National Center for Education Statistics

Amazon Hosted ESRI ArcGIS Servers

Project Final Report

National Center for Education Statistics

Description of Application

Operating Organization

The US Department of Education's (ED) The National Center for Education Statistics (NCES) is the primary federal entity for collecting and analyzing education-related data. NCES uses ESRI technologies to provide geospatial context to education data.

NCES participated in the FY 2012 Federal Geographic Data Committee (FGDC) GeoCloud Sandbox Initiative by hosting four ESRI ArcGIS Servers and supporting database servers in the cloud environment. ESRI's ArcGIS Server technology is a server based platform for centralized management of geospatial web services as well as geospatial data in relational database management systems. ED uses ArcGIS Server for its School District Demographic Data System and the Public School Boundary Collection and Verification Project.

Our objective for the project was to evaluate the cloud computing environment as a viable option for hosting publicly accessible education web applications and data successfully. The Department of Education evaluated performance, reliability, cost, and ease of deployment factors.

Community of Interest

The community of interest for this project was ED, NCES, and visitors to the NCES website. Visitors to the site consisted of private companies; students; researchers; governmental entities; and the public at large.

Operating System and Software Requirements

Windows 2008 R2 servers and the following software:
ESRI ArcGIS Server and Web Adaptor
ESRI ArcGIS Desktop
ESRI ArcGIS License Administrator
SQL Server Express, 2008, 2012
Microsoft Internet Information Server

Operational Requirements

ED hosts their website and supporting IT infrastructure in-house at the NCES Data Center which is managed by ED staff. There is one production ArcGIS server and one staging ArcGIS server that hosts the mapping applications and mapping web services. The in-house host production server has ArcGIS version 10.0 sp1 on a 32 bit Windows operating system.

ED management applied to the GeoCloud program to explore an initiative to move its GIS infrastructure to the more rapidly scalable platform that the cloud provides. This saved NCES the need to upgrade the existing infrastructure. In order to maintain sustainability of the application and to run ArcGIS Server 10.1, the server infrastructure at NCES would need to be upgraded to a 64 bit operating system requiring NCES to procure new hardware and software which would have taken several months to be fully functioning. The GeoCloud-II program was very timely with the new release of ArcGIS Server 10.1 in early 2012. Using GeoCloud-II instead allowed NCES to immediately realize the benefits of Cloud deployments. This includes increased efficiency in cost and time for server/systems provisioning and application deployment as well as improved support for testing or staging environments as the suite of tools grows.

ED specifically evaluated the experience of configuring and deploying servers in the cloud environment and getting a realistic estimate of costs, reliability, support options, and performance.

Deployment in the Cloud

ED used two existing ESRI ArcGIS Server images, ami-149d387d (ArcGIS 10.1 Server with SQL Express) and ami-c89d38a1 (ArcGIS 10.1 Server with SQL Server), released for production in June 2012 for the GIS software. For our sandbox testing of the application installations, we used two m1.Large instances. After migrating the servers to production, ED has a total of four GIS servers and two database servers hosting NCES supporting data as described in the table below. The AWS environment was not configured for auto scaling and load balancing.

The AMIs were launched from the AWS console. Specific ports were opened in the security group based on static IPs. Ports 27001 and 27000 were opened for the license server and port 3389 was opened for remote desktop use. Port 1433 was

opened for SQL Server and SQL Server Express. For the ArcGIS Servers, authorization files were obtained from the ESRI Customer Care Portal. The web adaptor was installed so that the mapping services are available through IIS and port 80. The data for the map services was transferred over to the cloud servers using Cloudberry. The map services were published using the ArcGIS desktop software installed as part of the AMI. On the database servers, SQL Server licenses were transferred from an existing purchase.

Operations in the Cloud

A description of the EC2s is listed in Table 1. The specifications for the instance types are in Table 2. Maintenance efforts were low once backups were scripted. Maintenance involved upgrading an instance type, increasing disk space on volumes, adding new volumes, opening ports for additional IP addresses at the security group level on the AWS console, upgrading the EC2Config service on one EC2, and installing windows updates. Windows updates were set to download automatically but install is manual. Restart is usually done when the weekly backup is performed.

Table 1: EC2 Descriptions

Server	Description	Region	Type	Volumes
SDDS - Production Server	AGS 10.1 sp1 Win 2008 R2 EC2	us-east-1d	m1.large	C drive 35GB D Drive 100 GB
School Boundaries - Production	AGS 10.1 sp1 Win 2008 R2 EC2	us-east-1d	m1.large	C Drive 45 GB D Drive 100 GB
SDDS – Production databases	Windows_Server-2008-R2 64 Bit Base	us-east-1d	m1.xlarge	C Drive 60 GB M Drive 1 TB T Drive 40 GB S Drive 60 GB Z Drive 1000 GB R Drive 1000 GB
SDDS - Production databases	Windows_Server-2008-R2_SP1-English-64Bit-Base-2012.10.10	us-east-1d	m1.medium	C Drive 30 GB D Drive 300 GB
School Boundaries - Staging	AGS 10.1 sp1 Win 2008 R2 EC2	us-east-1b	m1.large	C Drive 45 GB D Drive 100 GB
SDDS - Staging Server	AGS 10.1 sp1 Win 2008 R2 EC2	us-east-1b	m1.medium	C drive 35GB D Drive 100 GB

Backups were scripted using modified Powershell scripts available at http://messor.com. The scripts run from a remote Windows 2008 micro instance using the Task Scheduler. A weekly backup is performed that consists of creating an AMI of each server. The EC2 is shut down and the AMI is created. The EC2 is then

restarted and the IP reattached. Daily backups are performed that consist of taking snapshots of the volumes for each server. During the backups, cleanups are performed that remove older images and snapshots.

Table 2: EC2 Specifications

M1 Large Instance	M1 Medium Instance	M1 Extra Large Instance
• 7.5 GB memory • 4 EC2 Compute Units (2 virtual cores with 2 EC2 Compute Units each) • 850 GB instance storage • 64-bit platform • I/O Performance: High	• 3.75 GB memory • 2 EC2 Compute Unit (1 virtual core with 2 EC2 Compute Unit) • 410 GB instance storage • 32-bit or 64-bit platform • I/O Performance: Moderate	• 15 GB memory • 8 EC2 Compute Units (4 virtual cores with 2 EC2 Compute Units each) • 1,690 GB instance storage • 64-bit platform • I/O Performance: High

Basic monitoring of average CPU utilization, max bytes network in, max bytes network out, and average disk reads is available at the AWS console level but no additional Amazon monitoring features or alarms are enabled (Figure 1). Pingdom is used to test the ArcGIS Server EC2 that supplies the map services for the SDDS map viewer.

Figure 1: AWS Console Monitoring Interface

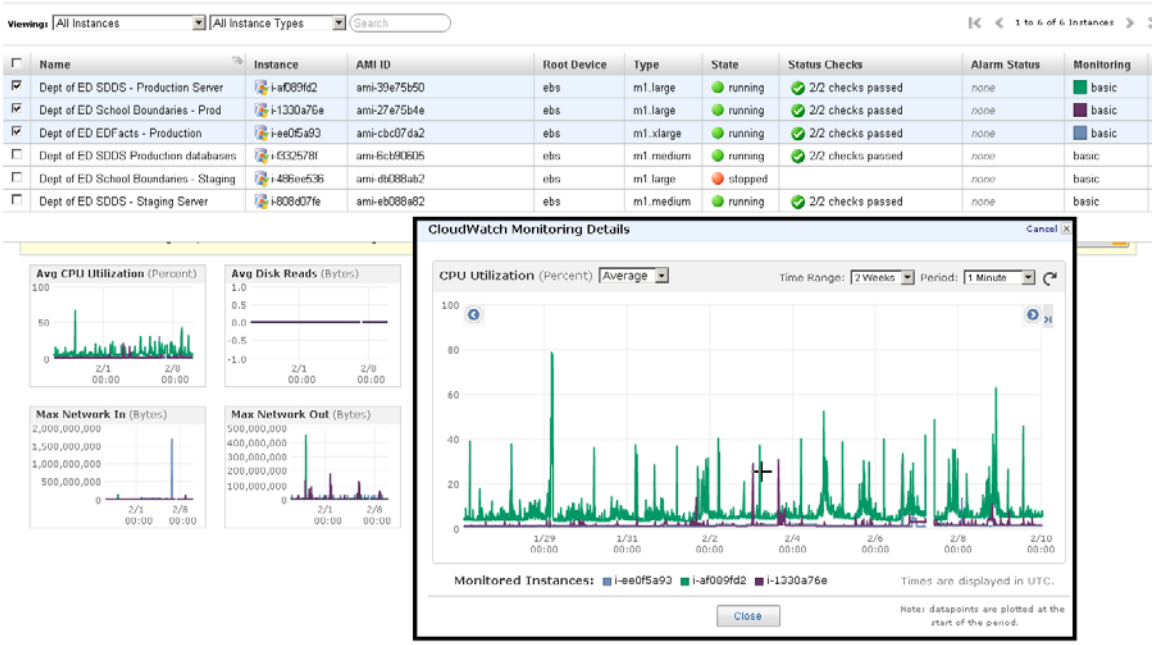

Operational Cost Comparisons

Table 3 shows the actual operating costs over the 6 month period operating in the production GeoCloud account. By January 2013, ED had launched all servers that were originally planned so the January and February costs leveled. In February, one server was turned off.

Table 3: Monthly Costs in GeoCloud Production Account

Month Year	Total (Dollar)	Amazon EC2s (# servers)		Amazon EBS	AWS Data Transfer	Load Balancer	S3	Cloud Watch
		Hours	Costs					
Sept 2012	**922.50**	660 (2)	891.90	29.23	0.00	1.38	0.00	0.00
Oct 2012	**1581.67**	2202 (4)	1450.53	123.24	0.53	2.27	5.00	0.90
Nov 2012	**1528.84**	2880 (4)	1363.68	159.84	4.29	0.00	1.02	0.00
Dec 2012	**2376.69**	3863 (6)	2107.70	266.50	2.12	0.00	0.36	0.00
Jan 2013	**3160.67**	4396 (6)	2752.37	402.61	3.41	0.00	2.28	0.00
Feb 2013	**2344.08**	3356 (5)	1880.56	451.43	3.56	0.00	8.53	0.00

Cost savings can be achieved by purchasing reserved instances to lower the charges paid for hourly compute capacity. Hourly compute capacity means that charges are accumulated based on the operating hours of the servers.

Table 4 displays the cost comparison per server for on-demand pricing and light and medium utilization reserved instance pricing. It also displays the number of operating months before savings are realized on the purchase of a reserved instance.

Table 4: Projected Operating Costs for 1 Year 24 hrs 7 Days/Week*

Pricing Category (# EC2s)	Windows instance-hour ($)	Upfront Cost Per Server ($)	Total/Server 720 hrs/mo 1 Yr Term ($)	Break Even Point (Month)
M1 Medium (2)				
Reserved Pricing (Medium Utilization)	0.088	320	1080.32	3.13
Reserved Pricing (Light Utilization)	0.118	138	1157.52	1.71
On-Demand	0.23		1987.2	
M1 Large (1)				
Reserved Pricing (Medium Utilization)	0.175	640	2152	3.12
Reserved Pricing (Light Utilization)	0.235	276	2306.4	1.70
On-Demand	0.46		3974.4	
M1 Xlarge (1)				
Reserved Pricing (Medium Utilization)	0.35	1280	4304	3.12
Reserved Pricing (Light Utilization)	0.47	552	4612.8	1.70
On-Demand	0.92		7948.8	
M1 Large SQL Standard (2)				
Reserved Pricing (Medium Utilization)	0.571	1670	6603.44	5.76
Reserved Pricing (Light Utilization)	0.624	1370	6761.36	5.44
On-Demand	0.974		8415.36	

* Assumes 720 hours in a month

From Table 4, hourly compute capacity costs (Running all 6 EC2s 24/7 for the whole year) at on-demand pricing is approximately $32.7K. A purchase of six (6) 1-yr term medium utilization reserved instances costs approximately $21.8K which is a savings of $11K. A purchase of six (6) 1-yr term light utilization reserved instances costs approximately $22.8K which is a savings of $10K.

Table 5 sums the operating cost values from Table 4 and the actual storage and data transfer values from February 2013 in Table 3 to calculate the approximate cost comparison among the different Amazon rates over a one year term for all 6 servers. One could expect to pay approximately $38.2K for the year at the on-demand pricing rate or $27.3K with the purchase of medium utilization instances. These totals include the hourly compute capacity for 6 servers and also include charges for EBS, S3, and Data transfer.

Table 5: Cost comparison of Amazon EC

	Cost		**Amazon EC2**
Fixed (one-time) cost			None
Monthly	Network cost		$4 for data transferring / month
	Storage		$451 /month
	Operating Costs	Not reserved	$2727 / month
		Reserved (medium)	$1819 / month
	Maintenance cost		None
Yearly	**Total**	**Not reserved**	**$38,184**
		Reserved	**$27,288**

Issues and Lessons Learned

Overall, deploying servers in the cloud environment was relatively easy. The reliability is slightly uncertain based on a severe weather event experienced during

the months of operation. Ultimately no server was lost but there were a few hours of inaccessibility. Regardless, servers in the cloud should never be considered fail-safe and backup strategies must be well planned and executed.

The best cost savings can be achieved by purchasing reserved instances for production servers operating in the cloud.

The greatest challenge still lies ahead. Where ED's goals fell short was being able to use a security-hardened AMI template, suitable for Government Agency use, when launching the instances.

Only a couple of weeks before termination of the GeoCloud-II program, a security-hardened template became available for the Windows 2008 R2 Operating System with ESRI's ArcGIS Server 10.1 and the SQL Express database. ED tested the template but at this time reconfiguring the 6 servers from this template is not an option because it would require a lot of configuration to get back to the current operational status. The template is only applicable to two of ED's servers. Also ED is using an updated version of ESRI's ArcGIS Server 10.1 that includes service pack 1. Using CloudBuilder, ED can't apply its custom security group to a site that is launched with this template. Not being able to apply an agency's own security group will be a problem for other agencies as well when trying to develop a common and re-usable template.

Using ESRI Cloud Builder and this template would be beneficial to those just starting out in this program. For already-configured EC2s, it would be very helpful to have access to the scripts and knowledge base used to apply the security policies to the Windows 2008 server. In that way, more thorough documentation would be available when writing the security plan and the individual agencies could gain experience in applying policies required for FISMA compliance. When alternate configurations are needed, such as ED's other 4 servers, the existing policies could be tweaked and expanded upon.