

CONTENTS

- ▶ Types of Cloud Computing
- ▶ Cloud Basics
- ▶ Hybrid Clouds
- ▶ Cloud Services
- ▶ Key Benefits...and more!

Understanding Cloud Computing

AMAZON WEB SERVICES, GOOGLE CLOUD, AND MICROSOFT AZURE

UPDATED BY DAVID POSIN
 ORIGINAL BY DANIEL RUBIO

ABOUT CLOUD COMPUTING

Cloud computing as an idea encapsulates a wide range of business and consumer services. At its most basic, cloud computing refers to using resources over the Internet that are housed somewhere else in the world for specific purposes.

Interacting with a cloud service is usually done through a web browser or command line interface. There is usually no software to install, no hardware to set up, and no data center to rent space from.

Google and Amazon were two of the earliest to the field of cloud computing, soon followed by Microsoft. The portfolio of products offered by cloud companies has matured and evolved over time, and so has the competition. Firms such as IBM and Rackspace, among others, have joined the fray. The field of cloud computing now includes a myriad of companies all offering services. This card will focus on Google, Amazon, and Microsoft since they are the market leaders at the time of this writing, and have the most mature product sets.

TYPES OF CLOUD COMPUTING

Cloud computing can be broken down into three main categories: IaaS, PaaS, and SaaS. Each is discussed below in more detail.

IAAS (INFRASTRUCTURE AS A SERVICE)

Infrastructure as a Service is what people usually think of when discussing cloud computing. IaaS emulates the hardware that companies traditionally purchased to create internal and external networks. The chore of installing hardware and connecting it to physical networks has been replaced with the task of creating robust configurations. Almost every part of a traditional network is offered as a cloud service.

EXAMPLES

- Servers
- Firewalls
- Routers
- Load Balancers
- CDN (Content Delivery Network)

BENEFITS

- **No Hardware:** No need to invest in equipment
- **Cost is a Function of Usage:** Initial investments are much cheaper than equipment acquisition
- **Scalability:** A network can easily scale, retract, or transform without taking a financial impact from leasing clauses or acquisition
- **Swift Management:** Ability to react to network problems or improve network configurations quickly

CONS

- **Learning Curves are Steep:** It takes time and effort to be able to use an IaaS service to its fullest potential
- **Outside your control:** The resources supporting the network are controlled by another company
- **Service Agreement:** The agreement controlling your use of the service might require legal review
- **Subject to Network Reliability:** Most IaaS vendors can boast of uptime in the 99% area so there is very little to worry about, but outages do occur

SAAS (SOFTWARE AS A SERVICE)

Software as a Service is so common in the modern Internet age that it is almost overlooked as part of the cloud computing field. Software as a Service describes any Internet-based product that provides a specific set of operations. This includes everything from email to contact management.

Interacting with a SaaS product had traditionally been done through a web browser. This has been slowly changing with the ubiquity of mobile devices. Most SaaS products now include web and mobile device clients.

EXAMPLES

- **Email services:** Gmail, Outlook.com
- **Customer Relationship Management (CRM):** Salesforce
- **Business Productivity:** Office.com, Dropbox, Box.com, Google Docs
- **Application Support:** InsideOps, Logentries, Google Forms



logentries
 BY RAPID7

Log These 15 Events,
 Spread DevOps Love

Download Whitepaper



The Fastest Way to Analyze Your Log Data

No complex setup. No waiting. Just answers.

Logentries' centralized logging service takes only minutes to setup and is free to use. Centralize logs from multiple sources and search using a simple query builder.



Collect & Centralize

Stream logs from all environments into Logentries' centralized log server in real-time.



Search & Analyze

Search using basic keywords or our advanced query language for deeper analysis.



Visualize & Report

Build charts, graphs and alerts to monitor performance, user behavior and more.

[Start a Free Trial](#)

BENEFITS

- **Managed Service:** Customers of a SaaS product are able to focus on using it; there is no need to worry about the software or hardware required to make it run
- **Quick and Regular Updates:** The product is the main focus of the company; the vendor is constantly improving, fixing, and patching to improve their tool's usefulness
- **Portability:** Enterprise-grade SaaS products will usually work on all modern browsers and a wide range of mobile devices
- **Dedicated Customer Support:** If something needs to be fixed, your internal resources don't have to slow their work down to fix it

CONS

- **Product Growth is Outside Your Control:** The roadmap of the tool's development and the features being added are outside your control. However, most vendors have a method for customers to request features
- **Service Agreement:** The agreement controlling your use of the service might require legal review
- **Customization Can Be Difficult:** Customizing the product to specifically meet your company's needs might not be possible or might be costly
- **Reliability:** Most SaaS products have uptime in the high 90% range, but outages do happen

PAAS (PLATFORM AS A SERVICE)

Platform as a Service products has received less attention than their SaaS and IaaS siblings until recently. The power that a platform can provide is becoming much more appreciated in recent days, which has led to a boom in PaaS products. PaaS products sit in the middle between IaaS and SaaS. PaaS services allow customers to create a customized experience on top of an existing set of tools.

EXAMPLES

- **Customized Interactive Running Tournaments:** RaceLink
- **Application Building:** Google App Engine, SAP Hana, Cloud Foundry

BENEFITS (SIMILAR TO SAAS PRODUCTS)

- **Quick and Regular Updates:** The product is the main focus of the company; the vendor is constantly improving, fixing, and patching to improve their tool's usefulness
- **Portability:** Applications built on enterprise-grade PaaS products will usually work on all modern browsers and a wide range of mobile devices
- **Customization:** The ability to use the PaaS product as a starting point means the applications built on it can be more closely tied to your business and its processes
- **RAD (Rapid Application Development):** Getting applications built, and in the hands of your employees or customers happens much more quickly than traditional application development

CONS (SIMILAR TO SAAS PRODUCTS)

- **Product Growth is Outside Your Control:** The roadmap of the tool's development and the features being added are outside your control. However, most vendors have a method for customers to request features

- **Service Agreement:** The agreement controlling your use of the service might require legal review
- **Reliability:** Most PaaS products have uptime in the high 90% range, but outages do happen
- **In-House Development Still a Factor:** Internal employees might need to spend time tweaking, fixing, or enhancing applications

CLOUD BASICS

WHAT IS THE CLOUD?

The Cloud is shorthand to describe the huge interconnected network of servers designed to deliver computing resources without a sense of location. In other words, The Cloud feels to users like a giant formless mass of computing power that runs everything from email to mapping to mobile applications.

In business terms, there is no "The Cloud." Cloud computing is really a collection of licensed services provided by different vendors. Cloud services take technology management and technology acquisition away, and replace it with products that are managed elsewhere and are only active when needed.

In order to better understand what cloud services provide, let's compare cloud to other traditional technology location options:

| | |
|--------------------|--|
| On-premise | Describes software and hardware that is housed and directly managed inside a business' physical location. On-premise hardware and software is usually always on and available. |
| Off-premise | Physical hardware and software tied to a particular location, often leased space. Off-premise hardware and software is usually on and devoted to the customer using it. This can refer to offsite data centers, tape archive locations, etc. |
| Cloud | Physical and virtualized hardware and software that is hosted as needed on any available resource inside a farm of resources. A company may be using a server to run an application but it is not a physical server leased by the company. The particular server being used changes as needed by the provider and can be in any number of locations. |

WHERE ARE CLOUDS LOCATED?

Consumers and businesses use different cloud services with such ease and for so many purposes that clouds can feel very ephemeral. When you begin using cloud computing with vendors directly, you are very quickly reminded that geography does matter. Although it may be true that most applications can be reached anywhere, unless geography is taken into account, it will not be the same experience everywhere.

Cloud computing products are generally region aware. All providers will give you some control over where to locate your environments. The location and scope of a region will differ per provider.

The number and location of regions you need will largely depend on the scope of your application, network, latency tolerance, and SLAs. The particular services required will also impact the choice of regions. Not all services are available in all regions.

The biggest factor that often creeps into conversations about regions

is latency. The amount of time it takes data and resources to travel from the server to the client can make the difference between an application being adopted or abandoned. An application serving static web content can probably be in fewer zones. The lighter data load means transit will be faster. However, an application serving large datasets should probably be located in the region closest to the user base. Larger data sizes traveling farther distances will accrue more latency.

Latency is an important factor for networks as well as applications. The number of hops that network packets need to take can quickly impact system response time. A load balancer can help with controlling traffic but if traffic has to go through several networks hops before hitting the load balancer, then the balancer's impact is minimal. Consider network latency as well as application latency when choosing regions to host your network.

Amazon Regions (as of 3/2017)

| | |
|----------------------|---|
| Asia Pacific | <ul style="list-style-type: none"> • Mumbai • Seoul • Singapore • Sydney • Tokyo |
| Canada | <ul style="list-style-type: none"> • Central |
| China | <ul style="list-style-type: none"> • Beijing |
| Europe | <ul style="list-style-type: none"> • Frankfurt • Ireland • London |
| South America | <ul style="list-style-type: none"> • São Paulo |
| US East | <ul style="list-style-type: none"> • Northern Virginia • Ohio |
| US West | <ul style="list-style-type: none"> • Northern California • Oregon |
| GovCloud | <ul style="list-style-type: none"> • GovCloud |

Google Regions (as of 3/2017)

| | |
|----------------------------------|---|
| Central US | <ul style="list-style-type: none"> • Council Bluffs • Iowa |
| Eastern US | <ul style="list-style-type: none"> • Berkeley County • South Carolina |
| Eastern Asia-Pacific | <ul style="list-style-type: none"> • Changhua County • Taiwan |
| Northeastern Asia-Pacific | <ul style="list-style-type: none"> • Tokyo • Japan |
| Western Europe | <ul style="list-style-type: none"> • St. Ghislain • Belgium |
| Western US | <ul style="list-style-type: none"> • The Dalles • Oregon |

Microsoft Regions (as of 3/2017)

| | |
|---------------------|---|
| Americas | <ul style="list-style-type: none"> • Virginia • Iowa • Illinois • Texas • West Central US • California • West US 2 • Quebec City • Toronto • Sao Paulo State |
| Asia-Pacific | <ul style="list-style-type: none"> • Beijing • Busan • Chennai • Hong Kong • Mumbai • New South Wales • Osaka • Pune • Saitama • Shanghai • Singapore • Tokyo • Victoria |
| Government | <ul style="list-style-type: none"> • Iowa • Virginia • US DoD East • US DoD Central |
| Europe | <ul style="list-style-type: none"> • Cardiff • Frankfurt • Ireland • London • Magdeburg • Netherlands |

HYBRID CLOUDS

Cloud computing has evolved a lot over the last few years. What was once an external addition to business' infrastructure has become an integral component. The concept of the hybrid cloud is a reflection of that evolution.

A hybrid cloud is one that integrates cloud resources with other on-premise, off-premise, and third-party resources. In most cases, these integrations are transparent to the user. For example, a local file server might host tomorrow's presentation, but employees working on it are using Office.com's cloud-based Powerpoint product.

Some cloud services even allow physical integrations to meet security, latency, and other business needs. Vendors often refer to these connections as connecting to a cloud's edge. Amazon's AWS Direct Connect, Google's Interconnect, and Microsoft's Express Gateway are all examples of this configuration.

CLOUD SERVICES

Cloud company product offerings tend to be wide, specific, and designed to fill niches. Collectively, it is possible to run an entire organization's technology stack on the cloud. The more common case is to use cloud providers to enhance existing pieces of infrastructure, or add pieces that

are missing. Google, Amazon, and Microsoft all provide a combination of IaaS, SaaS, and PaaS products for companies to pick and choose from.

The tables below attempt to identify the most widely applicable offerings in each of the big three that perform similar operations. Some of the categories may differ slightly as not all providers use all the same nomenclature. The table uses the category that matches the majority when there is a discrepancy.

Computing - Virtual servers, information processing, and event handling

| AMAZON | GOOGLE | MICROSOFT |
|---|--|---|
| Virtual Machines - including scaling | | |
| <ul style="list-style-type: none"> Amazon EC2 Amazon Lightsail | <ul style="list-style-type: none"> Compute Engine | <ul style="list-style-type: none"> Virtual Machines Virtual Machine Skill Sets |
| Virtual Machine Containers - i.e. Docker containers | | |
| <ul style="list-style-type: none"> Amazon EC2 Container Registry Amazon EC2 Container Service | <ul style="list-style-type: none"> Container Engine Container Registry | <ul style="list-style-type: none"> Azure Container Service Azure Container Registry |
| Web Application Support - ranging from supporting a ready-made application support platform to full PaaS application creation and management | | |
| <ul style="list-style-type: none"> AWS Elastic Beanstalk | <ul style="list-style-type: none"> App Engine | <ul style="list-style-type: none"> App Service |
| Event Handling - Serverless functions that can respond to specific event signals | | |
| <ul style="list-style-type: none"> AWS Lambda | <ul style="list-style-type: none"> Cloud Functions | <ul style="list-style-type: none"> Functions |
| Batch Processing - Batch computing and large scale processing that can be done off peak hours | | |
| <ul style="list-style-type: none"> AWS Batch | | <ul style="list-style-type: none"> Batch |

Storage - Data storage structures akin to disks and tapes

| AMAZON | GOOGLE | MICROSOFT |
|---|--|---|
| Object Storage - storage of object based data | | |
| <ul style="list-style-type: none"> Amazon S3 | <ul style="list-style-type: none"> Cloud Storage | <ul style="list-style-type: none"> Storage |
| Block Storage - storage of data in blocks and tracks. Ideal for virtual machine disk drives | | |
| <ul style="list-style-type: none"> Amazon EBS | <ul style="list-style-type: none"> Persistent Storage | <ul style="list-style-type: none"> Storage |
| File Storage - file system like storage | | |
| <ul style="list-style-type: none"> Amazon Elastic File System | | <ul style="list-style-type: none"> Storage |
| Archives - Services to store large amounts of data that don't require immediate and regular access | | |
| <ul style="list-style-type: none"> Amazon Glacier | | |

| | | |
|--|--|---|
| Hybrid - Storage mechanisms that bridge in-house storage with cloud storage | | |
| <ul style="list-style-type: none"> AWS Storage Gateway | | <ul style="list-style-type: none"> StorSimple |
| Automated On-Premise Backup - Automated backup of on-premise computers and server applications to the cloud | | |
| | | <ul style="list-style-type: none"> Backup |
| Private Cloud Recovery - Automated private cloud recovery services | | |
| | | <ul style="list-style-type: none"> Site Recovery |
| Data Migration Services - Move large amounts of data around the network, or collect multiple data sources into one place. | | |
| <ul style="list-style-type: none"> AWS Snowball AWS Snowball Edge AWS Snowmobile | | |

Databases - Wide range of organized data repositories from relational SQL to NoSQL to cloud caching

| AMAZON | GOOGLE | MICROSOFT |
|--|--|--|
| Relational Databases - includes relational database environments that support several database applications to fully managed, high availability, enterprise grade systems | | |
| <ul style="list-style-type: none"> Amazon Aurora Amazon RDS | <ul style="list-style-type: none"> Cloud SQL Cloud Spanner | <ul style="list-style-type: none"> SQL Database SQL Server Stretch Database (Combination of on premise and cloud SQL server) |
| NoSQL Databases - offerings range from simple NoSQL databases to fully managed, enterprise grade NoSQL instances | | |
| <ul style="list-style-type: none"> Amazon EBS | <ul style="list-style-type: none"> Persistent Storage | <ul style="list-style-type: none"> Storage |
| Caching Services - Redis-based or proprietary cache servers | | |
| <ul style="list-style-type: none"> Amazon ElastiCache | | <ul style="list-style-type: none"> Redis Cache |

Networking - Everything a modern small, medium, or large business needs to build and deploy robust networks

| AMAZON | GOOGLE | MICROSOFT |
|--|---|--|
| Virtual Networking - all three providers offer a wide range of network configuration options. Networks can be singular or multiple, can connect to each other, or be isolated | | |
| <ul style="list-style-type: none"> Amazon VPC | <ul style="list-style-type: none"> Cloud Virtual Network | <ul style="list-style-type: none"> Virtual Network VPN Gateway |
| Interconnected Services - Provides a direct connection from an on-premises site directly to the provider's cloud, usually through a third-party provider or dedicated hardware. | | |
| <ul style="list-style-type: none"> AWS Direct Connect | <ul style="list-style-type: none"> Interconnect | <ul style="list-style-type: none"> Express Gateway |

| CDN - Cloud based content delivery networks | | |
|---|------------------------|----------------------------|
| • Amazon Cloudfront | • Cloud CDN | • Content Delivery Network |
| Load Balancing - Automatically balance traffic amongst multiple cloud resources endpoints | | |
| • Elastic Load Balancing | • Cloud Load Balancing | • Load Balancer |
| DNS - Cloud based Domain Name Services | | |
| • Amazon Route 53 | • Cloud DNS | • Azure DNS |

Big Data and Big Data Analytics - Data warehouses and tools to analyze the data

| AMAZON | GOOGLE | MICROSOFT |
|--|--------------------------------------|---|
| Data Warehouses - Managed warehouse service | | |
| • Amazon Redshift | • BigQuery | |
| Hadoop Framework Services - Hadoop, and related open source framework, implementations | | |
| • Amazon EMR | • Cloud Dataproc | • HDInsight • Apache Spark for HDInsight |
| Datastore Management - Tools to clean, manipulate, transform, stream, search, and move large amounts of data through warehouses and for analysis | | |
| • Amazon Kinesis • Amazon Glue • AWS Glue • AWS Data Pipeline | • Cloud Dataflow • Cloud Dataprep | • Data Lake Store • Data Factory |
| Analysis Services - Analysis tools to provide searching or reporting on large datasets | | |
| • Amazon Athena • Amazon Elasticsearch Service • Amazon Quicksight | • Cloud Datalab | • Data Lake Analytics • Stream Analytics |

And That's Not All ...

| AMAZON | GOOGLE | MICROSOFT |
|--|--|---|
| Other Services - There are a wealth of offerings that are designed to complement the above services, are very niche in nature, or are implemented too uniquely to each service for generalization here. | | |
| <ul style="list-style-type: none"> • Developer Tools • Management Tools • Security, Identity, and Compliance • Artificial Intelligence • Mobile Services • Application Services • Messaging • Business Productivity • Desktop & App Streaming • Internet of Things • Game Development | <ul style="list-style-type: none"> • BigQuery • Machine Learning • Identity & Security • Management Tools • Developer Tools • Genomics | <ul style="list-style-type: none"> • Web + Mobile • Internet of Things • Enterprise Integration • Security + Identity • Developer Tools • Monitoring + Management |

KEY BENEFITS OF CLOUD SERVICES

PAY ONLY WHAT YOU CONSUME

Until a few years ago, the costs of providing off-site services were similar to a leasing plan; a certain amount of allocated resources and an incurred cost whether such resources were consumed or not. Cloud computing as it's known today has changed this. The various offerings provided by all major cloud computing platforms function on a consumption-based plan. The use of services is charged based on the applicable resources of that service. For example, a virtual machine will incur cost based on the computing power used and the server uptime. Meanwhile, a data warehouse service might charge based on the amount of data stored and the frequency with which reports are run against it.

This pricing structure may seem convoluted, but it greatly benefits the customer. Cost incurred for use means companies of all sizes can use resources that would be cost prohibitive if they had to acquire all the hardware, software, physical space, and connectivity on their own. An on-premise data warehouse can be extremely costly. A cloud based, consumption priced data warehouse is much more affordable.

ONE TIME & ALTERNATIVE PRICING

Cloud services are often subject to traffic spikes due to one time events (e.g. National broadcast exposure, Super Bowl commercial, etc.). Not only can this type of provisioning be expensive, but oftentimes difficult to achieve.

Cloud computing platforms allow applications to access additional resources immediately and as needed without additional configuration.

In addition to peak and one time pricing, some services allow for other alternative pricing structures. One example is pricing that becomes cheaper after peak hours. If operations can occur when there is a predictable lull in resource use, such as midnight to early morning for that region, then those operations can sometimes be done at a cheaper rate.

Alternatively, if operations have very flexible usage requirements, then some providers allow customers to bid on access to resources only when

they become available for however long they are available at a price calculated by current resource use. For example, a large batch process can be scheduled to occur when computing resources at that region become available for \$0.x/minute due to transitory low usage windows until resource usage returns to normal or the batch process completes.

Finally, there is sometimes the option of dedicated resources. These are often the most costly of cloud resources. For some resources, such as disk storage or servers, it is possible to engage a lease type pricing model where a company pays for a resource to be dedicated only to them. This is useful for:

- Mission critical applications, networks, and processes that cannot risk being affected by regular virtualization or co-hosted processes
- Regulatory situations where the benefits of cloud computing are necessary but co-hosting would break regulatory requirements

AUTOMATED GROWTH & SCALABLE TECHNOLOGIES

Cloud resources are virtualized and configured for specific requirements but not a physical piece of hardware. This means that it is very easy to launch more of any specific resource needed by the company. If a customer is in a period of rapid growth, then the cloud can easily increase the amount of resources available without intense installation or configuration. This is also true in reverse. If a company needs to reduce cost, then turning off resources or using less of a resource does not require much effort.

Scaling is also extremely easy to do since cloud services can be as permanent or temporary as needed. An application designed for high usage on a cloud provider can easily scale from a small amount of used resource to a large amount without little extra configuration. The cloud service will often instantiate additional resources as needed, for only as long as needed.

OTHER CLOUD COMPUTING PROVIDERS

Amazon's, Google's, and Microsoft's cloud computing platforms may dominate the cloud computing space, but other providers have also emerged. Some of these providers include:

- [OpenStack Public Cloud](#)
- [IBM Cloud & Bluemix](#)
- [Heroku](#)

Many of these providers rely on industry standard virtualization and operating system technology making them competitive to the other cloud computing platforms. However, comparing these other providers to Google's, Microsoft's, or Amazon's cloud computing platforms can be more difficult due to the maturity and robustness of the big three's platforms.

Nevertheless, some of these other cloud computing providers have carved out niche markets in the cloud computing market. Some do so by adopting more aggressive pricing structures, catering to the specific needs of certain communities (e.g. Ruby/Rails, or Linux), or providing better customer service than their larger rivals.

CLOUDS DO GET DISRUPTED (RARELY)

Customers should hold their cloud providers to a 99+% uptime. Cloud providers are well aware of customer expectations on this point, and do everything to meet that threshold. However, if the mind-boggling number of hours that services are available is considered, then even a 1% downtime is likely to affect you. No technology is perfect.

All Cloud providers offer status pages so customers can be aware of any disruptions that occur:

Google Cloud Status Page: status.cloud.google.com

AWS Status Page: status.aws.amazon.com

Microsoft Azure Status Page: azure.microsoft.com/status

IBM Cloud Status Page: ibm.com/cloud-computing/social/maintenance

Rackspace Status Page: status.rackspace.com

CLOUD COMPUTING TEAM BLOGS

In order to keep abreast on the latest offerings made by cloud computing providers, we recommend you consult each platform's team blog.

Google Cloud Platform Blog: cloudplatform.googleblog.com

AWS blog: aws.amazon.com/blogs/aws

Microsoft Azure team blog: azure.microsoft.com/blog

IBM Cloud blog: ibm.com/blogs/cloud-computing

Rackspace blog: blog.rackspace.com

ABOUT THE AUTHOR



DAVID POSIN has been involved in the Information Technology Industry for 2 decades. Fifteen years of that time was spent consulting with many companies in a wide range of industries to build solid technology stacks and robust application architectures. David has watched the Cloud and the World Wide Web grow from their infancy, and now spends every day fully entrenched in those worlds. Currently, David builds high-performance web applications and offers professional technical writing services.



DZone communities deliver over 6 million pages each month to more than 3.3 million software developers, architects and decision makers. DZone offers something for everyone, including news, tutorials, cheat sheets, research guides, feature articles, source code and more.

"DZone is a developer's dream," says PC Magazine.

Copyright © 2017 DZone, Inc. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by means electronic, mechanical, photocopying, or otherwise, without prior written permission of the publisher.

BROUGHT TO YOU IN PARTNERSHIP WITH



DZONE, INC.
150 PRESTON EXECUTIVE DR.
CARY, NC 27513

REFCARDZ FEEDBACK WELCOME
refcardz@dzone.com

888.678.0399
919.678.0300

SPONSORSHIP OPPORTUNITIES
sales@dzone.com