:07** ...as well as the enterprise geodatabase.

**09:10** We see the availability to actually deploy a highly available infrastructure...

**09:18** ...that you can do updates from out of the cloud, into the cloud.

**09:23** But let me clarify that, in business, the cloud is almost like you're renting that server hardware, right?

**09:28** You don't own it. Right? There's implications there too.

**09:31** So as we think about this from a budgetary perspective, as we think about these deployment patterns...

**09:37** ...think about the implications for your budget as well, right?

**09:41** Is it hard for folks to provision, hardware?

**09:45** I have found in my career it could be very difficult to provision hardware...

**09:49** ...because it comes out of a certain style of budget...

**09:51** ...whereas this may not come out of the same style of budget, so I'll come back to that.

**09:55** Okay, enhanced and dynamic processing.

**09:57** Large batch processing, geoprocessing large-scale analysis.

**10:02** Being able to blow out that infrastructure and then bring it back without taking ownership...

**10:06** ...or having to worry about actually provisioning all of that hardware.

**10:12** Some things we've seen going on already. Cache cooking. Do we like this term of "cache cooking?"

**10:18** So creating your cache and then moving it around.

**10:21** Something you have to do not so often, depending upon how often your data changes.

**10:26** But maybe you want to use some infrastructure that you don't own to actually create it...

**10:30** ...or store it, and then bring it back.

**10:35** So actually deploying your cache in S3 as well.

**10:38** And we'll talk about precisely what S3 is and what the price implications are later in the presentation.

**10:43** Large batch geocoding, or just geocoding in and of itself...

**10:47** ...do you want to support that on your own infrastructure or sort of farm it out?

**10:54** So the operational awareness pattern, in mobile as well. I kind of lumped these together on this slide.

**10:59** So growing out capacity as needed. This has been a barrier for many clients in the past.

**11:06** Or, simply, we didn't know how many folks were going to be utilizing web maps...

**11:12** ...or mobile capabilities, so we put something out there and said people will come to it.

**11:16** Well sometimes I say be careful what you're asking for.

**11:19** Put something out there and then all of a sudden on day two you've got 1,200 people coming at it at the same time, right?

**11:24** So how do you support that without sinking your infrastructure?

**11:28** This type of deployment, the Amazon cloud gives you very, very nice options here.

**11:34** So expanding your capabilities in near real time.

**11:38** Hopefully some of you were able to attend the earlier session...

**11:40** ...were able to see that it does not take a lot of time to provision a server.

**11:45** So one server access with ArcGIS Server, high availability. This is a big deal.

**11:52** We're going to start to see folks that in the past have not been able to actually deploy in...

**11:57** ...highly available environments, be able to utilize these deployment patterns to do so now...

**12:01** ...which means you can begin to support more mission-critical workflows.

**12:07** Same thing between the mobile pattern of usage as well as the operational awareness pattern of usage.

**12:15** So I wanted to throw some pretty simple architectural style diagrams at you today as well.

**12:21** What we're looking at is a traditional on-premise deployment, where we've got a single user, or a couple users.

**12:26** I don't mean to denote it's only one person, but basically there's a desktop...

**12:30** ...there's an ArcGIS Server, there's an editing environment on premise.

**12:35** On premise meaning you own the server. You've got to provision the server, right?

**12:39** You own it, right? You've got to administer it.

**12:47** So this one shows a cloud deployment of the same architectural pattern...

**12:52** ...where the user is accessing the same pieces of Esri technology...

**12:58** ...but now they don't own that physical piece of hardware anymore.

**13:01** It's actually on Amazon's infrastructure.

**13:04** The desktop can still exist there. ArcGIS Server can still exist there and you can do some editing in the cloud.

**13:11** You don't have to own any of it. This is one deployment pattern that we're seeing today.

**13:18** So expanding that out a little bit, let's talk about actually adding some redundancy to it, right?

**13:25** So here we've got an on-premise person, I'll just keep referring to it as a single person...

**13:31** ...but on-premise data management and analysis occurring, where you're still managing assets and you're doing geoprocessing.

**13:37** You're adding value to the data on premise, but you're pushing across and up into the Amazon cloud...

**13:44** ...information that you can serve as a web map.

**13:47** So you're not coupled to the infrastructure that you have on premise to serve a greater audience in the cloud.

**13:56** And we can utilize many, many different parts of the Amazon technology platform to manage growing this out.

**14:03** And we'll talk about that later in the presentation as well.

**14:07** So this is listed as ArcGIS Server 1 through N.

**14:11** This is a very, very nice architectural approach as well for decoupling your different environments.

**14:17** There is an SLA that's different for the operational awareness pattern...

**14:21** ...than it is for your data editing environment in this proposed conceptual architecture. The two are not connected.

**14:33** In this slide, we're demonstrating a bit of a hybrid, where we've got, an on-premise user.

**14:39** They may have an actual desktop there or maybe they're just tunneling in to Amazon...

**14:46** ...and accessing ArcGIS Desktop there, with ArcGIS Server there as well.

**14:51** And we can have the actual infrastructure supporting data management and visualization decoupled as well...

**14:58** ...with them both sitting in the cloud.

**15:08** We can synchronize between geodatabases in the cloud as well, with multiple ones existing.

**15:15** This has a lot of implications in terms of being able to support a large amount of users...

**15:21** ...for being able to actually synchronize a highly available geodatabase on Amazon infrastructure.

**15:28** This has typically been a pain point in the GIS domain for achieving maximum uptime of a redundant geodatabase...

**15:36** ...while trying to support the actual, the physical device itself.

**15:42** Hope this is becoming clear how it's playing out.

**15:44** There is a bit of...a degree of separation that we want to have for the types of patterns that you're using the software...

**15:51** ...and we want to correlate those with how you're deploying it.

**15:53** We want to see you correlate it with how you're deploying it.

**16:02** So in this one, you know we've got a publication instance actually in the cloud, editing actually going on on-premise...

**16:07** ...and we're synchronizing between the two and simply creating an active passive pair...

**16:13** ...where we don't require that high level of redundancy between the two databases.

**16:20** So again, let me rewind us a little bit and say it's very important for you and us to understand...

**16:26** ...how available the data needs to be to an application or an end user...

**16:31** ...before you choose your type of deployment pattern.

**16:35** And when does the business say that the data actually needs to be live?

**16:43** So in this case, you'll notice that the operational awareness pattern...

**16:46** ...which is the one all the way to the right where we're serving web maps, hasn't changed at all.

**16:50** It's completely different. So they are decoupled even though they're all in the cloud.

**16:59** With this slide, it's broken down even differently.

**17:03** Another message I think I should state for you is that there's a lot of different options now...

**17:10** ...and the ArcGIS platform is pervasive across all of it.

**17:14** So you have the ability to, even if you cannot have because of some sort of IT governance...

**17:20** ...where your organization does not allow things to exist outside of your data centers...

**17:25** ...or outside of your own owned hardware...

**17:28** ...maybe you might want to just have a development environment that exists on Amazon...

**17:32** ...whereas in the past you didn't have a separate development environment.

**17:36** You can build a development environment out very quickly and get rid of it when you don't need it.

**17:42** It's also very good for a QC/QA environment.

**17:45** This one is demonstrating a completely on-premise deployment, but only with a test environment in the Amazon cloud.

**17:54** So as you can imagine, there is a myriad of ways that we can mix and match the way that we deploy ArcGIS.

**18:04** And the way that you utilize data management is also very important.

**18:08** As Marwa will talk about later, there's implications for the size of data that you're going to be moving around.

**18:14** So if you don't need to move a certain amount of data from on-premise up, don't.

**18:21** 'Cause it may cost you more money.

**18:25** We'll come back to a pattern later in the presentation as well.

**18:29** So time to market. Here's another business reason that I wanted to discuss with you today.

**18:34** So getting a server provisioned quickly, getting an application up and running quickly.

**18:39** There's a bigger issue here too, erasing the limits of creativity with the ability to quickly respond to business needs.

**18:48** We see this with a lot of clients today.

**18:52** We see this in emergency response, where you quickly have to create an application to respond to some sort of need.

**18:59** So utilizing the cloud, the Amazon cloud, tremendously, tremendously increases your time to provision hardware.

**19:07** So rapid provision of ArcGIS Server equals less time to spin up your servers.

**19:13** They're easy to set up and administer.

**19:19** You can release a web mapping application extremely quickly, in under an hour.

**19:31** This completely removes restraints on innovation. Very exciting to me.

**19:36** It gives you the ability to test different deployment patterns all the while you're doing this, right?

**19:41** Maybe something isn't working for you...or the organization.

**19:47** It allows us to be simple as well.

**19:49** I love staying as close to off-the-shelf software as I can so I can switch between versions, I can upgrade...

**19:58** ...and allows me to scale very quickly, so if I do put something out there quickly and it's not correct...

**20:04** ...I can go back. I can rewind and redo it.

**20:13** Another concept I want to talk about, risk aversion.

**20:16** Large organizations, small organizations, any organization is typically going to be risk averse, right?

**20:25** You know, as you grow technology throughout your business and you meet and serve more needs...

**20:31** ...you should be risk adverse. You shouldn't just be spending a lot of money and dumping it out there.

**20:36** So we shouldn't have to limit ourselves due to cost ceilings as well.

**20:43** So as we're building things out and we make a mistake, it's good to be able to bring them back quickly.

**20:53** So what is the cloud appropriate for?

**20:56** There's some best practices that we've been talking about today a lot.

**21:00** I haven't called them out necessarily as best practices but they are.

**21:03** High availability, redundancy.

**21:06** What is it appropriate for, you know?

**21:08** Is it appropriate for your development environment?

**21:11** Is it appropriate for your staging environment?

**21:13** Is it appropriate for production?

**21:16** That depends upon your organization, right?

**21:20** By being able to do performance and scalability tests ahead of actually releasing things, we limit our risk.

**21:26** Maybe the cloud isn't reliable enough to meet your needs.

**21:31** This goes back to that service-level agreement stuff we talked about. Is it secure enough, you know?

**21:38** One of the main reasons why we see folks considering deployments to the cloud...

**21:44** ...is to maximize CPU utilization across infrastructure, right?

**21:49** Why do I have this under Limit Risk?

**21:52** I kind of toyed around with this one for a little while, because I have in the past bought...

**21:58** ...or I should say spent a lot of money on hardware, and then when we went back and evaluated it...

**22:02** ...we weren't actually utilizing as much CPU as we needed to across that hardware.

**22:07** So when I was reevaluated and I was going back for budget again...

**22:10** ...I had a bit of a problem justifying those big, honkin' machines that I had bought.

**22:16** Utilizing cloud deployments. We can limit that risk tremendously.

**22:21** 'Cause you can spread things out much differently than you can with that one big honkin' box that I just referred to.

**22:27** That's what I mean by real utilization of owned property, of owned equipment versus rented.

**22:32** Again, sort of touching on the concept of budgetary concerns.

**22:37** Compliance as well. This is interesting.

**22:39** You're able to comply with certain rules and regulations by being green in terms of computing environments utilizing the cloud.

[22:49](#) I'll take questions at the end.

[22:52](#) So budget. Growing the GIS into your business enterprise requires unique budgetary planning.

[23:02](#) We always need to limit capital expenditures, right?

[23:06](#) Especially nowadays. Flexible expense budgets can be used. What do I mean there?

[23:11](#) Can you use your expense budget for purchasing hardware?

[23:16](#) Ah! But are you purchasing hardware in the Amazon cloud or are you purchasing a service?

[23:22](#) You might be able to utilize different budgets to grow out capabilities within your enterprise.

[23:29](#) So reducing operational costs as well. No heavy lifting. This is interesting. I simply mean no heavy lifting.

[23:35](#) You're not loading servers into racks. You're not worrying about power, space...

[23:38](#) ...air conditioning, you know, any of this sort of stuff.

[23:41](#) Somebody else is dealing with that. That's built in to the cost of the actual service itself.

[23:46](#) And again, this time-to-market effect is very important.

[23:50](#) So literal cost and practice. We're going to talk about the literal cost and practice.

[23:58](#) We're going to talk about what an A-M-I is, or an AMI.

[24:01](#) We're going to talk about the cost of use for very specific scenarios. This could vary greatly.

[24:11](#) And we'll talk a little bit about billing and how it works as well.

[24:15](#) So I think at this point I'll turn it over for a couple of minutes.

[24:23](#) You might want to come over here.

[24:31](#) Sorry. So this section I'm going to be talking about how to estimate your costs.

[24:37](#) Some of the concepts have to do with the actual services that Amazon offers.

[24:41](#) We had a session earlier today where we went into those services, explained them in a lot of detail...

[24:48](#) ...so I'm going to do this a little bit briefly, while also giving you a good idea about what these services can do...

[24:56](#) ...and then how to understand how you're going to get billed for them.

[25:01](#) The...one of the most important concepts about the Amazon cloud is the instance types.

[25:07](#) What Amazon does is that they prepackage their computing power and their memory capacity in a preset configuration.

**25:17** This is maybe...

**25:20** So, for example, if you want to use a processor dual core versus quad core...

**25:26** ...you can't set this according to what you need at that moment.

**25:30** You're going to pick out of the list, and that list will contain a predetermined set of capacity configurations...

**25:37** ...and you will go with one of them. And depending on what you choose, you're going to be billed a certain rate per hour.

**25:43** So I put here the sizes that are typically published by Amazon.

**25:47** They have been modifying them and they've been known to add to them pretty quickly.

**25:52** The ones that are out there right now are the category standard, high CPU, and high memory...

**25:56** ...and from the name you can tell what has been the drive for that category.

**26:02** They want to offer machines that either have more CPU power or more memory capacity.

**26:08** I've also included a couple of examples.

**26:11** For the 64 bit, the example I have on the left side includes the smallest machine that they have...

**26:16** ...which is basically around 7 gigabytes of memory and it's a dual core with 4 EC2.

**26:25** And the way Amazon would do the computing unit, it's extracted to this term called EC2, which is an Amazon capacity...

**26:34** ...and it equates to around 1 gigahertz of speed for the processor.

**26:38** So their processors will vary depending on the different categories and depending on that number of EC2s.

**26:45** So in short, this means this is a dual-core machine that has pretty reasonable speed...

**26:51** ...not very fast processor, and it's only got 7 gigabytes of memory. It's 64 bit.

**26:57** Now on the other high end, they offer other machines that seem to be a little bigger.

**27:02** So, the high CPU extra large has more emphasis on the CPU power, so the CPU tends to be a little bigger or a little bit faster.

**27:12** They also don't have that much memory in it.

**27:15** One of their most recent releases was the high-memory category...

**27:20** ...and they have the largest and it's called quadruple extra large.

**27:24** They typically kind of rank them as small, medium, large, extra large.

**27:28** So this is the biggest of all the instances you can get...

**27:32** ...and it's got about 68 gigabytes of memory and what equates to 8 virtual cores.

**27:40** Now each of these instances, they come with different sizes and they also come with different prices.

**27:47** I included here another example for the standard type. This is the smallest.

**27:51** The smallest is around 48 cents an hour and that is what you will be paying if you run that much CPU and that much memory.

**28:02** It also comes in extra large. The price would be around 96 cents per hour.

**28:08** This is what's called the standard on-demand.

**28:10** So this is the pricing if you want to not sign up for something that is long term.

**28:15** You just want to use it for a number of hours and then you don't want to be paying again...

**28:21** ...and then when you feel like using it, you go back and start some other instance.

**28:24** If you...if you don't want to do that, if you know you're instance is going to be running all the time.

**28:30** So for this coming year, I am planning two instances that I'm just going to have running all the time.

**28:36** Now you can get something similar to like a bulk discount and that's what the reserved instances are for.

**28:43** You can see you would pay a certain year term that includes an up-front cost...

**28:47** ...but after that the usage rate per hour goes considerably low. It can go to savings up to 30 percent...

**28:55** ...so what most people would do is that for the instances they know they're going to be running all the time, they will use the reserve.

**29:02** And then for the other instances that they would like to use the elasticity model for, they would use the on demand.

**29:08** So that's how this model really works.

**29:11** And you can see considerable cost savings when you take into account how the reserve will be used and how the on demand.

**29:22** So to run an actual machine, if you want to compare a machine in the cloud to having it on premise...

**29:29** ...what are the costs to running a machine in the cloud?

**29:32** A lot of folks go to the Amazon website and they're not really sure how to figure out...

**29:37** ...which service is what and how things work out.

**29:40** As we just walked through the computing power in terms of CPU, and the memory has it's own hourly rate.

**29:47** There's also the rates that are related to the storage and typically the storage goes into the elastic block storage.

**29:55** So the elastic block storage is like an external hard disk that goes with your machine...

**30:00** ...so apart from the actual C drive on the machine...

**30:03** ...you're going to be paying for every gigabyte that you would like to attach to that machine.

**30:09** So this storage can vary quite widely.

**30:12** You might have a dataset that is just 5 gigabytes and then you're just paying for 5 gigabytes...

**30:18** ...or you can have a dataset that goes up to 2 terabytes or more.

**30:21** And then you're going to be having multiple EBS blocks, each contains different datasets...

**30:27** ...and all of them are attached to the same instance.

**30:30** And if you have that...the same dataset attached to multiple instances...

**30:35** ...you would paying for that times the number of instances.

**30:40** And I'll have a later on example that'll explain that at a little bit more detail.

**30:44** But for the time being, if you're trying to look at the costs of just one machine in the cloud...

**30:50** ...you have to consider the hourly rate that goes with the instance type that you've chosen...

**30:54** ...the storage capacity that you have attached to that machine...

**30:59** ...in addition to the network costs, which is expressed as data transfer in and out.

**31:04** I've listed here just some example for how the costs work.

**31:08** For the data transfer, Amazon has a promotion, so for the time being it's free...

**31:14** ...but what we understand is that it's not going to be free forever.

**31:18** As for the data transfer out, first gigabyte per month is free...

**31:23** ...but after that the first 10 terabytes are 15 cents per gigabyte...

**31:28** ...and after that as your load increases, the prices start going down.

**31:33** But that kind of gives you an idea that if you are going to really equate a machine running in your environment...

**31:39** ...and what you would get in the Amazon bill, these are the three things to take into account.

**31:45** And if you understand how long you're going to be running your instance, the storage capacity...

**31:51** ...how many users are going to be sending requests using the network...

**31:56** ...you can estimate the cost that you will see on your monthly bill.

**32:04** Speaking of instance, there's a very important concept just in case this is the first session you attend for the cloud...

**32:11** ...the AMIs are one of the important concepts that go hand in hand with using the cloud.

**32:17** It's a machine image, so once you've created an instance, it has the software...

**32:22** ...it has the data, it has everything you need, you create an image.

**32:26** And the image in the Amazon cloud is called an Amazon Machine Image.

**32:31** Esri has been working on creating AMIs that include our software, so we have two available.

**32:37** One with ArcGIS Server 10 and another one with enterprise geodatabase...

**32:42** ...and if you start from these AMIs and create your own by launching an instance from this one...

**32:48** ...and then setting it up with your data, then saving that as your own custom AMI...

**32:54** ...whenever you use that custom AMI, you have a ready-to-use machine...

**32:58** ...not just with the software, but with your data and with your applications as well.

**33:06** Other services and I'm going to come back and see how they relate to the patterns that Andrew's been talking about earlier...

**33:14** ...but just to understand, these are very important basic services and that's really what makes...

**33:19** ...these are the bits and pieces that make the cloud solutions, how you use it and how you deploy it.

**33:24** Another service is the Elastic Load Balancer. And this is an important service...

**33:28** ...because it plays the role of a load balancer that does the brokering between different instances.

**33:33** So if you have a distributed environment, you would rely on the load balancer to distribute the load between those different machines.

**33:41** The load balancer doesn't just work as a load balancer between a few machines in one data center.

**33:46** It works with the cloud, so it's elastic enough to go across different zones...

**33:52** ...across different data centers, and it will also grow with the load.

**33:58** So it works, in a sense, a little bit different from what a physical load balancer will do. It's more suitable for the cloud.

**34:04** So if you're doing something in the cloud, this would be a very important service.

**34:09** The way they charge for it is per hour, so you would pay, like 0.025, I think that's like a quarter a cent per hour...

**34:23** ...for using it, just for having it up and running. And then you would pay per gigabyte of data process.

**34:29** So if you've seen the previous prices for the data transfer in and out...

**34:34** ...if you are going to calculate a certain data going in and out throughout...through your instance...

**34:40** ...that would be the same amount of data that would be going through a load balancer as well.

**34:45** And the Amazon CloudWatch is a service provided by Amazon.

**34:49** The purpose of it is to make monitoring easier.

**34:52** So once you turn it on, you can monitor the progress of your machine resources.

**34:59** So it will monitor the processor utilization, the memory utilization, the network utilization...

**35:04** ...and you can see the graphs that express that on hourly rates.

**35:08** That service costs 0.15, I'm not sure, that's like less than a quarter dollar an hour.

**35:17** In general, that would be around \$11.00 per machine per month.

**35:22** So if you have like three machines, that would be 11 times 3 per every month.

**35:27** And it's one of the really handy services as well.

**35:31** Also worth mentioning, the elastic IP.

**35:34** So part of the nature of the cloud, when you start a new machine...

**35:38** ...it will be assigned a new IP address and it will include a new host name...

**35:42** ...and a lot of people typically configure their machines based on a known IP.

**35:48** So if you want to avoid this nature of changeable IP address...

**35:52** ...you can attach an elastic IP to your machine and from there on it will always have that IP.

**35:59** So that's a stable IP. It never changes.

**36:01** Once you create it, it sits there and you can attach it to any machine you start and that will become its IP.

**36:06** So that's the way to control the elastic nature of the cloud if you want to have a tied-down, known configuration.

**36:14** The way they charge for the elastic IP, there's no cost for having it up there...

**36:21** ...but you actually start paying for it if it's not attached to a machine.

**36:26** So as you see, the class model kind of varies widely between the different services.

**36:31** So I want to make sure you kind of understand the basic ones and then you should know that as you...

**36:37** ...kind...discover new services, always check what the price model is like. Don't make assumptions.

**36:43** So this one, you actually pay for it if you're not using it. But if you're using it, you're not going to pay.

**36:49** And the remapping, which is the reassigning.

**36:52** If you keep reassigning it to different machines, based on this reassignment, they will charge you.

**36:57** So the first hundred reassignments to the machine are free...

**37:01** ...but after that you start paying a certain charge for it, which is around 10 cents.

**37:10** One more service to explain really how the elasticity works in the cloud.

**37:15** So everybody hears the words AMIs, which is the machine image...

**37:20** ...and they understand an instance. But how does this elasticity really work?

**37:24** If you want to implement elasticity, the AMIs really play a very important part in it.

**37:28** Because what happens is that you want to automate the ability to see if these machines that I have running...

**37:37** ...for example, I have two instances, and if these two instances are having a CPU utilization higher than 70 percent...

**37:46** ...I would like to go and start another two instances.

**37:49** And you want something to do this automatically for you.

**37:52** That is the elasticity of the cloud. And that is really done through that service Auto Scaling service.

**37:58** So the Auto Scaling service will utilize the CloudWatch to monitor the CPU utilization...

**38:05** ...and you can define in that service that I would like to start two more instances if these ones are utilized that much.

**38:13** And then once these two new instances are up and running, I would like them monitored as well.

**38:18** So if they are not used that much, so if they are...if the processor utilization of the new

instances, the auto scaling just started...

**38:28** ...are used less than 10 percent, I would like the auto scaling to shut them down and go back to just having two running.

**38:35** So this ability to add new instances depending on the utilization can be done programmatically.

**38:42** You don't have to keep monitoring it and then from there do it manually.

**38:47** You can automate the whole aspect using the Auto Scaling service.

**38:51** And that is really how elasticity works in the cloud.

**38:53** It's the ability to use the AMI to create instances when you need them and shut them down when you don't need them.

**39:01** For that service, the way Amazon charges for it, it comes at no extra cost.

**39:05** The cost goes for using the CloudWatch service...

**39:08** ...as well as the time you pay for the instances when they're running, so the actual cost of the instances.

**39:17** Now a word of caution. When you are doing the elasticity, it's always good to understand what you expect...

**39:28** ...or at least do some assumptions around when you expect to have those extra instances running...

**39:32** ...because in the end you're going to get a bill. And that bill is going to charge you for every hour every machine ran.

**39:39** So if you don't know ahead of time that you're going to need to start four instances and keep them running for two weeks...

**39:47** ...it's good to maybe not have them run, but...that's part of the risk aversion.

**39:54** You want them to be running so users will find them, but you should know your bill will be affected.

**40:00** So always kind of have this trade off and know that while you can start the extra instances...

**40:05** ...and it's really great to avert the embarrassment of having your site go down because of extra load...

**40:11** ...always be aware that this is going to reflect on the bill...

**40:13** ...and be prepared to see that variation in the bill when you start those extra instances.

**40:22** So part of the beauty of the cloud is that it's pretty vast.

**40:27** However, the Amazon cloud has kind of enabled the ability to know where your software...

**40:34** ...or where your instances will be running.

**40:36** Because there are a lot of limitations or a lot of requirements that certain data would not leave certain borders.

**40:43** Certain content would be hosted in specific places.

**40:47** So from that perspective, Amazon has identified where the cloud is really running.

**40:53** They have four regions. One in the east coast, mainly in Virginia. One in the west coast here in California.

**40:59** There's one in Europe in Ireland and another one in the Asian Pacific realm. It's in Singapore.

**41:06** So you can pick where you want your machines to be running...

**41:10** ...and you can identify from there that's going to be the geographic location.

**41:14** Now within every region, there's a set of different zones, which represent different data centers.

**41:21** So you can be running all in the east coast and still be geographically redundant.

**41:26** Each of these zones is a completely different data center that sits in a different location within that state.

**41:38** So in general, kind of a recap, there are other services in the cloud and each charges in a different way.

**41:44** But these are maybe the basic ones that touch on the different deployment patterns that Andrew has talked about earlier...

**41:50** ...and I'll come back and kind of use an example to show you how they do.

**41:54** But in general, if you're trying to estimate costs, there are some kind of key rules.

**42:00** One of them is data size really plays a very important part in estimating the cost.

**42:05** Now if your data is not very large, say maybe 5 gigabytes...

**42:11** ...and you don't expect a lot of people are going to be using the services on this machine...

**42:16** ...say for example you have four or five users using that service randomly during the day...

**42:22** ...it's not expected that the majority of the cost is going to be going towards the data transfer.

**42:28** The majority of the cost you're going to see is going to be really in the hourly rate...

**42:33** ...the price to keep the CPU and the memory running for that many hours.

**42:37** The costs for the storage and the costs for the data transfer become pretty much negligible in comparison to the hourly rate.

**42:46** But if your data is very large, if we're talking maybe 1 or 2 terabytes...

**42:52** ...that is going to be a significant cost if you have many users using that data.

**42:57** So then that becomes the part where you really need to come up with certain assumptions...

**43:03** ...around how many users are really going to use that data and how much of that data will be moving around...

**43:10** ...whether through requests that are incoming or outgoing, or just for updates and maintaining it.

**43:16** It would be important then to take into account the life span and the life cycle of the data management...

**43:23** ...due to the fact that it will have a major impact on your cost.

**43:26** It will even probably be higher than the actual hourly rate.

**43:30** So it's always important to understand when you're deploying to the cloud, make certain assumptions.

**43:35** Understand a certain usage model or how the system is going to get deployed and how it's going to get used.

**43:43** And from there, make around these assumptions certain rules that you can use to estimate the cost for...

**43:52** ...each one of the services, to estimate the actual bill that you're going to get.

**44:00** Now to use an example, this is one of the deployment patterns that Andrew's talked about earlier.

**44:07** You can see here that we have multiple ArcGIS servers...

**44:10** ...so we've kind of discussed in detail in the earlier session this morning...

**44:15** ...how the AMIs for ArcGIS Server work and how the instances would work.

**44:21** So to kind of recap on that quickly, using the AMI you would generate one instance, set it up with your own data...

**44:29** ...with your own application, and then create your own custom AMI based on that.

**44:33** And then from that custom AMI, you can start as many instances that, as you launch them...

**44:39** ...they will have the software configured, they will have your data, they will have your application.

**44:44** And then from there, you assign the Elastic Load Balancer to all those different instances.

**44:49** So in the cloud, we're not separating the SOM and SOC in the deployments.

**44:53** The Elastic Load Balancer is the main tool used to distribute the load between the different instances.

**45:00** So from that perspective, you are using here two services.

**45:04** The Elastic Load Balancer, to do the brokering, as well as the instances that will run ArcGIS Server.

**45:11** Now, knowing your services and knowing how many people are going to be using it...

**45:15** ...you can estimate how many of these instances will be needed.

**45:19** And if you want to monitor them, you can use the CloudWatch service and from there you can keep an eye...

**45:25** ...and see if you need extra instances to be started when you need them, you go and start them.

**45:30** And when you don't need them you can shut them down.

**45:38** So this is what an Amazon bill looks like. They usually break down the number of hours from the different services.

**45:49** So in this case, based on a large Windows instance that was running that many hours...

**45:55** ...and of course, within a month it's impossible to run that many hours.

**45:58** This actually represents many instances. So they lumped it all into one item.

**46:05** So you see that many hours, around 2,000 hours by...multiplied by the rate 0.48, we get that cost, \$1,000.00.

**46:15** Then the storage, that's the EBS drive, and over here this is how they break the cost for it.

**46:22** So you could see the gigabytes per month, the number of gigabytes multiplied by the rate, and then the number of requests.

**46:30** These are I/O requests going to disk. And also the backups for this disk. And that's what a snapshot is.

**46:41** And from there, the number of requests to store that snapshot and retrieve it. So that's how the storage costs are.

**46:47** And you can see this is not a very high cost in comparison to the hourly rate...

**46:51** ...and it's because the data we're talking about is not very large.

**46:56** And then the Elastic Load Balancer.

**46:57** So if you have multiple instances and you were trying to distribute the load between them, that's what the service is.

**47:02** And you can see how the cost is broken down for that.

**47:05** And the Amazon CloudWatch, that's the one used for monitoring.

**47:09** So that's kind of the breakdown of the EC2 bill. But this is the actual bill of everything.

**47:16** So this contains everything that Amazon would bill you for. It includes the S3, which is the Simple Storage Service.

**47:24** This is like the big file server in the cloud where you can put your data and once you put it there, it's published.

**47:32** It's also good place to leave your data as a transition while you're setting up your environment.

**47:38** So in this case we haven't used S3 so there's no cost there.

**47:42** But if it is used, as you get access to it with an account, you will see the cost here.

**47:49** Also, the data transfer.

**47:51** You can see it's pretty negligible because there wasn't really much of a load on these instances that were running.

**47:57** But that's what the bill would look like and you can see it up above the title that this was marked for one month.

**48:04** So Amazon bills monthly, whatever you have running in the infrastructure, you will get that bill.

**48:09** And at the beginning, they ask you to give them a credit card to create an account.

**48:14** And put a limit on that credit card just to make sure that if they hit that limit, they notify you before they go and bill more.

**48:22** But what they will do given that mark, they will get the bill every month...

**48:26** ...and they will bill you to your credit card according to that.

**48:33** So a couple more things to take into consideration and going back to the time to market...

**48:40** ...beside just thinking on this as hardware, there's...

**48:44** ...you really need to kind of maybe think of it also as a different way to do things...

**48:48** ...because of the nature of the cloud, to make the best of it...

**48:51** ...maybe you need to learn a few new habits.

**48:57** Something to consider, if we compare the time in the past it took for procurement, hardware setup...

**49:02** ...IT setup, software setup, this has all been approximated in just going and launching an instance.

**49:09** So launching an instance takes about 15 minutes. And once you launched it, it's ready to use.

**49:13** You can start copying your data, copying your application, and in comparison to the number of days...

**49:19** ...or weeks sometimes it took to accomplish all of that, it's become really fast and easy.

**49:25** Next thing to do is to, once it's set up, you can go and create your own custom AMI.

**49:30** So even the time it took to set up your data and application loading...

**49:34** ...next time you launch an instance, that has been approximated. You don't have to do that except once.

**49:39** Once you launch the instance based on your custom AMI, it's got everything ready.

**49:44** So from that perspective, you really need to start thinking about maybe a different way to trouble shoot.

**49:51** You're not really tied to one instance. If you created an instance and it doesn't have a persistent problem...

**49:58** ...just for like, suddenly it stopped working, a lot of the times now it's okay to go and terminate that instance...

**50:05** ...and just start a fresh one that works. It saves you a lot of time from that perspective.

**50:10** Also, kind of the way you do updates. You don't have to take a machine down to do an update.

**50:14** You can go and create another instance, do the update, and then switch them.

**50:18** And when a machine is ready, you can put it behind the load balancer and now everybody can access it.

**50:23** You don't really have to make it available all the time.

**50:26** So, taking these things into account, and how you manage your time, and how you do things...

**50:32** ...you can actually be a lot more effective in very little time and do things very, very quickly.

**50:39** But with great power comes great responsibility.

**50:42** So there are a few things to keep in mind, and one of the most important recommendations I make...

**50:50** ...if you really are going to use the cloud, you really have to follow up on what you start and make sure you terminate it or you stop it.

**50:59** Because a lot of people are in the habit of once you get a machine up and running...

**51:02** ...it's running and then you go and do something else. Even though you don't need it anymore, it's still there.

**51:08** You don't recognize that you are going to get billed for it until you get the bill.

**51:13** So some of the best practices is really to remember. You just use the software and the hardware when you need them.

**51:20** When you don't need them anymore, you stop it. So you stop paying for it as well.

**51:25** And then when you need it again, you start it and you start using it again.

**51:29** So that's really one of the best recommendations to keep in mind.

**51:33** Also, using the AMIs.

**51:35** Like I just showed you, it takes away a lot of the things you used to spend a lot of time on unnecessarily.

**51:41** So now that you have a ready-to-use AMI, you just launch the instance and it's ready to use. It gets everything you need on it.

**51:47** Also, elastic IPs. If you think of a distributed environment that keeps changing its IPs...

**51:52** ...elastic IPs is one of those really good features to take into account...

**51:57** ...so once you've configured your environment, you don't have to keep reconfiguring it.

**52:01** Then having the AMIs with the elastic IPs and the Elastic Load Balancer, everything really works much easily...

**52:08** ...and it's a really easy job to maintain it from there on. Also, should add to automate wisely.

**52:18** So I've seen a lot of folks get very excited about the fact that we have web services in the cloud...

**52:23** ...and now they go and start building scripts to all sorts of very exciting things and very great stuff...

**52:29** ...but then you find that it really wasn't used much. It didn't really save them that much time.

**52:34** So it's always good to consider that there are web services to enable you to do all sorts of things in the cloud...

**52:40** ...but that doesn't mean you should just go and start building scripts unless you really need them.

**52:46** And like I just said, troubleshooting takes a whole different paradigm.

**52:51** So if a problem is not persistent, it's just a fluke or if it's just happening today...

**52:57** ...you don't really have to go and keep troubleshooting it to get that machine up.

**53:01** You could just kill that instance and go and start a new one.

**53:03** And that saves you a lot of time and saves you a lot of effort.

**53:09** Just a quick note, there are a few good references to take into account when you are doing this type of cost calculation.

**53:17** One of them is the Amazon Calculator on their website. It's a very nice tool. Excuse me.

**53:23** And also, keep an eye, because the prices in Amazon do change.

**53:28** Luckily they change to be cheaper most of the time, but they do change a lot.

**53:34** And also the services. They keep providing new services all the time.

**53:38** So maybe they release a new service that you find to be really good for you, works better...

**53:43** ...and then you can migrate to that rather than keep using the one you have.

**53:47** So this is something to keep an eye on.

**53:49** And also because it's a new environment, there's a lot of blogs out there...

**53:52** ...and a lot of people are willing to share their information and expertise through them.

**53:55** So if you're very interested to learn what other people are doing, you will find a lot of information out there.

**54:03** So give this back to Andrew and...

**54:07** I think maybe this will work.

**54:13** So yeah...we also want to make sure that we give you the message that...

**54:19** ...there is support and services today, if you're interested in doing this on your own.

**54:24** I'm sorry, not on your own. If you need assistance.

**54:28** There are certain licensing questions that will inevitably come up at the end of this session when we ask you to...

**54:33** ...give us some questions, and we'll address those, but there aren't easy entry ways into this.

**54:40** There are custom servers that we provide. We've got a jump-start package for you.

**54:46** There's architectural assistance, you know, through your distributor or out of your...through account management.

**54:54** There's a bundle that we're offering and there are experts very close by that can speak to the bundle.

**55:01** And there are premade AMIs that we've been working on for quite some time.

**55:06** And with ArcGIS 10 this is supported out of the box. So with that said, I'm going to do a bit of a review.

**55:14** We're not done, that's for sure, but there's a lot of information that's being thrown around at you today.

**55:19** I mean the reason why we wanted to have a business section, discussion I should say...

**55:24** ...in addition to technical ones, is, I mean, think about what Marwa just went through.

**55:28** Think about the implications for you setting up the wrong way. Right? How much it could potentially cost.

**55:35** So let's rewind a little bit and just think back to how we began the discussion today.

**55:39** You know, is the Amazon cloud right for you? It's Infrastructure as a Service. You need to understand that.

**55:47** The patterns and practices, yes, but what does ArcGIS technology do for you today?

**55:51** How can it better serve you in the Amazon cloud?

**55:55** We gave you a very brief introduction to an approach for beginning to dive in and investigate from a business perspective...

**56:05** ...how the platform should align with your business processes, and business cases, right?

**56:12** Sometimes folks do things just 'cause it's cool. Right? And maybe that's a valid reason.

**56:18** But I think what we're trying to...to get across to you is let's have a good reason to do it and then we'll do it.

**56:24** I'm just playing around in a development environment.

**56:26** That's a great reason, to be honest with you, 'cause then you get to know how the technology works.

**56:31** Think about all the different abstractions we talked about today. You need to become familiar with those.

**56:36** But, you know, the main points that we thought were important today, and there are many more, not in any particular order though...

**56:42** ...but are elasticity, I think this is a big deal in the GIS domain. I really do.

**56:49** I've seen limitations and failures with deployments because we just can't get to enough CPU when we need it.

**56:56** That barrier is going away.

**56:58** Time to market. Time and time again it has taken us a long time to get a technical workflow up...

**57:04** ...and supportable so we can support the business. This eliminates that as well.

**57:10** Risk aversion. You know, I love the concept that if something's not just working I can kill it. That's a neat concept.

**57:21** And then the budgetary one too. I mean I thought this was important to include today, because personally I've had success...

**57:27** ...in swapping around budgetary arrangements to pay for cloud deployments, and with clients as well.

**57:33** And that's simply...

**57:37** How many people capitalize their hardware on their budgets?

**57:42** I mean, do you own this hardware that you're deploying ArcGIS Server to with Amazon?

**57:47** You don't technically own it, which means it comes from a different expense budget.

**57:52** So the message here, again, you've seen this in the plenary, we've talked about it today...

**58:01** ...is that ArcGIS 10 is the enabling technology here.

**58:05** It's pervasive. It can be deployed on premise as part of your enterprise or locally, or in the cloud.

**58:12** Other things don't change.

**58:13** The capabilities are still there for visualization, for the creating your data, managing your data.

**58:19** This has huge implications for a collaboration message. Huge. Right?

**58:24** The discovery of data, managing your data, analyzing data.

**58:27** ...This is genuine transformation of the platform.

**58:32** So, ArcGIS 10 is a complete solution, as well as a complete system. So did we include our...? No.

**58:44** [Unintelligible] Yeah, we're going to do some Q&A now. Is that okay?

**58:50** I had some predefined topics that I was going to suggest to you, but apparently it was in a different presentation.

**59:00** So with that, no awkward pauses, let's open the floor. Yeah, we'll start with Lisa.

**59:09** What are the security implications to using the cloud? [Unintelligible]

**59:26** Sure.

**59:27** Do you get security with the cloud [inaudible]?

**59:31** Do you want to repeat the question first?

**59:32** Yeah.

**59:35** Sorry, so the question is what are the security implications for using the cloud?

**59:41** The instances that you use in the cloud typically have two levels of firewalls.

**59:46** One that is set by the cloud itself; it's called the Security Group.

**59:51** And the other one is a firewall local to the machine itself.

**59:55** So within the instance itself you have pretty much good control over the security.

**1:00:00** The part where the cloud gets to be more open is when you are exchanging the data or moving it around.

**1:00:07** And for that part, the recommendation is to either encrypt your data or to resource for something...

**1:00:15** ...like a virtual private network where it's a special connection that's very well encrypted...

**1:00:20** ...very well isolated from the rest of the cloud.

**1:00:22** So in terms of specific federal security regulations and certifications and accreditations...

**1:00:30** ...the cloud has only been certified and accredited for moderate level.

**1:00:34** So not for all sorts of federal certifications.

**1:00:38** So from that perspective, they're not considered something as efficient as an on-premise environment.

**1:00:45** So it's worth kind of knowing the details and what you need...

**1:00:49** ...what exactly is the security level you're trying to achieve...

**1:00:52** ...and from there looking at a specific service in the cloud that will match that.

**1:00:58** [Audience question] So do you have to manage the [inaudible] as far as on that specific [inaudible]?

**1:01:02** Yeah, the security groups are very easy to manage.

**1:01:05** In the other session that we have on the cloud, if you'd like to attend it, we demonstrate how you set this up.

**1:01:10** It doesn't take...it takes less than a minute. Very fast and very effective.

**1:01:16** So if all what you want is just to make sure nobody else is going to be accessing your machine, you can control that.

**1:01:20** Like I said, there are two levels.

**1:01:22** There's the security group, which includes many instances, but there's also the local firewall.

**1:01:26** You can turn that on and keep it enabled for your instance.

**1:01:29** And then if you are exchanging any data that you feel is sensitive, you can keep it encrypted and that will protect it.

**1:01:37** In a sense, there's sometimes not a lot of differences between different zonal approaches within your own business today.

**1:01:46** Not sure who's next.

**1:01:48** [Audience question] Can I, for example, implement a three-tier architecture...

**1:01:52** ...and keep the database server on premise with the web server, application server in the cloud?

**1:01:59** The question - and pardon me if I get it wrong - but I think your question is, Can I implement a three-tier architecture...

**1:02:04** ...where I have the database on premise, correct, and the web server and application server in the cloud?

**1:02:11** So is the presumption that you would be going back to your on-premise data in real time to draw it back?

**1:02:17** I don't think that would be a best practice, okay? You're creating a lot of single points of failure there.

**1:02:23** There might be a better way of taking the information, the data, out of that on-premise instance...

**1:02:28** ...and moving it onto the cloud in some way, shape, or form that's optimized for serving it back out.

**1:02:33** I don't know if you want to add anything to that.

**1:02:35** Yeah, I would say that would be the recommended approach is keep your database with your server.

**1:02:42** Because they communicate through TCP, and you don't really want that to be over the Internet.

**1:02:46** Yeah. Yes, sorry. Yes.

**1:02:49** [Audience question] Is there an educational cost schedule for universities and colleges?

**1:02:57** Is there an educational cost schedule for universities and colleges.

**1:03:03** Does your university or college, you know, have an enterprise license agreement with E-S-R-I, Esri?

**1:03:08** [Audience answer] Yes.

**1:03:09** Okay. So yes, there's a cost arrangement and agreement there.

**1:03:12** [Audience question] Is that on the Amazon side as well?

**1:03:13** No. As far as I'm aware today, the answer is no. I don't know if anybody else wants to chime in on that from the audience.

**1:03:21** I can confirm that, that these are two separate things.

**1:03:25** The way the AMIs work is you start them, but you bring your own license, so if your license is an ELA...

**1:03:31** ...you can then use it as many times as you need depending on the number of cores you're starting.

**1:03:36** But if your license is only four cores, you are bound to the same rules you are bound to in on-premise environment.

**1:03:42** Plus you'll get the bill that Marwa showed you, or the university will.

**1:03:47** Yeah, so these are two different costs right now.

**1:03:49** We're not bundling the licensing with the actual Amazon costs.

**1:03:53** And Amazon bills you for their costs, and you have your license...

**1:03:57** ...and you license the instances once you enable them based on that.

**1:04:02** I don't know if you want to add anything.

**1:04:03** [Audience participation] So Amazon does have a really neat grant program. Go to [aws.amazon.com/grant](https://aws.amazon.com/grant).

**1:04:10** Universities and other educational institutions can take advantage of their grants, and they provide free [inaudible].

**1:04:17** [Inaudible audience question]

**1:04:18** Yeah. [aws.amazon.com/grant](https://aws.amazon.com/grant).

**1:04:23** Yeah. Thank you, Neil.

**1:04:25** Any questions regarding our bundle too, this is a great resource over here, Neil Tomlinson.

**1:04:31** So we have some more hands. Yes?

**1:04:33** [Audience question] You mentioned that you have two Esri AMIs right now that you're testing.

**1:04:39** How does the individual organization create their own for deployment and use in the cloud?

**1:04:46** So if you're going to create your own AMI, we recommend that you base it on the Esri AMIs.

**1:04:53** So you probably want to build it because you have your own software stack that you would like to add to the Esri software...

**1:05:00** ...or are you talking about just replacing what Esri has done in their own AMIs?

**1:05:04** [Audience answer] Adding additional things. Adding data, having [inaudible]...

**1:05:09** Sure.

**1:05:10** ...[inaudible] certain use cases.

**1:05:11** Yeah. Start with ours is what we want. Yeah.

**1:05:13** [Audience question] You can start with yours, add to it and save it as your own.

**1:05:17** Yes. And that's the concept of a custom AMI.

**1:05:19** Now these AMIs are great to distribute internally within your organization, but according to our license setup...

**1:05:27** ...when you start sharing it with another organization, that you need to come back to Esri with that.

**1:05:32** [Audience question] If I have an ELA, I can use it in my organization.

**1:05:34** Yes. Yes, absolutely. Yes?

**1:05:39** [Audience question] My question had to do with the reserved instances and Amazon's pricing structure.

**1:05:43** Does that only apply to one instance, or are there like...

**1:05:48** Because I heard you say, well, you can just kill that instance and create a new instance, you know.

**1:05:54** How does that work with the reserved instances if you only get one, or can you get multiple instances with the reserved...

**1:06:00** What they say is you have two instances running for the duration of a year. Okay.

**1:06:04** Now you could stop and start and kill and start new instances.

**1:06:08** As long as there are two instances running all the time, you will achieve the cost savings that they promise you by doing that.

**1:06:18** Yes?

**1:06:19** [Audience question] Do you support these [inaudible] AMIs with the [inaudible] on premises?

**1:06:23** I'm sorry.

**1:06:25** [Audience question] Like running an [inaudible] against the cloud on premises...

**1:06:28** ...do you use those AMIs like you use other AMIs?

**1:06:31** Well, these AMIs are now built and optimized for the Amazon environment.

**1:06:36** They take advantage of specific capabilities in the Amazon environment.

**1:06:39** To build something similar to that for different virtualization software would require building it for that software.

**1:06:47** But right now, the ones we have work just with Amazon.

**1:06:50** [Audience question] How would you go about building them?

**1:06:54** So the question is how would we go about building that custom one?

**1:06:57** I think we could maybe just have a sidebar conversation with you.

**1:07:00** Yes, and also the other session we have, How to Use ArcGIS Server in the Cloud, we go into a lot of details.

**1:07:05** We had that one earlier today. There's going to be two other offerings for it Thursday and

Friday morning.

**1:07:13** We go into a lot of details on how you start your own instance, how to set it up with your own data...

**1:07:18** ...and just create your own custom EBS, custom AMI from there.

**1:07:22** If you would like, I could show you that also offline. Yes.

**1:07:29** [Audience question] Correct me if I'm wrong, you said that, in terms of like optimizing ArcGIS Server...

**1:07:33** ...you wouldn't want to add SOC machines; you would want to use the...

**1:07:38** Full AMI.

**1:07:39** [Audience question] ...the elastic cloud and just add SOC instances on that one machine?

**1:07:44** Now when you launch an instance from the AMI, it includes both the web applications, the web server...

**1:07:51** ...the SOC and SOM, all on one machine.

**1:07:53** So what you do, when you start multiple instances, is you enable the Elastic Load Balancer to seed these different instances...

**1:08:02** ...and it will do the brokering between them and...

**1:08:05** [Audience question] Isn't that more than one SOM on each machine then?

**1:08:09** Yes. In that case, the SOM isn't really doing any distribution.

**1:08:13** [Audience question] For the actual...

**1:08:14** Yeah, it's just running on one machine, so...

**1:08:18** [Audience question] But that's the difference between, like, real hardware...

**1:08:21** ...where you add another machine that you dedicate as the SOC machine, right?

**1:08:25** It's a different approach, but we've found that to be a very efficient approach for the cloud to kind of deploy it in that way.

**1:08:32** So you have one AMI and then you create multiple instances.

**1:08:36** The load balancer works very well if you're trying to do geoprocessing and stateful type of services.

**1:08:43** You can use the sticky sessions in the Elastic Load Balancer and that will allow you to do that kind of thing pretty seamless.

**1:08:49** The only drawback for it is that ArcMap doesn't really support the sticky cookies yet.

**1:08:57** The deployment pattern though that you're talking about is utilized, you know, not just in

the cloud.

**1:09:05** We can talk more about that later too, so whether it be a software load balancer or a hardware load balancer. Yes?

**1:09:12** [Audience question] So does that mean that I only have to have several SOC instances [inaudible]?

**1:09:20** With this deployment pattern, that is correct. Yeah.

**1:09:24** Yes?

**1:09:26** [Audience question] Go ahead and tell us about your jump start and your bundles.

**1:09:30** Okay. We'll start with the bundle, okay?

**1:09:33** We actually have the gentleman who architected this thing sitting right here...

**1:09:36** ...so that's why I'm looking at Neil, for the bundle anyway.

**1:09:39** You want to come on up? I'm glad you came.

**1:09:43** I'm not.

**1:09:50** So do you have a specific question about the bundle?

**1:09:54** [Audience question] How it works, cost.

**1:09:55** Okay. So basically it's a preconfigured instance that we have wrapped our term licensing around.

**1:10:03** It enables a lower point of entry for standing up ArcGIS Server, and it does take advantage of the AMI.

**1:10:09** So everything we've discussed up here is possible with the bundles.

**1:10:13** [Audience question] What is the cost?

**1:10:19** Yeah. Can we take that offline? Okay, great.

**1:10:22** Have you sync up with your account manager. What's that?

**1:10:27** [Inaudible audience question]

**1:10:29** Well, not offline. I just think that we... It has to do with his particular licensing.

**1:10:34** So the cost of a four-core AMI would be no different than a four-core license that you're deploying on premise.

**1:10:43** So if you're paying X amount of dollars for that four-core license, it's the same. You can take your licensing anywhere.

**1:10:50** You know, it's just within an audience, I don't want to presume what people are paying for their four cores of license.

**1:10:57** The retail cost is \$40,000.

**1:10:59** So for a four-core license, to deploy it in Amazon, you can take your licenses wherever you want to go.

**1:11:08** So I would encourage everyone, now we have the Managed Services Island available in the booth area.

**1:11:15** If you would like to stop by, we could give you an idea about the differences between the bundle...

**1:11:19** ...and the managed services packages that are all based on the cloud.

**1:11:23** So if you want to leverage that and have Esri do the management on your behalf, you can learn about that.

**1:11:30** You can learn about the costs and all the details regarding that.

**1:11:34** [Audience question] Which island?

**1:11:35** The Managed Services Island. Any other questions? Yes.

**1:11:46** [Audience question] The EBS, is that tied to an instance?

**1:11:49** Yes. Typically an EBS is attached to a specific instance.

**1:11:52** [Audience question] So if you upload data to an EBS, say if you had three instances running...

**1:11:58** ...do you have to upload that for each one separately, or should you upload it to one, kill everything, redo [inaudible]?

**1:12:06** Well, it depends. There are different approaches to that.

**1:12:10** Some folks prefer to put it on one EBS drive and share that like a Windows share, just like you do on premise...

**1:12:17** ...and from there have multiple instances all look at one EBS drive.

**1:12:21** Or another alternative is you can create one EBS drive per instance.

**1:12:26** So to do that, just like you create an AMI for a machine, there's a snapshot for an EBS drive.

**1:12:32** So you can create a snapshot and from there launch multiple copies and then attach it to each of the instances.

**1:12:39** The latter is actually the most recommended, which is considered best practice, because you can see with it the best performance.

**1:12:48** [Audience question] Is Esri looking at branching out into other cloud services? I know Windows has Azure and ones like that.

**1:12:55** I'm just curious to know, are you looking at setting up the equivalent of AMIs for those?

**1:13:02** Yes, we are. Okay.

**1:13:03** Yeah. But to elaborate a little bit on that, Azure doesn't have the infrastructure capabilities for services.

**1:13:10** As they do, we are going to definitely be interested, but right now you can use the API to build your applications...

**1:13:16** ...and do what you need with it using specific components that we have available on that.

---

© Esri 2013 <http://www.esri.com>