

# Using ArcGIS for Server in the Amazon Cloud

Join this session to learn how to leverage the Amazon cloud with ArcGIS for Server. You will learn how to setup ArcGIS for Server in the Amazon Elastic Computing Cloud, typical deployment patterns and workflows. We will also discuss how to estimate costs for your ArcGIS for Server deployment in Amazon and present several projects running ArcGIS for Server in the Amazon cloud.

<http://video.esri.com/watch/62/using-arcgis-for-server-in-the-amazon-cloud>

---

## Video Transcription

**00:01** My name is Ismael Chivite. I'm the ArcGIS Server product manager...

**00:03** ...and with me is Marwa Mabrouk who works in the Professional Services Division in Esri.

**00:08** And she has extensive experience deploying ArcGIS Server on Amazon EC2.

**00:13** So Marwa's going to present at the end kind of some experiences that she had with this...

**00:18** ...and I will kind of cover the product presentation, explaining what the product is and how...how you use it.

**00:23** So we will have kind of a step-by-step process so you can see how this thing really works.

**00:29** I want to point out that there is another session on ArcGIS Server on Amazon EC2, which covers the business aspects of it.

**00:37** Like why should I run ArcGIS Server on Amazon? Or why I shouldn't.

**00:43** This session is more about, okay, I think I...I know why I want to run on Amazon, so just tell me how this thing really works. Okay?

**00:51** So this is kind of a technical aspect of it, and then we have the business aspect of it.

**00:56** Although there is a little bit of overlap between the two, as you will see.

**01:00** Also I want to remind you that we have these surveys...

**01:03** ...that you should fill out and put in the boxes outside if you remember that at the end of the session.

**01:10** Please, if you have one of these, probably you don't because it means it's very old, but if you have one of these, you know what to do with it.

**01:17** Right? Just turn it off, and I'll do that as well. That's kind of embarrassing when...okay. Let's go into it.

**01:29** First, a brief introduction to ArcGIS on Amazon EC2, then basic concepts. This is the step-by-step process to get started.

**01:39** We will also touch on some of the key areas where people have problems with ArcGIS Server on Amazon EC2...

**01:46** ...specifically managing your GIS data; how do you get your data into the cloud and off the cloud as well.

**01:52** And also we will explain how you scale out and how you shrink ArcGIS Server deployments in the Amazon cloud.

**02:01** We'll go and look at the real-world examples and then hopefully have the Q&A.

**02:06** Let's just start with the introduction, and this will be very basic, just in case that you are not familiar with Amazon Web Services.

**02:14** Amazon Web Services is really a cloud infrastructure that is...

**02:19** Basically, these guys have data centers around the world, and they rent these data centers.

**02:26** You can rent computing, that is, servers. You can rent storage. You can also rent kind of network services for improved bandwidth.

**02:39** All these capabilities - computing, storage, network - you can rent, you can purchase through Web services.

**02:47** It's kind of a self-service approach to building an IT infrastructure.

**02:54** They have different Web services to access these capabilities; for example, EC2 is Elastic Cloud Computing...

**03:01** ...is the service that allows you to rent virtual servers.

**03:05** So you can say on...you can go and say, I want one server with two cores, or I want two servers with four cores. That's EC2.

**03:13** They have another one, which is CloudWatch. Is a service that allows you to monitor the services that you rented.

**03:21** Other services like CloudFront allow you to speed up download of files from the cloud.

**03:28** So you are storing your map tiles in the cloud...

**03:30** ...and you can enable CloudFront to kind of use edge caching to accelerate download of these map tiles into the clients.

**03:38** They have actually many, many different services, including other network services like virtual private network, storage, and so on.

**03:47** The unique characteristic of these data centers is that not only they offer you the self-service experience...

**03:54** ...you don't need to call anybody to get these resources; you just go to a website and you...you rent these things.

**03:59** But the unique characteristic is that you pay as you go, so storage, you pay...

**04:05** Well, you are storing 100 gigabytes, then every month, you will pay for this storage.

**04:11** And with computing, it's the same way. Are you using two computers or one computer?

**04:16** Oh, you stopped using one computer, then you can stop paying for that. You get the point?

**04:20** It's just pay as you go.

**04:24** We will sometimes refer to this as the Amazon cloud or Amazon EC2.

**04:32** For ArcGIS Server users, the advantage of running ArcGIS Server in this infrastructure are...simplified management is one of them.

**04:40** It's very easy to set up ArcGIS Server in this Amazon cloud infrastructure because we have preconfigured machine images...

**04:51** We'll talk about that. We call them AMIs.

**04:53** ....that are basically virtual images that we give you so you can pretty much select a size of the hardware you want to run ArcGIS Server on...

**05:02** ...and literally activate it. Say, Just make it happen; install ArcGIS Server. It will be preconfigured.

**05:08** And also, simple management because once the server is up and running, they have these very advanced tools to create backups...

**05:17** ...to monitor the CPU usage and disk I/O of your servers...

**05:21** ...to add more machines to a cluster of servers as needed, as the load increases, and also to shrink the deployment.

**05:30** And you do all of these doing this self-service experience.

**05:34** It's also interesting because Amazon gives you...I wouldn't call it real time, but almost real-time access to these resources.

**05:42** So if you need computing power, you don't need to wait for a machine to be delivered to your organization...

**05:49** ...and then you plug it in and then IT kind of goes and plugs all the things.

**05:55** I know of some organizations that it takes them six months, literally six months...

**06:00** ...to just get a machine where they can deploy their applications.

**06:04** And then in Amazon, once you have this self-service experience, it's kind of, you go there, you select...

**06:09** ...and in a few minutes, you have a server up and running on the Web.

**06:13** And that's a big, big reason why people move to the Amazon cloud, because they can go from idea to application very quickly.

**06:20** And often, it's not just about the technical aspects of these things, but just the red tape that gets you to the final solution.

**06:29** The same with storage and so on.

**06:31** And finally, as we mentioned before, you don't make a long-term investment into your infrastructure.

**06:37** People refer to this as pay-as-you-go model; you will also hear it...hear it as elastic...elastic computing in the context of servers.

**06:46** And the classic example is, I have an application and I honestly don't know how many people are going to use it.

**06:52** I'm going to try with one server. And that server is going and then more people use the application, right?

**06:59** As the load increases, response time decreases.

**07:00** Now you buy a second server for that increased load, and you buy a third server.

**07:06** What happens when people stop using your application?

**07:08** You wish you could go down to one server, but you already purchased two servers.

**07:14** In the Amazon cloud, you say now I want two additional servers; now I don't want them.

**07:19** And when you shut them down, you stop paying for that.

**07:21** And that is interesting from a financial point of view in some cases.

**07:26** It's not like going to the Amazon cloud is cheaper than deploying on premise, but in these scenarios where they have...

**07:32** ...you have these peaks of load, this can save you a lot of money.

**07:39** Here are some typical scenarios where people want to run ArcGIS Server on Amazon. One is development.

**07:46** For developers, EDN subscribers, starting with ArcGIS Server 10, have access to our virtual images in the Amazon cloud.

**07:54** And it's interesting from a development point of view because it saves you time.

**07:58** You can set up an ArcGIS Server 10 instance in just a few minutes as a developer.

**08:04** And you don't need to allocate a machine within your organization to run that server.

**08:08** Just go to the Amazon cloud, you pay \$50 for running 25 hours of your server, and then once you are done, you terminate the machine.

**08:18** So it kind of saves money in that regard.

**08:22** The other aspect is that you may want to run, on the Amazon cloud, your ArcGIS Server just

for development purposes.

**08:28** Maybe you are developing an application that you will deploy on your own premises servers.

**08:34** Right? Develop in the cloud; once you have done your work, you terminate the machines...

**08:39** ...and then you bring down your applications, and then you deploy them on premise. Or you deploy them on the cloud; that's your choice.

**08:46** The other one is [unintelligible] ad hoc GIS analysis. As you know, many GIS operations are very high CPU intensive.

**08:55** You need a lot of computing power if you want to get your job done in a small amount of time.

**09:00** So the easiest thing to do is to add more iron to these tasks; you add more servers.

**09:06** In the Amazon cloud, the advantage is...is that you have almost unlimited access to this computing power.

**09:12** So you can say, Give me 10 servers, and then boom.

**09:16** Ten servers are there, with ArcGIS Server, your services, your data...

**09:19** ...and you can split or break down your task across these different machines.

**09:25** We published a white paper some time ago, it's called ArcGIS Server in Practice, large batch geocoding...

**09:31** ...where we explain a technique by which you can take a cluster of peers, of machines that are running ArcGIS Server...

**09:39** ...and kind of break down your long tasks across all these machines.

**09:43** This technique works well on premise, and it works well, as well, in the Amazon cloud.

**09:52** The other one is I just want to have a Web mapping site up quick, so I go to the Amazon cloud...

**09:57** ...because I don't have to worry about having this machine 24/7 running here on premise.

**10:02** I don't need to worry about people using bandwidth in my network.

**10:06** It's all running in the Amazon cloud.

**10:08** And this links again with this concept of elasticity where I cannot afford to be down.

**10:14** This is a screen shot from an application that British Petroleum has for the oil spill thing.

**10:18** Probably the last thing they want at this point is to have their website down after what happened...

**10:22** ...so they are deploying their...this website in Amazon because in case that many people hit...

**10:27** ...they can almost immediately react to that and add more machines or drop them. Okay.

**10:35** Let's talk about some basic concepts of how you deploy ArcGIS Server in the Amazon cloud.

**10:43** And I will start with us giving you some terminology that we are going to use.

**10:49** The AWS, Amazon Web Services, management console. It's a web application.

**10:55** You open a web browser, you go to this URL, [aws.amazon.com](https://aws.amazon.com), and then boom. You can get into this console.

**11:02** It's basically a dashboard and a management tool. A dashboard because allows you to monitor the servers that are running for you...

**11:10** ...and a management tool because allows you to launch new servers...

**11:16** ...stop them, start them, terminate them, create backups, upload data, monitor the activity.

**11:22** And it also allows you to do that across any of the data centers that Amazon has around the world.

**11:30** That's what gives you this self-service experience that we were talking about.

**11:36** Another concept is this concept of AMI.

**11:39** AMI stands for Amazon Machine Image, and it's nothing but a virtual image that works in the Amazon cloud.

**11:46** It's kind of a cookie-cutter approach to putting an operating system, middleware or, in this case, ArcGIS Server...

**11:55** ...as well as the data that ArcGIS Server needs, into a server.

**12:00** So you get a server, you rent one, and then you say, I want to use this AMI.

**12:05** This AMI has ArcGIS Server, it has the data for world population and my services.

**12:11** And when you put that AMI within your server, your server is literally up and running; just can go and it works.

**12:19** We have two AMIs, we, Esri. We provide two AMIs to you.

**12:24** Basically, once you have licensed ArcGIS Server, you can call customer service and then tell them...

**12:30** ...Look, this is my Amazon account ID, and I want you to share the AMIs with me. So we share these AMIs.

**12:37** So when you go to this management console that you saw in the previous slide...

**12:40** ...you can select the ArcGIS Server AMI or the Enterprise Geodatabase AMI.

**12:45** The ArcGIS Server AMI, think of it as the GIS services server.

**12:51** We put...basically the...the GIS services install in there...

**12:55** ...map services, geoprocessing, geolocators, ArcGIS Server Manager, the services directory.

**13:02** All of these things are preconfigured in this AMI.

**13:04** It also has a SQL Server Express database.

**13:08** We initially put it there because it's kind of nice to have a...as a user store so you can secure your services.

**13:14** And we wanted to kind of simplify that experience, so SQL Server Express is there.

**13:19** But then, we also enabled ArcSDE. So if you have an ArcGIS Server Standard Workgroup license, for example, we'll give you this.

**13:28** And that gives you SQL Server with ArcSDE, ArcGIS Server.

**13:31** It also has ArcGIS Desktop installed in the machine, which you should use to administer your services.

**13:38** It's not for you to sit and start an edit session or do cartographic production because you will interact...

**13:44** ...as you will see, with the servers through Remote Desktop.

**13:46** So it's kind of a clanky experience when you are using ArcMap over Remote Desktop, but it's there for you to do things like...

**13:54** ...fixing the path of your map documents or your geoprocessing services...geoprocessing models before you create these services.

**14:02** We also attach 100 gigabytes of additional storage to these machines, then you can actually get rid of it or add additional storage if you like.

**14:09** But it's...it's there. So that's the GIS services AMI.

**14:13** And then we have the Enterprise Geodatabase AMI, which is basically a Postgres database with ArcSDE enabled.

**14:21** And the idea is that you launch the ArcGIS AMI and you have ArcGIS Server running...

**14:26** ...and then you launch another server where you put this Enterprise Geodatabase AMI.

**14:31** And you really don't need to log in to this machine unless you want to get into the table spaces management and so on.

**14:38** You can literally connect from your ArcGIS Server using ArcCatalog into this Enterprise Geodatabase AMI.

**14:45** So this again, completely preconfigured.

**14:47** Both AMIs run only on Windows 2008 64 bit. Only Windows 2008 64 bit for now.

**14:59** Another concept is EC2 instances. Amazon Web Services refers to servers as EC2 instances.

**15:07** An EC2 instance is a virtual server. Okay? Just think of it as just a server, and they have different sizes of them.

**15:15** They have the standard sizes, dual-core, four cores, and then they have the high-memory and high-CPU groups of EC2 instances.

**15:26** So this means that, well, you are really working with virtual images all the time, 'cause data centers have physical servers...

**15:31** ...but you only can instantiate virtual servers within those.

**15:35** And depending on how much of these physical servers you want to use, you will start a big EC2 instance or a small EC2 instance.

**15:43** Do you got the point?

**15:45** So now you can make the link between the AMIs and EC2 instances.

**15:50** If EC2 instance is my, let's call it the hardware, I put the AMI, that's the software.

**15:55** And together, it's a fully functional machine, right, with the software running.

**16:01** As you can imagine, these guys are making money somehow, right? So the bigger the machine, the more you are going to pay for it.

**16:08** But it's always an hourly price for these machines. It ranges from I think it's like 98 cents or so to about two-point-something dollars per hour.

**16:20** And you can...those are the prices on demand when you make a commitment for one hour.

**16:24** And you can actually rent these servers for one or three years. They call that reserved instances.

**16:30** And then the hourly prices drop, although you have to pay an up-front fee.

**16:34** This is all very well explained in their website. They are pretty transparent about pricing for their EC2 instances.

**16:41** Also, the prices vary with size, but vary also with the location.

**16:46** The European data center is slightly more expensive than the ones that we have in the U.S. and so on.

**16:57** Okay. Another concept. I'm going to throw some concepts, and then we...you'll start all the pieces getting together.

**17:03** Elastic IPs. So what happens when you start a new EC2 instance? The clock starts ticking, right? They start charging you per hour.

**17:13** At some point, say you are developing an application, you went...you may want to say, I want to stop my instance...

**17:18** ...because I want to go home to sleep, and I don't want to pay for this...this instance while it's running overnight.

**17:23** So you stop the instance. Now the next morning, you come and you start the instance again...

**17:28** ...and it will just have your applications as you left...as you left them the day before.

**17:33** But the thing is that because these are virtual machines that they hydrate and rehydrate when you start and stop...

**17:39** ...your virtual machine may show up in actually a physical server that is different from the one you ran yesterday. Right?

**17:47** It will still be a small machine or a big machine as you decided, but physically might be a different machine.

**17:52** And I'm oversimplifying things here, but essentially the...the IP of your machine may change. It actually changes quite often.

**18:00** So what you do for your applications, your applications are going to point to your IP, the IP of your machine...

**18:07** ...so if the IP changes, the applications will stop working.

**18:10** So these elastic IPs are basically IPs that you rent and you say, I want to reserve an IP, and they give you an IP that you can associate with an instance.

**18:20** So your applications will always go through this fixed IP regardless of the host name of the virtual machine. Okay?

**18:28** This is important for you to keep in mind, because our AMIs are not just, oh, we put ArcGIS Server in a virtual machine and there you go.

**18:36** No. These AMIs that we built and that you can get have special logic that makes that, on start-up, we check for these host name...

**18:45** ...and we do all the necessary changes in ArcGIS Server Services Directory, the cfg files, and so on, so ArcGIS Server runs.

**18:53** So when you start, ArcGIS Server is ticking, right there. Okay?

**18:59** Now, elastic IPs are good for this use case where you want to have a well-known IP for your applications...

**19:05** ...but there is also another interesting case for elastic IPs.

**19:08** And it's one where I'm going to start developing my application on a small instance because it's cheap.

**19:15** But...and then I put elastic IP on top.

**19:19** But when I go into production, I don't want to use the small machine; I want to use a bigger machine.

**19:23** So you may want to start a bigger machine and then associate your elastic IP with the bigger machine.

**19:30** But the software, the AMI that runs this bigger machine, is the same as the small one.

**19:35** So you're actually, under the covers, you are switching your server from a small machine to a big machine...

**19:41** ...and the applications don't even notice. Do you...do you follow the idea?

**19:52** Of course, they charge you for elastic IPs.

**19:55** Regions and availability zones. I said that Amazon Web Services has several data centers around the world. They call them regions.

**20:02** They have two in the U.S., one in Virginia, the other one in...is it California or Oregon? California.

**20:10** And then they have one in Singapore and one in Ireland.

**20:14** And that is kind of interesting because allows you to actually deploy your application across different data centers if you like...

**20:20** ...but it's also interesting because of legal regulations. Some countries will not allow you to get your data outside of your country.

**20:28** Like in the European Union, if you have confidential information, the data shall be hosted in the European Union.

**20:35** So you will go to that data center, right?

**20:39** Canadians don't like to...to store data in the U.S. You know that, right? Don't you know that? Yeah. It's by law in some cases.

**20:47** So it's good to be able to play with different data centers for those reasons.

**20:51** Another aspect is, aside from the regions, within each data center, which in...within each region, they have availability zones.

**20:59** These are sections of the data center that have different Internet access point, isolated Internet access points...

**21:09** ...as well as isolated...how do you call that? Like energy.

**21:15** Like if the power goes out, each of these availability zones has its own source of electricity.

**21:25** They are...okay. Within one region, they have independent data centers, and they call them these availability zones, right?

**21:32** So you will work with these availability zones if you want to build high-availability applications, right?

**21:38** You deploy everything in Ireland, but some of the machines in my cluster are in availability zone A; others are in B.

**21:45** And then you can kind of do network load balancing between these two availability zones.

**21:54** The AMIs that we provide to you can work in any region, any availability zone.

**22:00** So how to get started. I kind of explained this, but let's go through the workflow again.

**22:04** First, you...you make a commitment with Amazon. You create an Amazon account...

**22:11** ...because to us, Amazon is kind of another delivery mechanism for ArcGIS Server.

**22:16** You want to run your ArcGIS Server license in Dell, fine. Go to Dell and buy a machine. IBM? Go to IBM and buy the machine.

**22:23** You want to run on Amazon, go to Amazon and buy a machine.

**22:26** So you go to Amazon, and you create your Amazon account ID. At that point, you can start servers.

**22:31** But you cannot start servers with ArcGIS Server. You need to call Esri and tell them, This is my Amazon ID; it's a 12-digit number.

**22:39** And if you are licensed, we will share the AMIs with you.

**22:42** So the next time you launch a server, you'll see an option to launch ArcGIS Server.

**22:46** And at that point, you are off and going; you will have ArcGIS Server.

**22:50** So you just authorize normally. You authorize ArcGIS Server so it works, and that's pretty much it.

**22:57** We'll see that in another demonstration right now.

**23:12** Okay. So the first thing is I'm going to go to this [aws.amazon.com](http://aws.amazon.com). This is the website.

**23:21** Here you can learn about pricing, all the services they have, and so on.

**23:25** Today, we are going to click on the AWS Management Console at the top, and because Firefox is remembering my Amazon...

**23:34** ...I...I just logged in a few minutes ago. But basically, unless you have an Amazon account, you cannot get into this application. Okay?

**23:41** You will be prompted to type your region and password.

**23:45** This is not the typical Amazon account you use to buy books.

**23:47** This is an Amazon account from Amazon Web Services that has...is enabled to use EC2, Elastic Cloud Computing.

**23:55** It's fairly easy to create. I mean, not more difficult than...not harder than creating your account to buy books.

**24:02** And as you can see here, you can operate in different regions.

**24:07** Here are the different data centers. U.S. East, U.S. West, European Union, and Singapore. In this case, I'm going to work in U.S. East.

**24:16** And just to get things off the ground very quickly, I'm going to launch an instance.

**24:22** Now you remember, this EC2 instance concept we talked about, I'm going to launch one. Okay?

**24:27** And the first question that they ask is what type of software did you want to run on the server you are about to launch?

**24:37** These are public AMIs. So you can see Windows and Fedora virtual images in there.

**24:45** So let's say you want to start a server that runs Windows SQL Server 2008 and Windows Server 2008.

**24:51** Well, you come here, you say select. And that's the software that they will put into your machine.

**24:56** The ArcGIS Server AMIs won't show up here because our AMIs are not public; they are private...

**25:03** ...and we only share them with ArcGIS Server licensed users. Okay?

**25:07** So you will go to this tab, the Community AMIs tab, and then there are many community AMIs out there.

**25:14** You filter by either Esri or ArcGIS, and then you will see two.

**25:19** Actually, you see here a bunch because these are kind of the daily builds that we work...it's kind of a...this account is kind of a mess.

**25:25** But in your account, you will only see two AMIs, the ArcGIS Server AMI and the Enterprise Geodatabase AMI.

**25:33** And the process to start them up is very similar.

**25:35** You basically pick the software; in this case, ArcGIS Server 10. You select that and then, well, that's the software.

**25:43** But now tell me, how many machines do you want to start? Do you want to start 1, or do you want to start 5, or do you want to start 10?

**25:51** In this case, at this point, it doesn't really make sense to start 10; it makes sense to start only 1.

**25:58** Because this is an AMI that has ArcGIS Server, but it's not authorized, it doesn't have your data, and so on.

**26:03** But we'll see later, when you create your own custom AMIs, that that option is pretty powerful.

**26:08** Anyways, let's go with one.

**26:10** We are working on U.S. East, but which availability zone? So in this case, I have no preference; it's just for a demo.

**26:18** But in some cases, if you want to build high availability, you may want to start one here and the other one there.

**26:25** And this is a good one. What is the instance type? What is the size of the virtual machine you want to create?

**26:31** Is it a dual core? Four core? Eight core? You...you pick, right? You pick there.

**26:39** Amazon has way many EC2 instances that you see here, but because our AMI works on 64-bit, this is kind of a subset of their choices.

**26:48** So for development, when you are just starting, you can just go with a large...large, which is actually the smallest.

**26:54** This is about \$1 per hour, I think, in...in U.S.

**27:02** The next step is...shows some advanced instance options. And I just want to highlight one, which is this one, monitoring.

**27:09** This is an additional service that, when activated...

**27:11** And then I click Continue.

**27:13** ...will allow you to ping your servers and find out things like what is the CPU usage at that point, disk I/O, and so on.

**27:21** And that's key information that you may want to use when you are doing out scaling of your application...

**27:27** ...is to add more machines or fewer machines.

**27:30** It's a for-fee service; it's additional. In this case, I'm going to turn it off.

**27:37** They always give you choices to buy more, buy more, buy more.

**27:43** Okay. The next one is important. This step is important. This is a step where you define the firewall settings of your instance.

**27:53** And this is very interesting because when you start machines in the Amazon cloud, you have pretty much control over your machine.

**27:59** It's a virtual machine, but you can lock any port and open any port that you like...

**28:04** ...and you can restrict access through those ports to your machine from pretty much any IP.

**28:24** [Unintelligible] Continue... Okay. So actually I...I...I went too quick.

**28:31** Before the firewall settings, this is the...the key pair so the idea is when you launch an instance from scratch...

**28:38** ...from the AMIs we provide, we create a random password so nobody can get in your...in your server.

**28:45** And for you to decrypt this password, you need to create one of these key pairs.

**28:50** So...well, if you want to create one, just click here and then you download it. This...this just takes two clicks.

**28:57** But I'm going to use an existing one, and when you download the key pair, you will download this file.

**29:02** It's a PEM file, and it has this cryptic information in it. Okay? You will...we will use this later, as you will see.

**29:12** Anyways, you select that key pair and then this is the firewall settings.

**29:18** So in the firewall settings, let's create a new security group for you.

**29:22** You create security groups because normally the firewall settings are common to all your instances.

**29:27** So you create a security group and all these rules are preconfigured so you can reuse them in the future.

**29:32** So I'm going to say, okay, this is my security group for ArcGIS Server services, public-facing applications.

**29:41** Okay, so rules. First of all, I want HTTP access over port 80.

**29:47** And I want every machine in the world to be able to access this machine through HTTP 80.

**29:54** So I'm not going to restrict the IPs for access.

**29:58** And then I also want to open Remote Desktop Access, but I don't want anyone to remote desktop to my machine...

**30:06** ...so here you can type the range of IPs that can actually remote desktop to your machine for administrative purposes.

**30:13** So you would add the rule, and then you would continue. And actually, we can just continue, right? Those are my rules.

**30:20** And then you say, Launch. At this point, something is happening in U.S. East, and a machine, a virtual machine is created with me...

**30:31** ...is created for me with all these characteristics.

**30:33** So I can go to the instances page and see my instances. And actually, it went pretty quick.

**30:41** It goes from...oh, here it is. Pending. That machine is being built for me.

**30:46** This flag may turn red...I mean green, saying running, pretty quickly...

**30:50** ...but normally it takes between 10, 15 minutes to really have a machine up and running. And the times vary.

**30:58** I have seen them start in 4 minutes, in 20, in 25 minutes. So that's kind of a...an unknown. Sometimes take longer than others.

**31:06** It also depends on the size of the machine that you use.

**31:10** Bigger machines start faster, not because they are big but...but because there are more available; people tend to get the smaller machines.

**31:18** Anyways, if I go on, right-click here, I go...I can look at the system log.

**31:24** And the system log is empty; this is telling me that the machine is not ready to go.

**31:29** So, as I said, this process may take more time or less time, so before the presentation, I actually launched an instance...

**31:36** ...I gave it 15 minutes to start, and hopefully, it's up and running. Okay.

**31:41** So that's what you should be looking for. Ready to use. It says, Windows is ready to use.

**31:45** At this point, I can go into the machine and use it.

**31:50** Okay. So to get into this machine, the first thing I need is to get the Windows password, right?

**31:55** So I click here, and this dialog shows up.

**31:59** So now is when I come here, I take my key pair, I paste it here, and then I decrypt the password.

**32:10** That's the administrator password. And only I will know this password, so I just did copy it in my clipboard.

**32:19** And now I'm going to access to...the machine. You can see this option, Connect.

**32:25** Let me double-check that this is in fact the machine I want to use. Yes.

**32:29** And I say Connect. And there are two gateways into this machine.

**32:33** This machine is running right now with ArcGIS Server, but ArcGIS Server is not authorized yet.

**32:37** The gateway on the right, okay, this is the public DNS. This is the URL of my machine.

**32:43** If people want to access it with a Web browser, they will go there.

**32:46** And then the other gateway is this Remote Desktop file.

**32:50** So I click here, and this basically downloads a desktop file, Remote Desktop file, which I can use to connect to my machine in the Amazon cloud...

**33:00** ...and then I paste my Windows password, and I open a Remote Desktop connection to it.

**33:24** Okay. Okay. So this is...this is pretty impressive. I mean, this is the machine we just launched, and it's running somewhere in Virginia...

**33:35** ...and we are accessing it through Remote Desktop to control it.

**33:39** So now the machine is kind of starting and it shows up this home page.

**33:44** And you can see in the home page that it says that ArcGIS Server is not authorized.

**33:50** So now is where you need to authorize the software. And I'll do that very quickly.

**33:54** So first thing is, I need to get my authorization files, which I have here in my laptop.

**34:01** And I'm going to simply copy and then go to my remote desktop machine, and I'll paste on my desktop. Okay? Desktop.

**34:15** And at that point, I can go and say, okay, I have my authorization file, desktop and server, and this will authorize ArcGIS Server.

**34:35** That's it. Okay. In fact, if I now go back to my home page, you'll see that the message for the license went away...

**34:44** ...because ArcGIS Server is authorized. So this means that ArcGIS Server is running.

**34:49** So just to kind of have something that you...you would see that it's running, what we did was we created in these AMIs two sample services.

**35:00** So if I go back to Firefox, let's go with the other gateway we talked about, the public DNS. Okay?

**35:06** I'm going to use that and I'm going to copy that into my Web browser, Firefox, and I'm going to navigate to the services directory.

**35:21** It's the first hit, and the first hit, it always takes a bit more of time. Did that right, right?

**35:33** [Background conversation]

**36:00** Okay. Who knows. Anyways, that's my services directory, and it has two services up and running already, okay?

**36:07** It's just you normally will, when you take ownership of your machine, you normally get rid of these services, right...

**36:14** ...and you put your own, but they are actually up and running.

**36:18** The other thing you may want to do is have a look at the...at the service. Okay, that's the service. That's the machine we just started.

**36:30** Okay? So how much time we needed to do this? Around 20 minutes to get [ArcGIS] Server up and running on the web.

**36:36** Because if now...you saw the plenary, right? You know you can go with an iPhone now and look all this up.

**36:41** Yes, in fact, you could...you just connect to that, and you connect to it.

**36:44** Because my firewall settings allow you to do that.

**36:49** Okay. So I just want to finalize the authorization because of course [ArcGIS] Desktop is installed, but it's not authorized.

**36:58** So I'm going to authorize it. I open the administrator, desktop, authorize, same process, ArcInfo.

**37:15** In customer service, always they ask me why I authorize [ArcGIS] Desktop twice a day. Okay. That's it.

**37:25** So now [ArcGIS] Desktop will authorize here. Before I leave, I will change the password.

**37:32** Actually, I'm not going to change the password, but this...that's one of the typical things that you would do.

**37:37** You would authorize ArcGIS Server and then you change the administrator password that...to something that you will...you will remember.

**37:45** Okay. So that showed the process of starting ArcGIS Server in the Amazon cloud. Is that something interesting? Yeah?

**37:57** How much time do we have? How much time do we have left?

**38:03** Okay. So now we are going to talk about some of the...well, this is not...not nice.

**38:09** Uploading data to the Amazon cloud is probably one of the most difficult things. That's why we have a few slides here.

**38:16** There are two types of storage in Amazon Web Services, EBS and S3.

**38:22** EBS is called Elastic Block Storage, and it's ideal for fast access from your EC2 instances.

**38:27** That's where you put your GIS data so it can be accessed from your GIS services and your databases.

**38:33** And then they have S3, which is Simple Storage Service, and that's storage for the Internet.

**38:39** It's used as a backup, so if you have GIS data in EBS, you can create a snapshot and put it into S3...

**38:46** ...and that's very, very, very reliable. Okay? That's where you make your backups.

**38:51** Also, it's kind of a staging place where you can upload data to Amazon S3, and then once the data's in S3...

**38:58** ...you can download it to your EBS drives, which are really the drives that are attached to your EC2 instance.

**39:04** And there are some people who are also using S3 to store their map tiles.

**39:11** Just the map tiles that you create with [ArcGIS] Server, you put them into S3 because S3 has some additional services...

**39:16** ...like this CloudFront we talked about, that accelerate download of map tiles to clients.

**39:23** So let's talk about these EBS drives in your EC2 instance.

**39:26** When you launch the instance, you have two drives. You have the C drive and the D drive.

**39:30** The C drive only has 30 gigabytes, and you should not use it for storing your data.

**39:36** Two reasons. One, it's too small; it's only 30 gigabytes...

**39:39** ...and second, it's a space that the operating system itself uses for temp, for virtual memory, for many things. So don't compromise that space.

**39:49** The other reason is that it's not a safe place because if you terminate your instance, that drive will go away; you will lose it.

**39:57** However, if you go to your D drive, which is called GIS Data, that's a safe place.

**40:03** It's also an EBS drive attached to the machine, but it's bigger. It has 100 gigabytes.

**40:10** And as I said, it's ideal for your GIS data. The machine will access the data very quickly.

**40:15** Think of it like a virtual or a portable...portable hard drive for these data centers.

**40:21** You can attach this EBS drive somehow to different machines. And it's easy to back up.

**40:28** You can literally take an EBS volume and say, Create a snapshot, and that will take that, store it in...in S3 so you can reuse it in the future.

**40:36** If you need additional storage attached to your EC2 instance, you can create EBS drives that are up to 1 terabyte...

**40:43** ...and you can attach up to 12 drives to an EC2 instance on Windows.

**40:51** But by default, as I said, it ships with 100 gigabytes on D:\GIS data.

**40:58** Now that tells you where to put the data for your GIS applications. It doesn't tell you how you get the data into your machine.

**41:06** And we are going to discuss four different options.

**41:09** One you already saw. It works for very small files like this copy and paste or drag and drop with Remote Desktop.

**41:17** It's not a safe way to move sensitive data, but it's kind...it's...it's handy.

**41:22** I mean, when, you know, we have 30 megabytes or 50 megabytes, I just do it that way. It's...it works fairly well.

**41:29** The only problem is that you cannot interrupt this copying of files, like the data's in your clipboard...

**41:34** ...and it's being moved to your machine, if you use Ctrl+C again, it will...the process will be interrupted.

**41:40** So you need to leave the machine alone.

**41:43** So we say around 100 megabytes. I mean, if you go beyond that, who knows? You may need to try a couple of times before you succeed.

**41:51** It's quite convenient.

**41:52** And you can see that if this option is not enabled, when you open the Remote Desktop connection...

**41:57** ...if you go to the Advanced options, you can actually make sure that in fact the Clipboard option is enabled.

**42:05** The other one is FTP. FTP works with much larger files. And there are two strategies.

**42:09** One is to start your ArcGIS Server instance in the Amazon machine and have it access an FTP server that you may have somewhere else.

**42:17** That's pretty...you know, pretty standard. You already have an FTP site, and you just pull the data.

**42:22** Or if you want to do a backup copy or whatever from the Amazon cloud, you can push to the FTP site.

**42:27** The other...other strategy is to set up FTP server in your own Amazon machine, but I want to give you a warning.

**42:35** It's not really safe to enable FTP sites on production systems, so you should do that in an isolated machine.

**42:42** You start a new machine, a small one, you enable FTP, and you put all your data on that FTP site.

**42:50** And then ArcGIS Server is running on a separate EC2 instance. So if people want to attack this instance, it's fine; it's just a staging area.

**42:57** How do you get your data, once it's in this FTP instance, to your ArcGIS Server instance?

**43:03** Well, because the data is in an EBS drive, in your FTP instance you can detach it and then attach it to your ArcGIS Server instance.

**43:12** Or you could simply create a snapshot, S3 snapshot, and then create as many volumes, EBS volumes, as you like...

**43:18** ...to attach them to the...to your other ArcGIS Server instances. Is this making sense? Yeah?

**43:24** So just that warning there; be careful with FTP servers, right, running on your production systems, and that's how you get around the issue.

**43:32** S3 is another strategy to get data into the cloud. So you can use different tools, like S3 Organizer is a Firefox plug-in. Right?

**43:42** And it allows you to pretty much drag and drop data from your local drive into this thing called S3...

**43:47** ...which is this storage for the Internet.

**43:50** And this works well also, like FTP, for very large files. Like if you have hundreds of gigabytes, you could actually go this way.

**44:01** And, well, it always takes time; it's the Internet, right?

**44:04** There are some third-party business partners of Amazon that actually accelerate this upload of data to S3...

**44:10** ...and they are, you know, quite efficient, but of course, this is kind of an additional service.

**44:17** So you get your data into S3 and then with that same S3 Organizer plug-in for Firefox, you can download from S3 into your EBS drive.

**44:27** Okay? So S3 is kind of a staging area for the data that you move to the cloud.

**44:36** Import/Export is an interesting service. Basically, you take your drive, you mail it Amazon...

**44:43** ...to the data center where you have your EC2 instances, and then they plug it, and they...

**44:47** ...they move that information into a S3 so you can download it. That's probably the fastest way to do it.

**44:53** If you are going to do a bulk data loading, kind of I have 1 terabyte of data - Import/Export.

**44:59** Send it, takes a couple of days to get there; one day if you rush, and then you have it in S3 pretty much there.

**45:10** And it works the other way around as well. You can tell them, Why don't you give me all this data I have in the Amazon cloud...

**45:15** ...you send the drive to Amazon, and they put the data into your drive, and they send it back to you for recovery, whatever.

**45:24** Just a few tips. Organize your data for data transfer.

**45:28** Make sure that you are using relative paths to your map documents, to your geoprocessing services, and so on.

**45:33** Eliminate unnecessary data; that's obvious.

**45:35** Compress whenever possible. Compress for the purpose of transfer the less amount of information possible.

**45:42** Break down your compressed files into smaller chunks so if something goes wrong, you can try again but not start all over.

**45:49** Be aware of S3 bucket restrictions. This is well documented in Amazon.

**45:54** You cannot use just any name for your compressed files. You need to be careful with that.

**45:59** If ...if you are transferring terabytes, either you have a lot of time and patience, or you should go import and export.

**46:05** And then create a data transfer strategy. This is something where you will need to deal with moving data to the cloud.

**46:12** Not only with the first initial load but also incremental updates.

**46:16** So just think about how you are doing...going to do the incremental updates.

**46:21** Think about geodatabase replication, for example, as another mechanism to push data from geodatabase to geodatabase...

**46:26** ...straight through geodata services.

**46:34** Yep. Let's talk about elastic deployments of ArcGIS Server. Let me introduce the concept of a custom AMI.

**46:42** You saw how I launched that AMI that Esri provides, it has ArcGIS Server.

**46:46** Now the great thing is that once you have your AMI running, let's say if I go back to my...

**46:56** If I go back here to my ArcGIS Server instance, I could now start ArcGIS Desktop and I could download data to this machine.

**46:59** I could create my own services. I could put my own applications in this machine in the Amazon cloud.

**47:11** And at some point, you have a server that is functional, meaning that it has your own applications, your own services, and so on.

**47:19** And that's the point where you want to save your work, and you do that by creating a custom AMI.

**47:25** So you move all the data here, and then from the management console, you can select your instance and say...

**47:35** ...save or create an EBS AMI. This will save your AMI.

**47:41** You should do this with the instance stopped, so nobody's monkeying around with the instance.

**47:46** You stop it, then you create the instance.

**47:48** And your AMI will display here. Okay? These are AMIs I created for all the demos for the Users Conference.

**47:55** And the great thing is that now that you have your custom AMI, you can...you can come here and you say...

**48:01** ...Launch the instance, and the same wizard that we used before shows up.

**48:06** And I can say, Why don't you create 10 instances of this project A AMI. Right?

**48:18** Once you have these custom AMIs, you can launch as many servers with your services as you like...

**48:24** ...and then you can do things like horizontal scaling.

**48:29** You add 10, 15 servers running in parallel, all running the same services because all...they come from the same AMI...

**48:36** ...and then you pull...you put this thing called Elastic Load Balancer on top.

**48:41** So when a request comes, gets to the Elastic Load Balancer, and the request gets directed to different EC2 instances.

**48:48** That's how you scale horizontally.

**48:51** Custom AMIs can be also used to scale vertically, meaning that once I have my AMI running...

**48:57** ...I can start a small machine and have people hit that.

**49:01** But I can launch a bigger machine with the same AMI and then swap the EC2 instance and say, okay, get rid of this small one; put a big one.

**49:11** Right? It's the same application, but you make the machine bigger as opposed to adding more machines.

**49:20** Okay. So this slide is important 'cause it tells you how you scale up ArcGIS Server in the Amazon cloud.

**49:27** Traditionally, you add more SOC machines. In the Amazon cloud, you don't do that.

**49:33** You don't want DCOM stuff going across the data center.

**49:37** You want self-contained machines that have everything they need to run for the Web, the SOM, the SOC, and the data.

**49:43** Actually, the data might be living in a separate EC2 instance running ArcSDE with the geodatabase AMI.

**49:49** But your map services will point to that.

**49:52** And then once you have that, you have these clones, this cluster of peers, where all the machines have the same services...

**49:59** ...and then you put Elastic Load Balancer on top.

**50:02** Yep. That's how you scale in the Amazon cloud. And it's fairly efficient if you work well with these custom AMIs.

**50:11** A machine goes down; don't even remote desktop to that machine.

**50:14** In fact, in production, when you set the security groups, you probably don't even want to enable Remote Desktop.

**50:19** You just open 80P...HTTP 80. Just don't remote desktop. Just kill the machine, start a new one, and put it under the ELB.

**50:26** That's how you fix things, right, in the Amazon cloud, because that AMI is like, yeah. Ten minutes, boom.

**50:32** Your machine is up and ready, and then you put it into the cluster.

**50:37** So this leads to the concept of Elastic Load Balancer.

**50:40** As you can imagine, it's just a load balancing that works in the elastic cloud.

**50:44** It supports sticky and nonsticky sessions, so you can imagine in this configuration, if you have an asynchronous geoprocessing service...

**50:51** ...unless your client is always going to the same machine, you will lose the job, right?

**50:56** So for that, you need to enable sticky sessions.

**51:00** Unfortunately, sticky sessions don't work with ArcMap yet because the stickiness works with browser cookies...

**51:07** ...and ArcMap doesn't use browser cookies.

**51:08** So we'll try to fix that.

**51:10** But if you are using Web browsers to run asynchronous processing, it will work fine as long as you set sticky sessions.

**51:17** Let me give you a quick demonstration with...I think we're running out of time, right? Yes. No demonstration. No demo. Okay.

**51:29** CloudWatch is that service I talked about which allows you to monitor.

**51:33** Very important for this environment because you may want to add and remove instances from your ELB as the load goes up and down. Okay?

**51:43** And things to remember. I already talked about these issues with sticky sessions and...and so on.

**51:51** No demo. Maybe at the end if we have time.

**51:55** Anyways, before...before I...we jump into the real scenarios thing, we kind of followed a workflow through the presentation.

**52:04** This is what we call the Amazon EC2 workflow. First you launch the ArcGIS Server instance, then you take ownership of it...

**52:11** ...you authorize it, you change the password, you add your services, you add your data.

**52:16** Then you do your functional testing on that AMI; you test how the application works.

**52:21** And if it works, you save that work as a new AMI.

**52:24** And then you launch that AMI in a small machine, in a big machine, and you do nonfunctional testing, performance testing.

**52:32** That will give you an idea of how many concurrent users a machine of type A or type B can hold.

**52:38** That's information you need when you need to scale up. That's what we mean by measure the throughput.

**52:44** And then you just put your machine in production; you have an estimate of the load, you know

how much load a machine can hold.

**52:52** And then if things go wrong because you have a lot of users, you can enable CloudWatch and add more machines to that deployment.

**53:01** Or remove them. Almost as important as add is to be able to...to remove.

**53:05** So, many people actually...we are releasing this with ArcGIS 10 as you know...

**53:10** ...but we have been working with ArcGIS Server on Amazon for quite a while.

**53:14** So I'm going to let Marwa talk about some of these experiences from real clients.

**53:21** Thank you, Ismael. [Unintelligible] the other part.

**53:30** So I hope everybody's familiar now with these concepts for Amazon.

**53:35** What I'm going to talk about next is a few examples of real projects that are hosted right now in the cloud environment.

**53:44** They're out there happening today. I always have this problem, so I'll just hold it.

**53:52** So what I'm going to talk about next is a number of projects and applications that are running in the Amazon cloud.

**54:00** They're hosted there today based on the AMIs that Ismael has been talking about...

**54:05** ...using the same kind of services and concepts he's explained already.

**54:11** So I want to...I...I just have a few examples.

**54:14** We actually have a few more running today, but I wanted to use these because each of them has a unique story.

**54:21** This application, UDS Mapper, was developed by a partner of ours; their name is Blue Raster.

**54:27** They've built a very nice application, and they've set up the data behind it.

**54:32** And what they wanted from us to help them with was hosting the ArcGIS Server part.

**54:38** The application had a unique life span. It was intended to be running for...actively running for a period of 90 days.

**54:47** It was going to be public facing. They expected there would be heavy load during the first part of the hosting period and the end of it...

**54:55** ...but not really very heavy during the middle of that period. And due to this variation, and being a public site...

**55:05** ...it was very hard to predict how many people were really going to come and use it at a certain time.

**55:10** So from that perspective it wasn't really clear.

**55:12** Should we buy very large machines or very small machines, and how are we going to handle that?

**55:17** So the cloud seemed perfect due to the limited duration as well as due to nature of it being public facing...

**55:24** ...and we weren't really sure how this is going to scale.

**55:28** So for this application, what we decided to do was host both the application and the tile data content on Esri.

**55:39** And then we started ArcGIS Server instances based on a custom AMI that we created for that specific project.

**55:45** So we just launched a couple of...an instance from the original ArcGIS Server AMI, set it up with the data...

**55:53** ...as well as the service, and then created a custom AMI for that.

**55:57** And from there, we launched two instances, set them up with the Elastic Load Balancer, and that was it.

**56:04** The application was a Flex application that is using both the S3 tiles as well as the services that are being published by ArcGIS Server...

**56:12** ...and it's running on the Internet. It's available for everybody to use.

**56:17** Now, because of the nature of being public facing, what we decided to do was to start with smaller-sized instances...

**56:25** ...a couple of them, because we wanted it to be high availability, so if one fails, it kind of goes to the other one.

**56:31** So we started with two standard large, but what we did was we used the CloudWatch to kind of monitor the utilization of these instances...

**56:39** ...and if we saw that the utilization was going beyond a certain extent, we were prepared to go and start more instances.

**56:46** So we decided to go with the horizontal scaling, and that kind of worked really well for that project.

**56:55** Another example is another project we worked on with USDA.

**57:01** And in this case, USDA was actually hosting a number of applications already themselves, so they felt pretty confident about how things ran.

**57:09** However, the problem they was after press releases, especially big ones like ones that go on the White House website...

**57:17** ...or something that goes on CNN or other news channels as a big application, they found that for a very short duration of time...

**57:25** ...they would get very high loads, and that caused a lot of issues and problems.

**57:31** Then after that, the loads go back into very low, normal ones, maybe not even a lot of usage.

**57:38** So this seemed to be like the perfect opportunity to leverage the cloud.

**57:42** What they wanted to do, rather than invest in a major infrastructure, to meet public demand at these peaks...

**57:48** ...was to have a cloud setup that will go high when we need that load and then when we don't need it, we shut down all these extra machines.

**57:57** So what we did in this case, again, we started an instance from the AMI, set it up with the data...

**58:02** ...and this time put the application as well on this instance, and then saved the custom AMI.

**58:08** And from the custom AMI, we launched two instances that were registered with the Elastic Load Balancer.

**58:14** The application is using ArcGIS Online for the tiles, the basemaps and between the application being hosted...

**58:22** ...and the ArcGIS Services being hosted in the cloud and using the tiles from ArcGIS Online, they have the application running in the Internet.

**58:30** And whenever we expect a certain load, they notify us there's going to be a press release.

**58:35** So we keep an eye on it, and if we see that utilization exceeds a certain limit...

**58:39** ...that's when we go and prepare ourselves to launch more instances based on these custom AMIs.

**58:44** So the beauty of it is that in a matter of 15 to 20 minutes, if there is a load, we have the custom AMI...

**58:51** ...we can go and add more instances, two or three machines depending on how much we need.

**58:57** And when we don't need it, we just shut it down.

**59:02** Also, I thought this would be a good example. This is the public-facing PB [sic] application that's hosted by PB on their website.

**59:10** The story for this application, I'm sure everybody's pretty familiar with what happened for this pretty sad disaster.

**59:18** But what we did was we were on the phone with them around three p.m. one day...

**59:22** ...and they said that they really needed to go and put an interactive application on the website that the public can get access to.

**59:31** And they wanted to make sure this information was publicly available at any...in an application where users can do some analysis themselves.

**59:40** They've already had this data on their website as PDF files, so what they wanted to do, rather than just keep publishing the PDF files...

**59:45** ...they also wanted in addition to that an application to show the content.

**59:53** And it was, of course, pretty quick turnaround. They wanted to see as soon as possible a prototype application up and running.

**1:00:02** And what was really great was that within that same night... So we kind of finished the call around three-thirty.

**1:00:10** In a matter of a couple of hours or three hours later, we actually had that server up and running with the prototype application.

**1:00:17** And they were hitting on it and kind of thinking, oh, you know, we need to change a few things here and there...

**1:00:24** ....and they started modifying a few things.

**1:00:26** And within less than a week, we had this up and running and it was available for public.

**1:00:32** There was a few challenges related to this application.

**1:00:37** One of them was the data was constantly being updated.

**1:00:42** What we kind of opted for to start things really quickly was to set up an FTP server on a staging machine...

**1:00:49** ...where we go and upload the data updates to that server, and from there disseminate it to all the production machines.

**1:00:57** They also wanted to make sure this application is going to be public facing so we can scale it depending on demand.

**1:01:04** So we started pretty big instances, so this is where we went to start with vertical load matching.

**1:01:14** So we started pretty...pretty big instances, about a couple of them, and then added a few more...

**1:01:21** ...and then we keep them monitored as well.

**1:01:23** When we see utilization is going to be needing more instances, we can add more, depending on the custom AMIs.

**1:01:29** In this case, the data is not part of the custom AMIs due to its nature of being updated all the time...

**1:01:36** ...so the custom AMI was mainly targeting setting up ArcGIS Server with the number of services in addition to the application.

**1:01:43** Then after we launched these instances, we go and get the latest update and make sure they are up and running...

**1:01:49** ...before we make them active behind the load balancer.

**1:01:53** And that's one of the nice uses of the load balancer is you can have your instances up and

running...

**1:01:57** ...but that doesn't mean they're accessible.

**1:01:59** You can detach them and attach them to the load balancer depending on when you think it is the right time to make them available for end users.

**1:02:07** So while we're doing the updates, we detach them from the load balancer, and then after we're done, we register them again.

**1:02:15** This was kind of done in a very big hurry.

**1:02:17** What we're moving to right now is we're starting to think about doing replication using Enterprise Geodatabase AMI...

**1:02:24** ...so what we're going to do is we're going to replicate with the BP database...

**1:02:29** ...get updates, a more real-time nature, to that staging server, and from there disseminate it to the production machines.

**1:02:38** So that's pretty much the number of examples I wanted to highlight quickly and next we want to do the QA.

**1:02:46** You want to do that, Ismael?

**1:02:47** Actually, just throw any questions that you may have. Yes?

**1:02:53** [Audience question] [Inaudible] support [inaudible] enterprise for SQL Server Express on the AMIs?

**1:03:01** Okay. So the question is if we're going to provide AMIs with SQL Server Enterprise and ArcSDE?

**1:03:07** No. The issue we had was...is not a technical issue; we could do that, but we...

**1:03:12** We are working with Microsoft and with other database providers so we can create AMIs for you.

**1:03:17** The issue is that we cannot right now redistribute anything but SQL Server Express and Postgres in our AMIs.

**1:03:25** However, you can...you can install ArcSDE yourself from the cloud and create your custom AMI.

**1:03:33** But one note there. When you run ArcGIS Server in the Amazon cloud with the AMIs we provide, you get full tech support...

**1:03:39** ...meaning that you can call Tech Support, they will take your calls, and if you have any issues, we will be able to repro.

**1:03:44** Because we have the environment you are running on.

**1:03:46** But if you install ArcSDE on your own, then we don't have that environment set; we have not certified that AMI.

**1:03:53** So you can still call and we'll try to help you...

**1:03:55** ...but if we cannot repro on one of the certified environments, we won't be able to address the issue.

**1:04:02** So we are working on that, but right now, only Postgres database with the AMIs. Questions? Yes?

**1:04:12** [Audience question] Yeah. That [unintelligible] process of starting up your instances [unintelligible] can be automated in any way?

**1:04:20** Okay. Can you automate the process of launching and dropping AMIs? Yes.

**1:04:26** As we said in the first slide, all of this is accessible through Web services.

**1:04:30** It just happened that we were using a Web application to do things manually, but there are toolkits for Java, for .NET.

**1:04:38** They have the SOAP web services, and you can automatically, you know, monitor your machines, attached drives...

**1:04:45** Scaling service. There's also the Auto Scaling Service, so this is a service that comes with no extra charge.

**1:04:52** They only charge you for starting the instances and the monitoring.

**1:04:55** And it will use the CloudWatch to monitor the utilization on the instance and automatically start more instances as needed depending on a criteria.

**1:05:03** Like for example, the processor is 70 percent utilized, so go and launch another instance based on this AMI.

**1:05:10** It also can...allows you to limit the number of instances you will need to start so you don't end up with a bill that's completely unexpected.

**1:05:20** [Audience question] Can you tell us something about the licensing on the Esri side?

**1:05:24** Okay. The licensing on the ArcGIS side. Okay. Good question.

**1:05:28** So people say, oh, this idea of elasticity is great because if I don't use the server, I don't have to pay, and that is true for hardware.

**1:05:38** You have to license ArcGIS Server. Okay? So if you start 10 servers with four cores, you need 10 licenses of ArcGIS Server.

**1:05:48** Now, even if you eliminate the machines, you have already paid for the software.

**1:05:54** So what we have done is, right now, you can...you can buy a subscription...you can license ArcGIS Server for one, three...

**1:06:05** Sorry. ...for 3, 6, and 12 months. So it's not an hourly rate; it's a monthly rate for now.

**1:06:12** I have to say that this ... the elasticity with ArcGIS Server on the Amazon cloud realistically only comes true...

**1:06:21** ...when you have an enterprise license agreement where you have unlimited access to ArcGIS Server licenses.

**1:06:27** If you have just an ArcGIS Server license, you may want to run in Amazon because it's quick to set up...

**1:06:32** ...because you don't need to fight...you don't need to fight your IT department, but not because of elasticity.

**1:06:40** Because you will have one license.

**1:06:41** So you can have ArcGIS Server running, and it's pretty affordable and it's pretty good, but the elasticity...

**1:06:47** ...either ELA or this monthly termed licenses of ArcGIS Server. Good question.

**1:06:54** [Audience question] Well, if you throw money at it, you could scale up too.

**1:06:56** Oh, of course. Let me...let me honest...let me be honest with you. Forget about the ELA; just purchase the ArcGIS Server licenses.

**1:07:03** So if you start 50, just 50 licenses. I'm just kidding.

**1:07:09** [Audience question] So you talked about the benefits. Besides loading data, what are the downsides of using Amazon?

**1:07:17** The downsides of using Amazon. Well, I guess one you just highlighted. It's like, moving data to the cloud happens frequently over the Internet.

**1:07:26** So if you...this bulk...this bulk update of data can take time and effort.

**1:07:33** And then for incremental updates, if you don't...if you don't have big incremental updates, this geodatabase replication works well.

**1:07:40** The FTP thing that you described.

**1:07:42** I mean, you can get around, but if you have, you know 100 gigabytes of data changing every day is...is really a big challenge for you to do that.

**1:07:51** I'm not sure if you want to add.

**1:07:53** There's a few more. I mean, for example, folks that work with Image Server, that is, we're talking about terabytes and terabytes of data...

**1:08:01** ...and management and updates, it is...it is kind of a challenging process.

**1:08:05** We're trying to figure out what is the best way to deal with it now so we can come up with a few best practices.

**1:08:09** [Inaudible audience question]

**1:08:10** Yeah, and kind of how to move the data around, how to manage the updates and things like that.

**1:08:15** Also worth mentioning, because of the nature of the cloud and it being a shared environment, we've seen sometimes the performance...

**1:08:24** ...you can benchmark it to be a certain performance, and then when you run it on another day, you can see that vary a little bit.

**1:08:32** So these type of variations in the performance, the capacity of these resources has been something to keep in mind when you're planning.

**1:08:42** They do have different models between the on-demand, reserved, and the spot and the different regions...

**1:08:48** ...so there are variations between the different regions in terms of their capacity, in terms of their performance, and things like that.

**1:08:55** So that kind of comes with some experience.

**1:08:58** And, you know, a lot of people share their experiences online. so you can always learn about that as well.

**1:09:05** So just to give you an example, we uploaded, right before the User Conference, 1 terabyte of imagery to Amazon.

**1:09:12** It took, I think, four days to get the 1 terabyte of data from...from my machine into the EBS drive.

**1:09:19** [Audience question] With a DSL connection?

**1:09:21** I don't know exactly what we...what we used.

**1:09:23** Yeah. It's one of the tools [unintelligible].

**1:09:27** Okay. We can give you...we should connect with the image service guys to just give you the details.

**1:09:31** Another example is the same machine, the same data, the same testing, 30 percent variation in performance.

**1:09:38** Why? It wouldn't...it wasn't the Internet, because we were testing from inside the Amazon cloud; it was the machine.

**1:09:43** The machine itself was not as fast as the day before. And then the day later, it was faster. Thirty percent variation.

**1:09:50** [Inaudible audience question]

**1:09:51** We know...Amazon told us that if you have an eight-core machine, and eight-core virtual image, you are going to own one physical machine.

**1:09:59** So the machine is yours, and the performance is going to be pretty consistent.

**1:10:03** If you purchase or you rent one of these smaller machines, you are sharing the machine with others and...

**1:10:09** Well, it depends on how busy the other guys are and also...

**1:10:13** Although these virtual machines are all sharing kind of the same port to communicate, and that's where the bottleneck happens.

**1:10:19** So when you are renting small machines, you may observe these variations in...in performance.

**1:10:24** And between regions depending on the [inaudible].

**1:10:26** Of course, between regions, yeah. Between...yeah, real...yeah.

**1:10:29** So always try to do performance testing and measure to what extent...what each EC2 instance can give you. Right.

**1:10:39** [Inaudible audience question]

**1:10:46** Well, actually, no. We...we published a white paper not long ago; it's called Building High-Capacity Map Services.

**1:10:54** It's in the white paper section of esri.com.

**1:10:57** And it explains how to use network...network load balancers with ArcGIS Server following this peer architecture...

**1:11:03** ...not using the Web SOM/SOC kind of traditional thing.

**1:11:06** And it has many advantages to it. So this is for the Amazon cloud, but it also works perfectly on premise.

**1:11:12** In fact, we have people using VMware on premise with virtualization in VMware.

**1:11:17** Kind of the same technique. You know, I create my image, I start it, put ELB on top, and I just forget about SOMs and things like that.

**1:11:26** [Inaudible audience question]

**1:11:27** Well, read the white paper. There are good things and bad things about them.

**1:11:31** But it's definitely a good pattern for some use cases.

**1:11:35** [Inaudible audience question]

**1:11:44** Okay. So extensions. When you get the AMI, we have installed pretty much everything in there...

**1:11:49** ...and then when you use your license file, you unlock one feature or the other. Right?

**1:11:54** So all the extensions work, including Image [Server] extension, which as you know is part of ArcGIS Server in 10.

**1:12:01** There are some...exceptions. JTX is not installed and Geoportal toolkit is not installed.

**1:12:10** [Inaudible audience question]

**1:12:18** Yeah, we use Amazon ourselves. If you look at the effort...relief effort kind of websites for the earthquake in Chile and the Haiti thing...

**1:12:28** ...everything was running on Amazon for some of the reasons we explained.

**1:12:32** And also, ArcGIS Online itself, our service is running Amazon machines using this technology that you saw. Yes?

**1:12:44** [Audience question] Amazon charges by the hour for these machines to be up. How do you address the situation...

**1:12:49** ...where you have a service that needs to be available but doesn't need to be available all the time?

**1:12:54** [Unintelligible] pull licenses so that they can share [unintelligible]?

**1:13:01** Do you understand the question? We don't understand the question. Sorry.

**1:13:09** I think...Is it a licensing question regarding if you want to run multiple services on the same server?

**1:13:16** [Audience question] No, but... You don't want to run the server 24 hours a day for a user to be able to...

**1:13:22** ...you know, randomly throughout the day be able to [unintelligible] service.

**1:13:25** But could you pool other users and start up one machine and use it to run your own services on [inaudible]?

**1:13:37** Once you start the machine, you have a machine running with ArcGIS Server for just four hours a day.

**1:13:42** You will pay for four hours of hardware and the storage that you need and all of that.

**1:13:46** From a software point of view, you have to license an ArcGIS Server license for four cores...

**1:13:50** ...regardless whether you use it 24/7 or just four hours a day. That's how it stands right now, okay?

**1:13:58** The granularity is at this...you know, we...we mentioned this 1-month or 3-month, 6-month, 12-month subscriptions.

**1:14:06** That's the level of granularity we can give you right now.

**1:14:09** Hourly rate for ArcGIS Server, right now we...we are not there.

**1:14:16** If we didn't get your question right, you can just come over at the end. Any other questions? Yes.

**1:14:24** [Inaudible audience question]

**1:14:32** Yeah. I mean, not only 8, but 24. If you have a machine with 24 cores ArcGIS Server, it's going to be beauty in there.

**1:14:40** [Inaudible audience question]

**1:14:47** Yes. The question is about licensing.

**1:14:49** As we...we mentioned before, if you do outdoor scaling, you need to license every ArcGIS Server that you are running.

**1:14:56** So if you, at some point, start 10 servers with four cores each, you need 10 licenses of ArcGIS Server.

**1:15:02** That's where we mentioned the idea of the enterprise license agreement. That's the way to go for elastic.

**1:15:15** Ok. Thank you very much. If you have additional questions, just come over. Thank you very much for your time.