



THE CLOUD AND ITS IMPLICATIONS FOR ENTERPRISE MOBILE DEVELOPERS



MOTODEV
FOR ENTERPRISE

developer.motorola.com/enterprise

TABLE OF CONTENTS

INTRODUCTION	1
SIGNIFICANT DEVELOPMENTS IN SERVER AND STORAGE TECHNOLOGY	2
THE ROLE OF VIRTUAL MACHINES	3
WHAT IS THE "CLOUD"?	4
LEVELS OF SERVICE FROM THE CLOUD	5
THE CLOUD'S SIGNIFICANCE FOR MOBILE	6
ARCHITECTING FOR THE CLOUD - SERVICE ORIENTED ARCHITECTURE	6
WEB SERVICES	7
COLLABORATION IN THE CLOUD	7
THE FUTURE OF CLOUD COMPUTING	8
BOTTOM LINE FOR MOBILE DEVELOPERS	9
MORE ON CLOUD COMPUTING	10

INTRODUCTION

Consider two, perhaps surprising, statistics. In 2009, technology analysts were still writing articles with titles like “Is the Cloud Hype?” But today, more than 900 million people are comfortably using the cloud (whether they realize it or not) through their membership in Facebook, Netflix, Gmail™, Amazon, and similar services.

The second statistic compares the relative popularity of computing devices. In the last quarter of 2011, the number of mobile devices sold outnumbered the number of PCs sold. Mobile phones and tablets, already widespread, have displaced PCs as the most popular computing platform.

This white paper examines how these two trends of “increasing cloud use” and “increasing mobile use” strengthen and reinforce each other. It explains how developers can access and use the cloud inside mobile apps.

Cloud computing and cloud storage today represent the latest results in a series of incremental technology developments. A straightforward way to understand the cloud is to review past milestones in structured storage. Going back to the 1981 launch of the IBM PC, essentially all disk storage was local and attached to the CPU. In the mid 1980s, the arrival of local-area-networked workstations fostered a new paradigm of client/server storage. Storage need no longer be directly attached to the user’s PC. It could now be centralized on a large server, and accessed remotely via requests from many different clients over the local subnet (at first) or over the Internet (eventually).



SIGNIFICANT DEVELOPMENTS IN SERVER AND STORAGE TECHNOLOGY

Early developments were about storage but with the rise of online web systems, processing power has also become a specialized resource. Computers that serve web pages are dedicated to that task alone. The trend of setting aside hardware for specific tasks grew more pronounced, with different servers assigned to load-balancing, to application execution, and to database hosting. Multi-tier app systems had arrived.

Table 1 below shows storage technology milestones, the incremental “enabling” changes, and some of the benefits delivered.

TECHNOLOGY MILESTONE	BENEFITS
<p>Client/server file sharing: Delivering file storage across a network</p>	<p>Centralization, cheaper administration, larger capacity, bulk storage</p>
<p>Data center with “islands” of SCSI disk arrays: Centrally-housed servers each with a “disk farm” often dedicated to individual apps</p>	<p>Scalability, cost reduction, data integrity (RAID)</p>
<p>Storage Area Network (SAN) and Network Attached Storage (NAS): SCSI disk “islands” absorbed into the fabric of a high speed switched ethernet. Storage is exported as raw blocks (SAN) or a filesystem (NAS)</p>	<p>Greatly simplifies storage administration and provisioning, parallel hardware supports very high data transfer rates</p>
<p>Multi-tier web system: Dedicating different hosts to different tasks - IP load balancer, web server, app server, database server</p>	<p>Higher performance and scalability</p>
<p>Virtual machine images: One or more completely isolated guest operating systems, running as a software image on a host computer</p>	<p>Isolation, security, ability to replicate and achieve scalability, ability to move to adjacent servers, operating system standardization, task dedication</p>

By running apps on dedicated servers (instead of individual desktops), updates and upgrades are greatly simplified. You only have to administer a few dedicated servers in the data center, not software on hundreds or thousands of client PCs on desktops throughout the enterprise. The rise of dedicated app servers was the point at which CPU power, too, became a centralized server resource. The next section examines the significance of virtual machines, and how they are used to manage CPU resources in the cloud.

THE ROLE OF VIRTUAL MACHINES

The next development was to virtualize the hardware. By running your web server in a virtual machine, you separate it from physical hardware, and allow better granularity for matching tasks to computing resources. Here's how that works.

A virtual machine (VM) is a completely isolated guest operating system, running on a host computer. The Android™ emulator used to test apps is a VM. The Android VM emulates an ARM cpu that runs Linux and the Android platform, on top of a desktop computer running MacOS, Windows or Linux. When using VMs, there are at least two operating systems - one OS running on the host PC, and another OS running on each VM. Each VM can run the same OS as the host, or a different OS.

A VM can also allow a single machine to act as if it were many independent machines (by running many VMs). Power users create a separate VM for each Windows app - one VM image for Internet Explorer, another for Word, a third VM for Microsoft Silverlight. By running in separate VMs, an app crash will only halt that specific VM, and not all instances of Windows apps. It costs a bit more in memory and licenses, but it saves hours of recovery, set-up, and restart time. At the time of writing (summer 2012) an hour of cloud cpu time in a VM at Amazon costs 8 cents and there are cloud providers who offer compatible interfaces for even less (see aws.amazon.com/ec2/pricing/ for pricing).

 Parallels VM software product supports a VM environment on Mac OS X, providing Mac users with access to legacy Windows applications.

The technique of “run each app in a VM” really comes into its own on servers in the cloud. Here, the standard approach is to run server-based apps, not directly on a physical server, but inside VMs that run on the server. By packaging the server-based app within a VM image, the data center gains significant system administration advantages:

- **Security** - the application is isolated in a known, controlled, locked down environment
- **Stability** - the environment won't change unless you change it (no unexpected OS upgrades pushed from the OS vendor, and no system library updates caused by installing additional applications)
- **Ease of administration** - hours of re-installation or upgrades can be completed in seconds by simply copying a virtual image file
- **Scalability** - additional instances of the app-in-a-vm can easily be started on this or additional physical servers, to meet heavy demand (scalability)
- **Defect isolation** - a crashed VM won't take down other VMs running on the same server




- **Lower cost** - it's cheaper to lease a VM running on a server than to lease the entire server

The advantages of packaging each server app inside a VM to run it are so compelling that this is the standard "best practice" within a cloud.

i Applications are run as a service that clients access through a URL. The app service is delivered from within a VM running on the physical host.

WHAT IS THE "CLOUD"?

Some people use the term "cloud" in the general sense of network transport infrastructure. In this white paper, cloud means the combination of storage and computing resources that support scalable software services delivered over the network. It's the data center, software, and virtualization, not the wiring. There are many clouds - the Facebook cloud is different from the Flickr cloud which is different from the Google™ cloud. A cloud may be open to the public, or completely private for use only by employees of a particular organization.



How you see a cloud depends on the hat you wear. To a user, a cloud is an amorphous resource, typically accessed through a browser. It can provide flat storage of books, music, and other files. Or it may provide computing resources, such as when you use Google [App Engine](#) framework to run a servlet in the cloud. It often does both file storage and computing, for example when you use Google Maps™ with Navigation (Beta) to display a driving route which is calculated and updated in real time.

To an IT architect, the cloud is the service that delivers computing and storage resources, and meets written Service Level Agreements, to a mixed community of end users.

To a data center operations manager, the cloud is:

- One or more data centers
- Each containing file servers with network attached storage
- Combined with compute servers
- Each running virtual machines
- With each virtual machine dedicated to one or more instances of a single application

To an app developer manager, the cloud brings to your mobile device:

- Almost infinite CPU capabilities
- Almost infinite storage
- Easily accessed software services running on backend servers
- The ability to upload storage or processing to servers, and easily access them later
- New opportunities for users to collaborate with each other

LEVELS OF SERVICE FROM THE CLOUD

Cloud computing is available at several levels of service. The most basic cloud model is known as “Infrastructure as a Service” (IaaS). Providers offer basic hardware, network and storage elements, usually in the form of virtual machines, accessed via IP address blocks. The utility computing resources are provisioned on demand from large pools in the data center. The client is billed per resource use, and is responsible for installing and administering the operating systems and applications. Additional service levels include Platform as a Service (PaaS) and Software as a Service (SaaS).

CLOUD SERVICE	DESCRIPTION	EXAMPLE
Infrastructure as a Service (IaaS)	Delivers basic hardware, storage, delivered over the network. The customers for this type of cloud are IT groups.	Rackspace Cloud
Platform as a Service (PaaS)	Delivers virtualized hardware with an installed operating system, programming language execution environment, database, and/or web server. Typically supports on-demand scalability - additional server instances will automatically be deployed to meet incoming demand, and turned off again when no longer needed. The customers for this type of cloud are software developers.	The Google™ App Engine framework, Salesforce.com’s CRM platform (Heroku)
Software as a Service (SaaS)	The complete stack including application software is hosted in the cloud. Users typically interact with the application software using a browser. The customers for this type of cloud are end users.	Gmail™ Google+™

Table 2
Three levels of cloud service

The Wikipedia report on SaaS rightly notes that “SaaS has become a common delivery model for most business applications, including accounting, collaboration, customer relationship management, management information systems, enterprise resource planning, invoicing, human resource management, content management, and service desk management. SaaS has been incorporated into the strategy of all leading enterprise software companies.” (en.wikipedia.org/wiki/Software_as_a_service).

THE CLOUD'S SIGNIFICANCE FOR MOBILE

The key insight is that cloud makes enormous storage and computing resources available very cheaply, and as a commodity or utility. These resources are easily accessed through URLs over the network. When the https protocol is used, all communications are securely encrypted. For developers and enterprises, the benefits of using the cloud include:

- Easy deployment of new software
- Automatic scalability to meet demand. New servers are brought online as needed.
- Freedom from provisioning or managing data centers
- Https protocol uses strong encryption to secure data

The significance for mobile device users is you can connect to that network from anywhere your device has a data or Wi-Fi connection. That means that you have access to the vast resources from almost anywhere you are likely to be. You can use storage and computing resources wherever you need them (as long as you have a network connection). The benefits are seen across the whole spectrum from individual end users to multinational enterprises and adoption by industry has been swift. There are indications that the smaller the company, the more likely they are to use cloud-based services for all their IT needs.

ARCHITECTING FOR THE CLOUD - SERVICE ORIENTED ARCHITECTURE

There is a strong trend in IT at present towards “Service Oriented Architectures” (SOA). SOA is a set of design principles for decomposing business systems into interoperable software components accessed over the network (often through URLs in a browser). The components are well-defined business functions (such as “enroll a customer” or “take payment for a sale”) that can be reused whenever that action is needed. Components can be strung together as part of a larger transaction. Components can pass data to one another in the form of XML files, but the newer JSON notation is simpler and quicker to get running. You write the server-based code in the form of servlets, just as we have done for the last dozen years.

SOA embodies a number of IT best practices including separation of concerns, information hiding, and consistency. SOA is strongly connected to use of clouds, to the point where some industry analysts regard cloud computing as the logical invention of SOA.

 Perhaps one of the strongest benefits of SOA is that it completely decouples the service from the client, and allows either one to be radically changed, redesigned, upgraded or mobilized without affecting the other.



The **REST** (Representational State Transfer) architecture guidelines are commonly used to implement SOA designs, displacing the older, unduly complicated SOAP approach. The whole world-wide web is an embodiment of the RESTful principles, recognized and codified in retrospect. REST involves clients and servers exchanging synchronous, cacheable messages and responses. It is particularly helpful for achieving scalable server components - a key attribute of the cloud.

→ There are few clear descriptions of REST, but an accessible one is: rest.elkstein.org, while other people like: tomayko.com/writings/rest-to-my-wife.

WEB SERVICES

Web services are built using REST principles. The term “web service” refers to a unit of code, hosted on a web server, that will be executed when someone browses a particular URL. The results of running that code are delivered back in the form of renderable HTML (when requested by a person) or XML/JSON (when requested by a computer). Web services are used to implement SOA.

You can write additional web services that live in your company’s cloud, and invoke them from your apps. A web search will find a number of free sample web services, like [character recognition](#) from image files, [stock quotes](#), sending a [free SMS](#) message, sending a [PDF as a fax](#) or [rendering a graph](#) from numbers. Any of these can be built into your app, expanding its features with new server-based functionality. It’s conceivable that some mobile apps can be reduced to a thin client for UI only, getting all semantic content by invoking web services in the cloud. Such an app is quicker and cheaper to write, compared with an equivalent app that does everything itself.

💡 You can easily use Java’s http libraries to connect to existing web services in your mobile apps.

COLLABORATION IN THE CLOUD

Software as a Service (SaaS) opens up new possibilities for collaboration, which has rarely featured in desktop software to date. Users of Google Drive™ are familiar with multiple people editing text documents and spreadsheets at the same time. Figure 2 shows a text document open in Google Drive.

From a browser each user can open a document in Google Drive and place their own cursor in the document to make changes or comments. Each user’s changes are immediately seen by everyone. This can be a fast way to refine and edit a document together. Figure 2 shows editorial comments of two co-workers as this white paper was under development.






Figure 2
Example of collaboration on a document in the cloud

PUSH FROM SERVER

Android™ 2.2 (API level 8, Froyo) introduced a new service to help developers send updates from servers to their applications on Android phones. The initial service was known as Cloud to Device Messaging (C2DM). In June 2012, C2DM was replaced by Google Cloud Messaging™ for Android (GCM). GCM makes it easier for mobile applications to sync data with servers. The cloud can now reach out to the mobile device and say, “Hey, sync yourself to get up to date!”

The name “push” is used because the server pushes data onto the client, rather than waiting for the client to ask for it. For example, email servers can use GCM to inform devices when new email has arrived for them, rather than waiting to be polled by the device. Pushing data notifications from the server conserves battery life better than the alternative of a background service repeatedly polling the server.

➔ Push technologies are a natural addition to the cloud framework. More information about Google Cloud Messaging for Android is available at:
developer.android.com/guide/google/gcm/index.html
developer.android.com/guide/google/gcm/gcm.html

THE FUTURE OF CLOUD COMPUTING

The cloud is good for nearly everyone with a good network connection: end consumers, and enterprises large and small. It delivers compelling economic benefits through sharing IT infrastructure. The cloud should be like any utility - it comes to you through a socket in the wall and is so reliable you don’t notice it. Enterprises are free to operate their own cloud, or to lease cloud facilities from a commercial provider, like Amazon.

STILL GROWING

Cloud use is currently big, and getting even bigger. As part of ongoing growth, we expect to see increased use of Service Oriented Architecture and web services. Forward-looking organizations will adopt an SOA approach to designing new systems to keep their options open, regardless of cloud plans.

➔ Gartner Group predicts that SaaS sales will total more than \$22 billion by 2015:
infoworld.com/d/cloud-computing/gartner-saas-market-grow-179-percent-145-billion-189583.

MOBILE CLOUD

Some analysts have proposed that there is a mobile version of cloud that differs from non-mobile cloud. They hold that “mobile cloud” is more than just cloud computing, because it converges cloud computing, cloud networking, cloud services, open wireless, seamless mobility and mobile virtualization into one common platform for future user-centric (contrasted with the current carrier-centric) mobile devices.



Mobile cloud is still an emerging technology. You can read more about it at: en.wikipedia.org/wiki/Mobile_cloud_computing and mobilecloudsummit.com

BIG DATA AND ANALYTICS

The cloud is an enabler for storing “big data” (datasets so large, they are awkward to store in older, conventional storage silos). And “big data” is an enabler for big data analytics, where data is analyzed with statistical models to extract trends and derive insights for improved business decisions.

NETWORK PERFORMANCE

The cloud depends on network access to servers. Cloud performance is limited by the latency and bandwidth of the network connection between server and client. The move to mobile devices will spur pressure for faster data rates. Speed is currently blocking home users from moving all their data into the cloud and never again having to worry about backups, sharing and remote access.

At the time of writing, North American carriers are in the middle of a transition from 3G to 4G LTE data networks, with Verizon further along than AT&T or Sprint, and T-Mobile just getting started. The carriers use different frequencies in the radio spectrum so their 4G solutions are mutually incompatible, but the speeds are standardized. The LTE specification provides downlink peak rates of 300 Mbit/s and uplink peak rates of 75 Mbit/s, compared with prevailing 802.11g Wi-Fi maximum bit rate of just 54 Mbit/s. Reading from the cloud is about 6 times faster over 4G than over WiFi. The LTE and Wi-Fi standardizing bodies both have faster standards under development, but the pressure for even more is relentless.

DATA CENTER DESIGN

The cloud is also driving some surprising innovations in data center design. In an attempt to improve density, Google™ recently patented [stackable containers of datacenters](#). With the hope of power generated through wave action, Google envisions [floating containers](#) located offshore on barges.

BOTTOM LINE FOR MOBILE DEVELOPERS

Software as a Service is now well established in IT. It is cost-effective, elastic (grows and shrinks with changing demand), and highly reliable (guaranteed service levels). Cloud computing makes large amounts of storage and processing power available on demand. Mobile devices get access to those resources from anywhere with a Wi-Fi or data connection. The increased use of cloud services from mobile devices flows from the way the two technologies complement each other. Savvy systems designers are taking full advantage of both technologies as they mobilize their enterprises.



MORE ON CLOUD COMPUTING

Visit MOTODEV for Enterprise for more resources to help you develop cloud-based mobile apps, including:

- **Podcast:** [Building Cloud-Based Enterprise Mobile Applications](#), discussing topics such as migrating legacy apps to cloud-based mobile apps, web services and mobile security.
- **Technical article:** [Introduction to Cloud and Mobile](#), outlining how app developers can access cloud services from mobile devices.

MOTODEV for Enterprise offers a wealth of resources for developing mobile apps for the enterprise. You'll find technical documentation, best practice guides, webinars, podcasts, tech support and more to support you throughout the development lifecycle.

For more information, visit: developer.motorola.com/enterprise.

