



**DavidChappell**  
& Associates

# GIS IN THE CLOUD

THE ESRI EXAMPLE

DAVID CHAPPELL

SEPTEMBER 2010

SPONSORED BY ESRI

## CONTENTS

<b>Contents</b> .....	<b>2</b>
<b>Cloud Computing Basics</b> .....	<b>3</b>
Cloud Applications and Cloud Platforms .....	3
An Example Cloud Platform: Amazon Web Services .....	4
<b>GIS and Cloud Computing: Examples from ESRI</b> .....	<b>5</b>
ArcGIS Server on Amazon EC2 .....	5
ArcGIS.com .....	7
ArcLogistics .....	9
Business Analyst Online .....	10
<b>Conclusion</b> .....	<b>11</b>
<b>About the Author</b> .....	<b>12</b>

## CLOUD COMPUTING BASICS

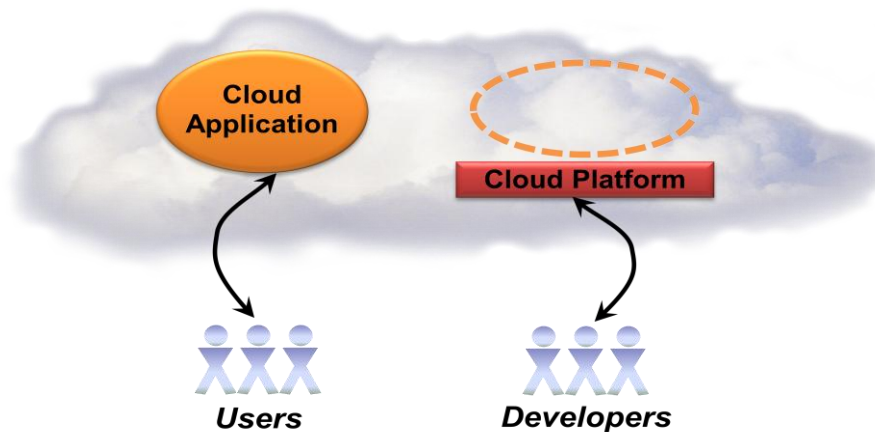
The rise of cloud computing just might be the most important thing happening in information technology today. While not everything will move into the cloud, it's fair to say that nearly every organization will use this new approach in some way. Given this reality, cloud computing has the potential to change many aspects of our world.

For users of geographic information system (GIS) technology, the cloud opens a number of new possibilities. But what are those possibilities, and why might they be better than what you're doing now? In fact, why should you care about cloud computing at all?

The goal of this paper is to answer these questions. After a quick introduction to cloud computing, we'll look at some concrete examples by exploring what Esri is doing today to bring GIS to the cloud.

## CLOUD APPLICATIONS AND CLOUD PLATFORMS

The phrase "cloud computing" means different things to different people. This discussion focuses entirely on public clouds, in which code and data live in Internet-accessible data centers owned by Amazon, Microsoft, Google, and others. In this context, it's useful to distinguish between two broad categories: cloud applications and cloud platforms. Figure 1 illustrates this idea.



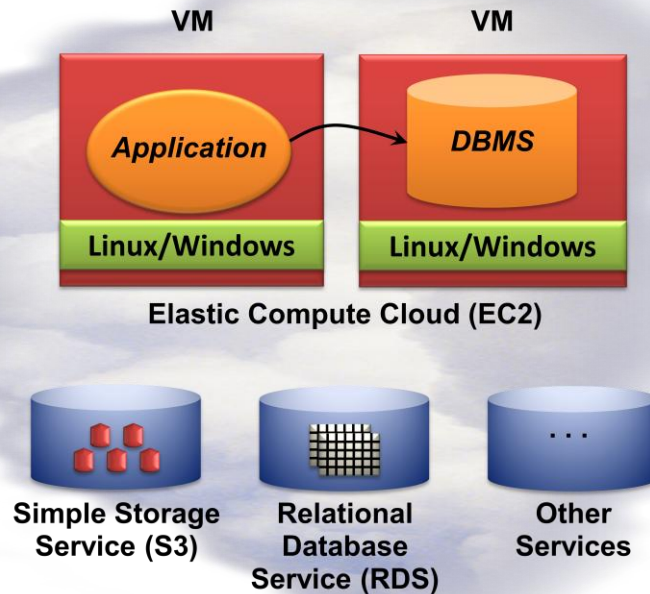
**Figure 1: Cloud applications are accessed by users, while cloud platforms are used by developers.**

Like the applications that run in your data center (sometimes referred to as *on-premises* applications), cloud applications have users. The big difference is that unlike on-premises applications, cloud applications run in a shared data center that's accessed via the Internet. This application style is sometimes called *Software as a Service* (SaaS), and it gets more common every day. Some visible examples of SaaS applications include the customer relationship management provided by Salesforce.com CRM, the document creation services provided by Google Apps, and the business-oriented email service offered by Microsoft Exchange Online. Going forward, more and more organizations are likely to rely on cloud applications for a range of business functions.

As Figure 1 shows, however, applications aren't the only cloud computing option. It's also possible to create platforms that run in Internet-accessible data centers rather than inside your own firewall. These cloud platforms can run applications, store data, and provide other useful services. And while one important role of cloud platforms is to support cloud applications, cloud platforms themselves are of interest to developers rather than end users, as the figure shows. Several providers offer public cloud platforms today, but the technology was pioneered by Amazon. Their offering, called Amazon Web Services (AWS), remains a leader in this space, and it's worth a closer look.

## AN EXAMPLE CLOUD PLATFORM: AMAZON WEB SERVICES

Like an application running in your data center, an application running on a cloud platform needs compute services and storage services. Figure 2 shows how AWS provides these things.



**Figure 2: Amazon Web Services provides virtual machines, services for storing unstructured and relational data, and more.**

To run applications, AWS provides virtual machines (VMs) through its Elastic Compute Cloud (EC2) service. Developers can request a VM (also called an EC2 *instance*) through the AWS Web site, install any software they like in it, then begin using it. The VM is actually located at an Amazon data center, and its creator is charged for each hour the VM runs. This approach is commonly called *Infrastructure as a Service* (IaaS), since it provides a basic foundation for running software.

The creator of a VM specifies an Amazon Machine Image (AMI) from which this instance should be created. An AMI can be based on either Linux or Windows Server, and many different AMIs are available. As Figure 2 suggests, for example, some AMIs contain preconfigured database management systems, while others contain Web servers and other supporting software for applications.

To store data, AWS provides several different options. The most commonly used choice today is the Simple Storage Service (S3), which stores unstructured data (known as binary large objects, or blobs). S3 also provides storage for Elastic Block Store (EBS) volumes, which allow persisting the contents of an EC2 VM's file system. Another AWS storage option is the Relational Database Service (RDS), giving developers managed relational storage based on MySQL. And although they're not shown in the figure, AWS provides other choices as well, including support for storing very large structured, non-relational datasets.

The basics of cloud technology aren't hard to understand. Yet by itself, this technology has no value—the value comes from how it's used. One important example of this is providing GIS services in the cloud, as described next.

## GIS AND CLOUD COMPUTING: EXAMPLES FROM ESRI

Esri uses the cloud today in several different ways. The currently available options include the following:

- The ability to deploy ArcGIS Server on Amazon Web Services.
- ArcGIS.com, a Web site offering tools and shared data for GIS applications.
- ArcLogistics, a cloud application for optimizing routing, such as for delivery vehicles.
- Business Analyst Online, a cloud application for geographic analysis of demographic, consumer, business, and other data.

These examples illustrate how cloud computing can make life better for GIS developers and users. Each one is worth a closer look.

### ARCGIS SERVER ON AMAZON EC2

ArcGIS Server is a platform for delivering GIS services to software on other systems. These capabilities are exposed as RESTful services, via SOAP, and in other ways, and they can be consumed by clients written using various technologies, including JavaScript, Adobe Flex, and Microsoft Silverlight.

Today, customers typically deploy ArcGIS Server on a computer running in their own data center. It's also possible, however, to deploy ArcGIS Server in the cloud using AWS. As Figure 3 shows, ArcGIS Server itself can run in an EC2 VM running Windows, while a relational database holding GIS data runs in a second VM.

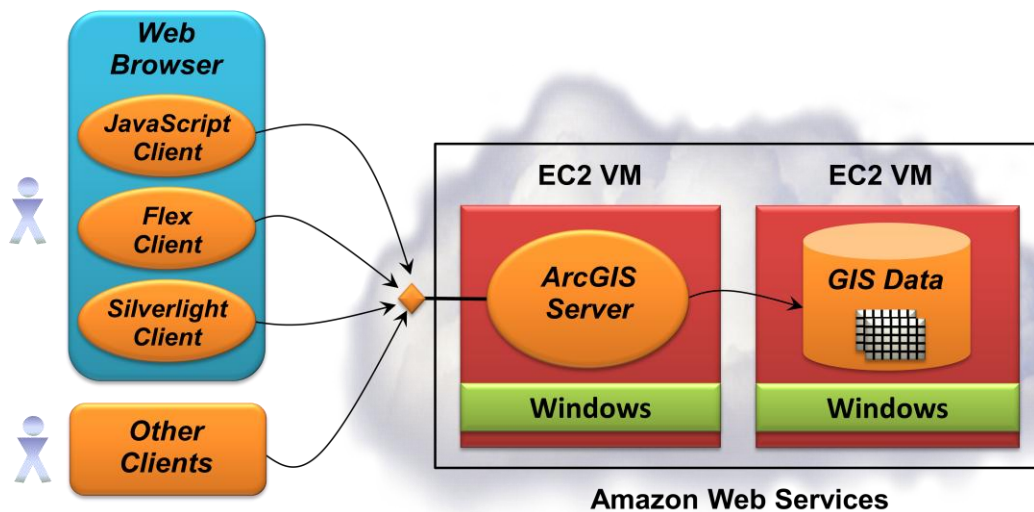


Figure 3: ArcGIS Server can be deployed in an EC2 VM on Amazon Web Services.

Deploying ArcGIS Server on AWS provides an alternative to deploying it solely in your own data center. But why would anybody do this? What are the advantages of choosing this option? There are several possible answers, including the following:

- Easier deployment: Esri provides a preconfigured AMI containing ArcGIS Server. Rather than install and configure the product on a machine in your data center, you can just create an EC2 VM from this AMI.
- Faster deployment: In many organizations, making a server machine available to run new software requires following a multi-step process, one that can take weeks or more. With AWS, by contrast, anybody with a valid credit card can get an EC2 VM deployed in a few minutes. If getting ArcGIS Server—and the application that uses it—up and running quickly is important, using the AWS option can make sense.
- Lower cost: As mentioned earlier, AWS charges customers for each hour a VM is running. Depending on usage and costs in your own data center, running ArcGIS Server on AWS might be a less expensive option. This is especially true for applications with elastic (i.e., widely varying) demand. If, say, an application typically uses a single instance of ArcGIS Server, but needs ten instances for occasional peak loads, the AWS pay-as-you-go model lets you pay for this higher usage only when you need it.
- Broad availability: Because software running on AWS can be accessed by anybody with an Internet connection, GIS services exposed by ArcGIS Server in an EC2 VM can be broadly available. While this raises the security bar, it can be an attractive approach for providing some kinds of information.
- Better performance: ArcGIS Server allows running various kinds of analysis on GIS data. If you're creating an application that performs complex data analysis, you might choose to deploy multiple instances of ArcGIS Server in multiple EC2 VMs, then use all of them in parallel to work on the same data. Applications running on AWS also have access to large amounts of bandwidth, which can significantly improve performance in some cases.
- Simpler development and testing: Because the AWS environment can be essentially identical to an on-premises Windows environment, a development team can build and test an ArcGIS Server application in the cloud, then deploy it either in the cloud or on premises. Given that creating (and paying for) EC2 VMs can be more flexible than creating the same environment in your own data center, this can make the development process easier.

Running ArcGIS Server in the cloud isn't always the right solution, of course. In many organizations, for instance, storing sensitive information outside the firewall is frowned upon. Still, this alternative deployment approach can sometimes be the best option.

## ARCGIS.COM

Using the cloud to share GIS data and applications is an attractive idea. ArcGIS.com, run by Esri, provides a good example of this. Figure 4 shows the site's home screen.

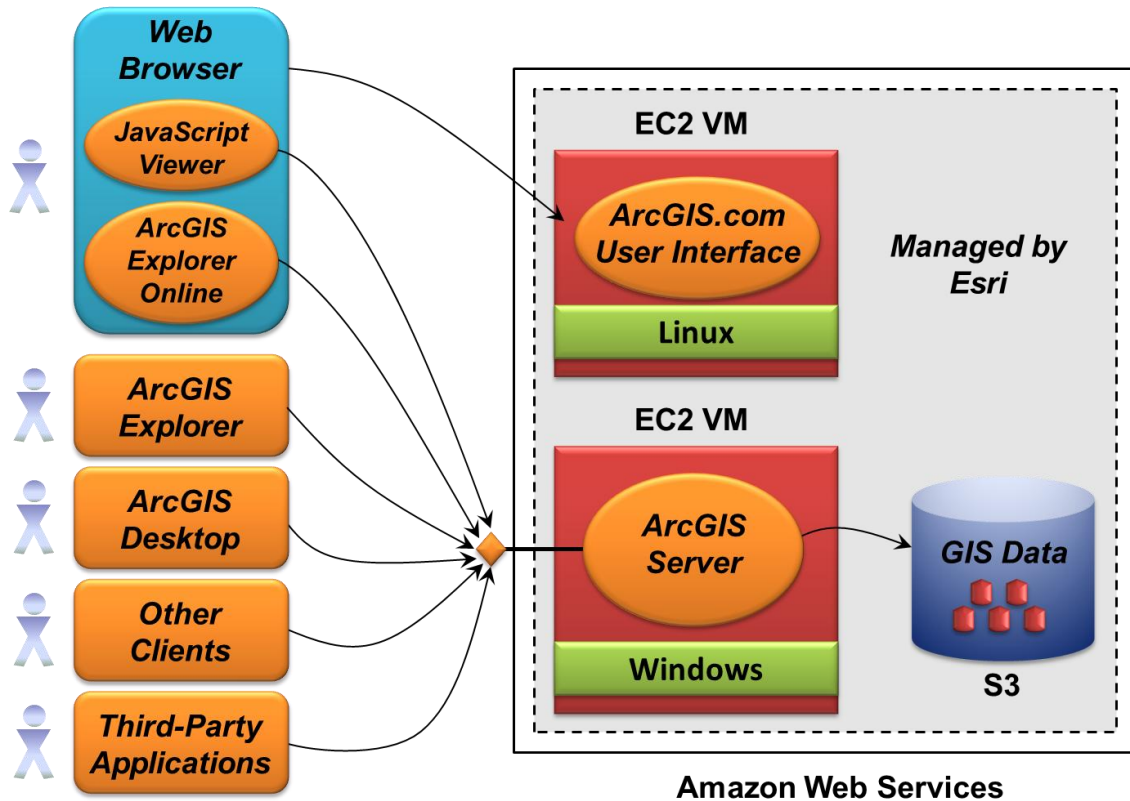


**Figure 4: ArcGIS.com is a Web site providing access to GIS data and applications.**

As this screen shot suggests, ArcGIS.com provides a number of services, including the following:

- Storage of publicly available maps and other GIS information. (Esri provides some of this data, including things such as layers and basemaps, but users of the site also contribute.)
- Esri-created applications for examining and working with the site's information, such as tools for creating maps.
- A central site for finding and accessing GIS applications that use the data on ArcGIS.com, including applications that run on computers in non-Esri data centers.

Figure 5 shows the main components of ArcGIS.com.



**Figure 5: Run by Esri, ArcGIS.com provides Internet-accessible GIS data and applications that can be used by many clients.**

ArcGIS.com provides tools designed to let non-expert users create and work with maps. As Figure 5 shows, those tools include a JavaScript viewer for browsing GIS data, along with the Silverlight-based ArcGIS Explorer Online, a more powerful tool for creating and working with maps. Both of these tools access data via services exposed by ArcGIS Server. These services can also be accessed through other Esri-provided clients, including ArcGIS Explorer (the more functional on-premises version of ArcGIS Explorer Online), ArcGIS Desktop, and mobile devices such as the iPhone. As mentioned earlier, it's also possible for non-Esri applications to access this data through the services exposed by ArcGIS Server.

As the figure shows, Esri has built ArcGIS.com today on Amazon Web Services. The site's logic, including its user interface and ArcGIS Server, runs in EC2 VMs, while the GIS data is stored as blobs in S3. This dependence isn't visible to the site's users, of course—they just see the user interface and other services that the site offers—but it's another example of how Esri is using cloud technology.

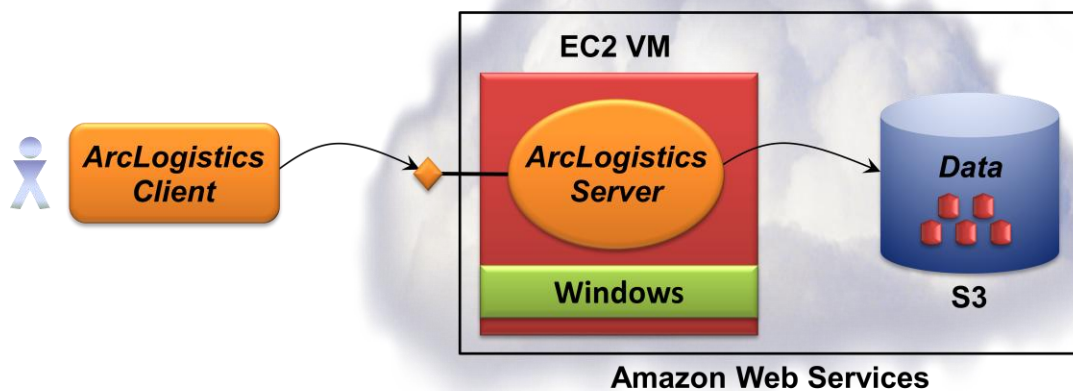
One final question worth addressing is this: Is ArcGIS.com a Web site or is it a cloud application? The real answer is that it's both. Some Web sites, such as those that provide only static pages to their users, clearly don't qualify as cloud applications. Others, though, especially sites offering useful tools such as ArcGIS.com, just as clearly do fit into the category of cloud applications. While not every Web site is a cloud application, it's probably fair to say that every cloud application can be seen as a Web site.



## ARCLOGISTICS

Think about an organization that maintains a fleet of vehicles. Maybe it's a large delivery company that sends out hundreds of trucks a day, for example, or a home healthcare service that must coordinate pick-ups and drop-offs of a dozen people. While the details differ, these firms face a similar problem: They need to route their vehicles as efficiently as possible. Add in other important variables—vehicle capacity, the times that customers can accept delivery, and more—and the problem gets even harder.

The goal of ArcLogistics is to help them solve this problem. This cloud application lets a user enter the number of vehicles together with the stops they must make, then get back optimized routing for these vehicles. Figure 6 illustrates the application's components.



**Figure 6: ArcLogistics is a cloud application for creating optimized routes.**

Once again, this application is built on Amazon Web Services, with logic running in an EC2 VM and data stored in S3. (No sensitive client data is stored in the cloud, however—it's sent in, used for computation, and deleted, which minimizes security issues.) Note that the ArcLogistics client isn't a Web browser. Instead, it's a custom Windows application built using Windows Presentation Foundation. And although it's not shown in the figure, the ArcLogistics server component is actually built using ArcGIS Server.

Esri also provides a desktop version of ArcLogistics. The application's cloud incarnation has some significant advantages over its on-premises predecessor, however. Those advantages include the following:

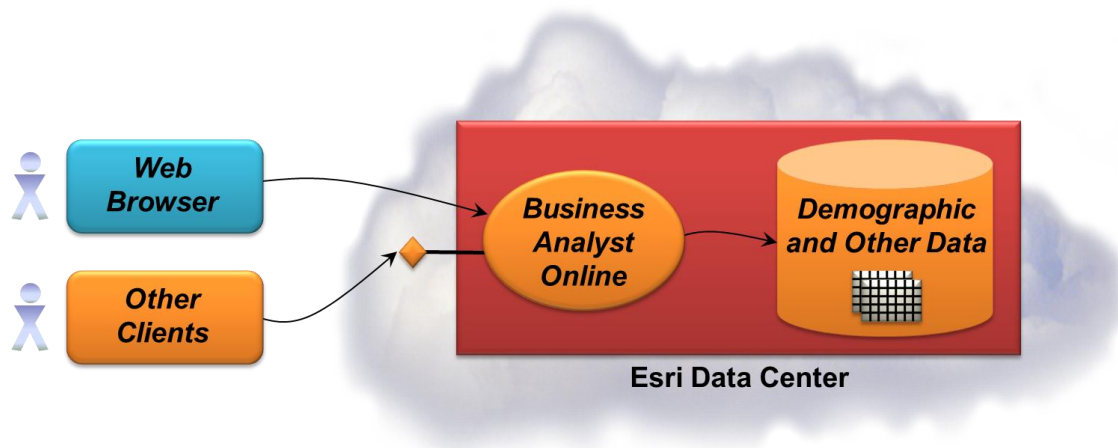
- Lower cost of entry: Rather than the fixed price of the desktop version, which is high enough to be feasible only for larger organizations, ArcLogistics in the cloud is priced based on the number of vehicles a customer has. This makes getting started cheaper, and it also opens the application to smaller organizations with only a few vehicles.
- Try before you buy: Rather than buy an on-premises application, install it, learn how to use it, and only then determine whether it actually has business value, ArcLogistics in the cloud lets a potential customer try the service on a small scale. The customer can then adopt (and pay for) the application only if it delivers significant value.

- No on-premises data or software updates: To do its job, ArcLogistics needs an up-to-date picture of the geography over which it's finding optimized routes. Since the world changes, this database must be updated periodically. With the desktop version, each customer must regularly install updated data sent out by Esri. With ArcLogistics in the cloud, this is no longer required. Instead, all customers rely on a single shared database that's updated centrally. Everybody always has current information with no need to install updates themselves. And because there's just one copy of the application, Esri can also update this code centrally, making new features immediately available to all users.
- Better customer support: Since all customers of ArcLogistics in the cloud are using the same code with the same underlying data, the application's support people can do a better job. With the desktop version, different customers might have different releases, older data, or both, making support significantly more challenging.
- Better performance: Computing optimized routes is CPU-intensive work. Doing this on a cloud platform such as AWS lets ArcLogistics create as many VMs as needed for whatever the current customer demand might be. It also lets customers avoid dedicating their own compute resources to this task.

ArcLogistics is a good example of an application that probably makes more sense in the cloud than it does in your data center. It's compute intensive, relies on a large and changing database, and doesn't need to store sensitive information outside your firewall. It's also a good example of how cloud computing can improve our lives.

## BUSINESS ANALYST ONLINE

Suppose it's your job to choose a new location for a children's clothing store. To make a good decision, you'd like to know how many families with young children live within, say, a 20-minute drive of each potential location, along with their median income. Answering questions like these is the purpose of Business Analyst Online. By providing demographic information and other data, along with tools for working with that data, this cloud application can help business owners, planners, realtors, and others make better decisions. Figure 7 shows its main components.



**Figure 7: Business Analyst Online is a cloud application for working with demographic data, consumer data, and other information.**

As the figure shows, Business Analyst Online can be accessed from an ordinary Web browser. The application also exposes a Web services interface that lets it be used by other clients. Esri provides an iPhone app, for example, and third parties can also access these services. Notice that this cloud application isn't built on Amazon Web Services or another cloud platform. Instead, it runs in Esri's own Internet-accessible data centers. While cloud platforms can be quite useful, there's no requirement that a cloud application be built on one.

Like ArcLogistics, this application also comes in on-premises versions, including both a desktop and a server product. In some cases, such as when an organization wishes to merge its own data with the standard data this application provides, one of these on-premises versions is a better choice. In fact, the on-premises versions of the product allow more sophisticated analysis than their cloud siblings (today, at least). Yet there are a number of reasons why Business Analyst Online makes sense as a cloud application. They include the following:

- More approachable for infrequent or non-expert users: Many potential customers need the services of Business Analyst Online only occasionally. Rather than make them install on-premises software that they don't often use, putting this application in the cloud makes it simpler and quicker to access when needed. And because Business Analyst Online is designed to be easier to use than the on-premises version—its users aren't expected to be GIS professionals—using it only occasionally isn't especially difficult.
- Lower cost of entry: Like most cloud applications, Business Analyst Online charges its customers based on usage. While a frequent user of the service might buy an annual subscription, a customer is also free to buy just a single report and never use the service again. In either case, you pay only for what you use.
- No on-premises data or software updates: All Business Analyst Online customers use the same database, which is updated regularly by Esri. And because it's a cloud application, the code is also updated centrally. This makes software updates simpler—they're done by Esri—with new features available immediately to all of the application's users.
- Better customer support: Because all Business Analyst Online customers use the same code and the same data, Esri's support people don't need to worry about differences in configuration or installed data. This can potentially help them resolve customer support questions faster than with an on-premises product.

In some ways, GIS applications are an especially good fit for the cloud. Because they rely on large and changing data sets, putting both this data and tools to work with the data in the cloud can make sense. Business Analyst Online provides a good example of the value this provides.

## CONCLUSION

There's no denying it: Cloud computing is here, and its effects will be widespread. In the GIS world, Esri provides clear examples of how these new technologies can be used. Those examples include:

- Using a cloud platform to provide new deployment options, such as running ArcGIS Server on Amazon Web Services.
- Using a cloud platform to support Web sites that provide broadly usable GIS data and tools, as illustrated by ArcGIS.com.
- Offering cloud applications that provide useful alternatives to their on-premises siblings, as shown by ArcLogistics and Business Analyst Online.

The move to the cloud is all but certain to have an impact on your organization, both in GIS and other areas. Why wait? The time to start understanding this shift is now.

## **ABOUT THE AUTHOR**

David Chappell is Principal of Chappell & Associates ([www.davidchappell.com](http://www.davidchappell.com)) in San Francisco, California. Through his speaking, writing, and consulting, he helps people around the world understand, use, and make better decisions about new technologies.