Leveraging Cloud for Enterprise Mobile Services Delivery

IEEE International Conference on Cloud Computing 2013

Steve Mastrianni
IBM Research
TJ Watson Research Center
Yorktown Heights, NY 10598
stevemas@us.ibm.com

***Abstract*— In this paper, we discuss the advantages of leveraging the cloud to provide a cost-effective and secure platform for the delivery of enterprise mobile services. We describe the various components and functions required for the delivery of those services, along with the components and functions that comprise cost-effective and secure cloud-based mobile services solution.**

I. INTRODUCTION

In today's global economy, companies have realized that it is possible to gain a competitive advantage by supplying key personnel with mobile devices such as mobile smartphones and tablets. These devices allow employees to conduct business from any location where they can easily find a suitable connection utilizing technologies such as Bluetooth, WiFi, WiMax, 3G/4G, and LTE. These devices have become very powerful, and current smartphones now have the processing power and storage capacity to run vertical enterprise applications. The value proposition for deploying these devices is greatly enhanced, in part due to the availability of high-speed wireless connectivity.

Deploying and managing these mobile devices can be a challenging task, especially for smaller businesses that might not have their own IT staff or support organization. While a small company can hand an employee a device to user, a large company must consider not only the cost of maintaining the infrastructure to support the device but the possible security risks should the user's device be stolen or otherwise compromised.

A viable and cost-effective solution is to outsource the deployment and management of the mobile devices, allowing the company to concentrate on its core business. Cloud based device management solutions make it possible to provision, deploy, and support desktops and mobile devices from any location in the world. Because the services are cloud based, the customer is not required to invest in the infrastructure to support the deployment of the mobile devices. Client support, backup, disaster and data recovery, data resiliency, patch management, application licensing, security, governance, and compliance can all be handled by the service provider.

While the benefits of deploying mobile devices are apparent, many companies have been hesitant to deploy their enterprise applications on smartphones. There are three primary reasons for this.

One reason is cost. In the current challenging business climate, cost often trumps other concerns such as security and privacy, even though the effects of a security breech could be much more costly. Deploying and supporting mobile devices requires a substantial investment in hardware and software as well as the personnel required to design, install, configure, and support the systems. Although some level of support can be handled by a self-service web portal, support for more complex issues requires skilled personnel that users can call to help resolve their issues. Keeping the support staff trained on the latest hardware and software is an ongoing effort and can represent a significant investment over time. Outsourcing mobile deployment to a third party minimizes the initial investment and helps lower the cost of entry.

Another reason for not deploying mobile devices is complexity. Assembling, configuring, and managing the various components necessary to deploy, manage, and support mobile devices and applications can be a daunting task. Many issues must be considered such as what devices should be deployed, what applications should be installed on the devices, how the devices should be backed up, and how to ensure the devices and communications are secure. A mobile device deployment needs to be designed, configured, and implemented correctly from the bottom up to insure the proper level of security and compliance.

A third reason for not deploying mobile devices, especially on a large scale, is the concern over security and data leakage. Over 60 million mobile devices are lost, damaged, or stolen each year (PR Newswire, 2011). The ability to control the distribution of sensitive information across the many types of mobile devices can be a daunting task. The ability of the device to restrict the information from being accessed by unauthorized individuals is very important, and the ability for the management system to be able to reach out to the device and disable or erase that sensitive information is a necessary component of any mobile management infrastructure. However, a majority of management systems do not have the ability to do this in a

simple, unified fashion. The mobile management infrastructure should have the ability to render any of the managed devices unusable and provide IT administrators with a triangulated position of the device if possible.

Most corporate executives now carry some type of mobile device to allow them to remain in contact with their company. Other employees such as drivers, repair technicians, government workers, police and fire officials, and military personnel also carry and use mobile devices as part of their job. Many of these same workers use a mobile device in their personal lives as well to stay in contact with friends and relatives, to access social media sites, and for location-based services. Most users have a favorite device that they are most comfortable with, and prefer that same device for business and personal tasks. The challenge for the business is to embrace those various devices but at the same time insure that they meet the required security and compliance policies.

A cloud-based mobile device management solution can help companies and organizations deploy mobile devices quickly, which can be invaluable for first responders such as, law enforcement, hazardous materials response teams, and the military. An advantage to cloud-based services is the speed in getting users deployed (Messmer, 2012). Mobile users can begin using their devices few minutes after receiving them.

II. Service Provider Architecture

In general, the service provider maintains the subscriber company's configuration information in the cloud, although the service provider must provide the ability for the subscriber's data to be co-located behind the customer's firewall. The reason for this is that some companies do not want to trust their confidential or sensitive financial data to a cloud-based service provider. There are at least two reasons for this, the first being that many cloud-based service providers have come and gone, and no company is going to trust their important information to a provider that they fear could go out of business. The second reason is that many customers are uncomfortable about having their data available to outside parties, even if it is encrypted. For this reason, the cloud-based service must be flexible enough to permit the partitioning of the data across various locations in a complete or striped form.

The service provider must provide an 'on demand' architecture that automatically grows or shrinks to meet the subscriber's needs within the contracted service level agreement (SLA). Computing resources should be allocated to provide the agreed-upon response time and resource availability. Unused resources should be coalesced in some fashion to save energy costs and to provide a capacity pool from which resources can be allocated by any subscriber within the parameters of their existing SLAs.

Providing a system and infrastructure that grows and shrinks based on demand is made possible by virtualization. Virtual machines can be created and destroyed as needed, allowing critical computing resources to be allocated to insure that SLAs are met.

The cloud-based architecture should be secure as to permit subscribers to log on without allowing them to compromise the system in any way or to affect existing SLAs.

III. Secure Access

The architecture must be multi-tenant, that is, it must be capable of simultaneously supporting multiple subscribers while providing an impenetrable barrier between subscribers. No subscriber should be allowed to view or access another subscriber's data. Cloud-based virtual systems can provide a virtually impenetrable barrier between virtual machines and virtual machine instances by employing intelligent routers and gateways, and by providing private VLANs to and from the customer premises. Using private VLANs allows one-way, secure access to the customer's directory servers that remain under control of the company's local IT organization. Because the data is encrypted, the service provider has no access to the data, and no passwords are ever sent to the service provider.

IV. Mobile Device Deployment

Two similar, but distinct forms of mobile deployment can be employed. In the first model of deployment, an enterprise supplies devices that are preloaded with the necessary applications installed and security policies configured. The user receives the device and temporary credentials, then logs in the first time and follows the installation instructions. The cloud-based service configures the mobile device for the correct level of security and compliance, and automatically installs the applications required by that user.

In the second deployment model, the user is first entitled to use the service and is sent an email with credentials to grant them access to the system. Users log onto the system, and once authenticated, the device is checked for compatibility and compliance. If the security and operational requirements for the device are met, the required software is then downloaded and installed on the user's mobile device.

When the user initially logs on with either a company-supplied device or a personal device, the cloud-based management system verifies the device type and operating system level as one of the supported device types. Companies can specify a list of the supported device types, and devices that do not meet the requirements are not

eligible to become a member of the network. The management system then verifies that the device meets the proper security and access requirements such as encryption level, password strength, installed patches, software levels, processing ability, memory size, and other parameters to insure that the device is able to run the required applications. The centralized management capability of the cloud-based system provides a single point of reference where device requirements can be specified.

V. Mobile Device Management

Key to a successful cloud-based device management implementation is the maintenance of a consistent software and hardware configuration across all deployed systems, and strict adherence to a set of policies and procedures. There is a strong contrast between large and mid-sized enterprises when it comes to implementing systems management practices, policies, and procedures, and the cloud-based management system provides a single place where global, local, or regional policies can be specified. This flexibility is especially important in geographies where local regulations can be dramatically different.

Computer systems in large-enterprise environments are tightly controlled to prevent the installation of programs or components that might affect the proper operation of the systems, and to minimize the chances of a conflict with any future application or operating system update. The IT department of a large enterprise is of sufficient size and skill to establish policies for computer access after conferring with management, and can deploy these policies effectively across the company's network. In addition to locking down its computer systems, the company might also attempt to limit potential problems by limiting users' access to certain web sites or external networks by use of deep packet inspection or blacklisting.

While desktop computer systems in an enterprise are tightly controlled, mobile devices deployed in small and mid-sized business present a unique set of problems for the cloud-based service provider.

Mobile devices such as smartphones are not usually attached in a wired fashion to the corporate network. Mobile users are often in buildings where connectivity is an issue, or on the road where they encounter areas of intermittent or no connectivity. Providing services from the cloud (or indeed from any management server) is dependent on the ability for the device to establish some form of connectivity.

While some mobile devices may be owned by the enterprise, users often ask to use their personal devices because they are familiar with them, and often have other applications, contact lists, or media files already on their device. The user could decide to install other software on the device, or download a music file. Since the smartphone is not owned by the company, it can't be locked down in the same fashion that a company-owned smartphone would be. Downloading an application or music file to the smartphone could allow a malicious application to find its way into the corporate network the next time the user connects to company's internal network.

While current smartphones have sufficient processing power and capacity to run many applications, they don't have the resources to run antivirus or malware detector software in real time while handling other smartphone functions without a noticeable degradation in performance and response time. This may require that an image of the mobile device software be sent to the cloud-based server so the image can be scanned for viruses or Trojans. that the device may be unable to detect itself.

Current mobile device and smartphone processors do not have Data Execution Prevention (DEP) that prevents the execution of a program in the data region of memory. This could allow malicious code to be injected into buffer space and executed. While this is a problem unique to the device and not the cloud server, it places an extra burden on the cloud-based server resources and could impact the overall allocation of server resources.

Removing malware or a virus could take many steps and involve user interactions. A wrong step or response could render the device unusable. It is therefore desirable to prevent the virus or Trojan from being installed rather than trying to remove it. This places an extra burden on the cloud-based server in performing pro-active scanning of e-mail, SMS, instant messaging, and video content to prevent infection.

It is possible that the device could be used for an extended period without being connected to the management server. Once the device is reconnected, the server should recognize that the device has been used without being connected to the server, and perform an extensive check for viruses, Trojans, or other malware without requiring any interaction with the user.

VI. Managing Power Consumption

Mobile devices are always starved for power. Newer devices are configured with dual-core processors, lots of memory, and a variety of sensors and peripherals such as high-resolution video cameras, Bluetooth adapters, WiFi radios, GPS transceivers, RFID detectors, barcode scanners and stereo speakers. While these devices provide desirable and useful features, they also consume battery life. Users are often faced with manually disabling these features when not required, and enabling them when the need to use them. This requires even casual users to learn to navigate sometimes complex and detailed menus on the device. It is not always obvious which features are enabled and which features have been disabled. Forgetting to disable WiFi and Bluetooth can reduce the battery life of

some newer smartphones to just a few hours before they require recharging.

A cloud-based mobile device management system can monitor the usage of the devices, gathering heuristic data on usage patterns for the device. The data collected is then used to manage the mobile device power by insuring that device features that consume the battery such as the Bluetooth radio are placed in the 'off' condition. Users don't have to keep track of features that are enabled or disabled to conserve battery life. The central management system determines the relevance of the features based on usage patterns and specific needs. The system establishes a usage profile for each device and user, building on that profile over time. It then uses this heuristic data to make sure that the device is using its power wisely.

Devices that have a GPS transceiver can provide the cloud-based management system with their location, and the system can enable certain devices based on that location and its associated usage pattern. For example, if a user goes to the local bookstore every Tuesday evening, the user could indicate that the current location has WiFi and the user would like it used at this location. When the user arrives at that location on Tuesday evening, the management system enables the WiFi radio and allows the user to get connected. The system then turns off the WiFi access at a specified time later that night, conserving battery life.

The user can enable WiFi at any location, whenever they want to, providing the policy in place at the management system allows it to be manually enabled. When enabled, the device contacts the management system sending it the GPS coordinates and time of day that can be used as data for the historical usage profile.

VII. Managing Passwords and Remote Access

Managing passwords in any environment has always been a challenge, and it is just as difficult in the mobile space. Security is more of an issue in mobile computing since the devices are small and are often lost or left behind at a restaurant or in a cab. Devices used to process or store sensitive data must be managed in a way that protects sensitive data from unauthorized access. What makes this even more difficult is that the device is often the user's own personal phone or mobile device, and the device can't be erased or disable en masse without also losing the user's personal information.

IT organizations in general support fairly large unique passwords that provide the greatest degree of entropy. While this works for a traditional desktop or server with a full keyboard, enforcing such a requirement on mobile devices with tiny keyboards represents a large inconvenience for the user. The user is required to enter this long password while walking through an airport, riding on a train, or while a passenger in an automobile. Because these passwords contain special characters and numbers, the user must often use the shift or control key to enter the characters. This sometimes takes two hands, and requires a great deal of dexterity to enter the correct information.

The mobile infrastructure must provide a way to manage strong passwords for the device while at the same time allowing the device to be used with a minimal amount of keyboard entry. The cloud-based infrastructure can accomplish this by providing an authentication mechanism based on a rolling password that changes frequently and is a result of the hash of the device, location, time, and user-supplied salt. This mechanism insures that should the device be lost or stolen, that the data will remain secure and erased should the device be incorrectly activated.

VIII. IT Process Automation

All of the steps to deploy, entitle, and maintain devices and user accounts should be scriptable and have the capability of running in an asynchronous or synchronous mode. The various processes should be described in an industry-standard form so the inputs and outputs of those processes are available, and the processes can be run and monitored by a process engine.

The management system should oversee the installation of software upgrades, firmware upgrades, policy updates and configuration changes. The operations should be orchestrated and tracked with a process automation engine. This insures all of the required changes have been successfully implemented and that the devices are operational after the changes. If an update or change causes the device or software to no longer work, the user should have the ability to roll back those changes by contacting the management server via the device or through a secure connection from an external web site. Whenever possible, the user should never be left with a device that no longer operates.

Providing IT automation is necessary to the overall cost-effectiveness of the solution. The more features that can be automated, the more cost effective the solution becomes. A significant level of automation for mobile device management can be accomplished by using a stored repository of problem signatures and the steps to remediate those problems. As problems are successfully fixed, the steps used to provide the solution are aggregated in the central knowledge base (Chefalas & Mastrianni, 2008). The information in this shared repository can then be applied to future problems and solutions, eliminating the need to diagnose similar program signatures. The solution can be later scored by the user, and that score used to select the most popular or most successful solutions from the catalog.

IX. Detecting and Eliminating Malware

In the past several years, spammers and Internet hackers and have come to the realization that installing spyware on a computer is much more rewarding than installing a virus. While viruses can render the system unusable by removing critical operating system files, spyware can be used to capture user IDs, passwords, credit card numbers, and personal information without the user's knowledge.

Key loggers, popup generators, and email zombies operate silently, stealing information and sending it to offshore servers where it is harvested and sold. It is nearly impossible to locate the thieves, and even more difficult to prosecute them because of jurisdictional and legal limitations. By the time the source of the malware is found, the perpetrators have usually packed up and left.

So far, mobile devices have not been heavily targeted for malware and spyware. Hackers have spent the majority of their time focused on the more pervasive and vulnerable Windows desktops that are then used as Spam bots and for coordinated DDOS attacks. Hackers infect the desktops by utilizing a variety of methods including flaws in the operating system, applications, open ports, phishing, and by sending infected attachments.

The tremendous increase in the deployment and use of mobile devices and smartphones is certain to generate a lot of interest among the hacker community, providing them with a huge untapped target for their malware. Some of the traditional antivirus companies such as Symantec® and McAffee® have efforts underway, but this will require users to install the antivirus software on their smartphone, and like their desktop counterparts, the antivirus programs will require constant updates to maintain protection. A cloud-based solution relieves the user from having to remember to update their protection by providing automatic updates whenever the device is connected.

Like any virus, the best protection is to not get infected in the first place. This requires the use of a cloud-based service to insure that any software being installed is free from viruses or spyware. The mobile device performs all of its connectivity tasks such as browsing and email using a cloud-based proxy. The proxy server uses a deep packet inspection to verify that the data being downloaded or sent is free from viruses and spyware.

Apple® employs this type of technology in the iTunes® ecosystem. Data loaded onto an Apple device such as an iPhone® or iPad® is scrubbed and checked for malware before allowing it to be installed on the device. Programs and other software that can be downloaded from the iTunes site are scrutinized by Apple technicians to insure that the application is not, for example, gathering sensitive data and sending it to a third party. RIM provides a similar service by routing all traffic to and from its BlackBerry® devices through a gateway that provides secure encrypted communications. Because all data flows through the gateway, it represents a single point of failure. If the gateway fails, no data can flow to or from the devices. It is therefore extremely important to provide redundancy at the gateway.

Routing data through sites such as iTunes and the Apple APNS gateway provides an added level of security. Browsers and applications delivered through sites such as iTunes can be customized to permit only certain types of data and to limit access to harmful sites. The cloud-based infrastructure facilitates the ability to centrally manage the content delivery and provides a strategic mechanism to insure that viruses, Trojans, and other malware are never delivered to the user's device. Once a particular site has been recognized to deliver malware, access to that site can be quickly disabled preventing further infections.

X. Application Execution

A cloud-based infrastructure can help companies deploy their applications to heterogeneous platforms by providing a cloud-based virtual mobile runtime. Users are not required to run platform-specific versions of business or enterprise software directly on their device, but instead interact with a virtual instance of the application running on the cloud server. This not only provides a platform-agnostic environment for running the application, but also insures that the correct version of the application is used and that no viruses or Trojans are present. The application can be kept current on the server, eliminating the need for every user to upgrade their local copy of the application.

Desktop applications can also run in the virtualized environment, providing a virus-free environment and access to desktop applications and data without the fear of data leakage. The cloud-based service should set policies about what types or forms of data can be viewed, analyzed, downloaded, or printed, further protecting sensitive information. The policies would also dictate what types of information could be accessed based on time of day, location, or other parameters. For example, an attempt to access certain sensitive information from a location that is not close to a corporate headquarters could be declined.

XI. Leveraging Analytics

Analytics can provide valuable information to a an enterprise by providing answers to questions such as how efficiently the mobile devices are being utilized or how well the software is able to solve customer problems. Data sent to and from the mobile devices are tagged and logged with information including the user, the device, date and time, application events (such as start and stop), and

location information. The information is then analyzed to determine certain usage patterns or user experience. Analytics can also be used to gather usage patterns that can be used to evaluate the speed of networks, infrastructure, and servers to help identify bottlenecks. The information can also be used to indicate areas where productivity can be improved by adding hardware, software, or network resources.

Monitoring operations in the form of application usage, network activity, data usage patterns, and system events can provide valuable information to the service provider in helping to identify problems or areas of concern. Event data from the devices or from other infrastructure devices can provide more detail and can help narrow the focus of the problem determination process.

Employing analytics in a mobile services offering is essential to contain costs and to provide a more enjoyable customer experience. The use of analytics is even more important for a mobile services offering where services are often delivered to a broad range of device types and diverse user community. The ability to respond quickly to system outages, application problems, and connectivity issues is made possible by monitoring the operation of the mobile devices and responding to problems immediately. To enhance the user experience, the use of analytics provides a way of recognizing patterns of activity that indicate that problems are likely to occur or may occur in the near future, and to take corrective action before the problem is detected.

Phone calls to tech support can be expensive for a service provider, consuming valuable resources and time. A cloud-based infrastructure can leverage the power of back-end computing resources to perform deep mining and analysis of problem trends. Once recognized, the system can take immediate action, eliminating the need to handle calls and problem tickets from end users. The proactive nature of these services can substantially reduce the cost of supporting the large number of users usually found in an enterprise deployment by reducing the number of help desk tickets or phone calls. The analytics engine and analysis models are capable of recognizing patterns of events that result in trouble tickets. A feedback mechanism can be employed to rate the success or failure of certain corrective actions, and the rules associated with those actions can be automatically updated to reflect the latest successful solution. This mechanism allows the analytics system to acquire or update certain corrective actions based on their effectiveness without the added complexity usually associated with traditional learning systems or artificial intelligence.

Analytics can also be used to identify potential problems or future service interruptions through the examination of usage patterns and temporal statistics. An inference engine with substantial resources can be used to continually analyze and correlate events and notifications over time and either make recommendations or provide programmatic remediation of problems. Further analysis of the data can also be used to provide justification for equipment or infrastructure changes or improvements.

XII. Mobile Device Security

Mobile device security represents a significant challenge for enforcement at the enterprise level. There are two models for the distribution of mobile devices. In one model, the user is given a mobile device that has been preloaded a standard set of corporate-sanctioned applications. Preloaded applications might include a corporate email client, or an application that reads a barcode and looks up the price and availability of a part. The mobile device may be a specialized hardware platform that includes a signature pad or biometric hardware.

In the other model, the device is the user's favorite smartphone that the user's desires to carry because they already have a service carrier and have their personal applications and data already installed on the device. The corporate IT administration system should be flexible enough to accommodate a wide range of device and operating system types. The system must also insure that the devices meet the IT security guidelines, and must be able to enforce the security policies without affecting the normal use of the device. Because the mobile device will likely contain sensitive corporate data in addition to personal data, there devices should provide the ability to segregate personal data from corporate data, and may provide the ability to delete the corporate data should the device become lost, stolen, or the user is no longer authorized to possess the data.

XIII. Content Filtering and Data Management

The ability to inspect data that flows to and from the mobile device helps safeguard the potential release of sensitive information. It is certainly easier to prevent data from getting on the mobile device than to remove it after the fact. A server or proxy should provide the ability to monitor data using some type of content filtering or deep packet inspection, and to provide triggers to notify IT or corporate management should there be a potentially harmful release of data to the mobile device. Classified or sensitive objects should have encrypted, attached security descriptors or metadata that defines the allowed usage or distribution policies of those objects. The security data should be used to restrict or control the flow of the sensitive information. The management system should insure compliance by verifying the security metadata or descriptors before allowing the data to be sent to the mobile device. The management system should also provide the ability to receive sensitive data that has been tagged or classified as sensitive, and refer to the IT security policy to insure that the data remains secure.

Deep packet inspection and filtering are important aspects of the management system, and can insure the integrity of the users' devices and data as described in section IX.

XIV. BYOD

In the past, enterprise deployment of mobile devices meant deploying a specific device such as a BlackBerry smartphone to users. Today, users are not as likely to adopt this strategy. Most users have already selected their own device based on personal preferences, and most already use their phone to receive e-mail and text messages. These users are not likely to carry another smartphone unless it is the same type or model as their personal smartphone. Businesses that wish to deploy mobile devices to their users must be able to accommodate the user's existing device. The management infrastructure must be flexible enough to support a wide range of devices, providing, of course, that the device can be configured to comply with the company's security and access policies. No company can allow a non-compliant device to access their network or computing resources, so the challenge for the device management platform is to provide support for these various devices while still maintaining the integrity of the corporate network.

The system must be able to enforce the security policies without affecting the normal use of the device. Because the mobile device will likely contain sensitive corporate data in addition to personal data, the device should provide the ability to segregate personal data from corporate data, and may provide the ability to delete the corporate data should the device become lost, stolen, or the user is no longer authorized to possess the data. A cloud-based management platform is not limited or constrained by power or storage, and thus can be easily designed to accommodate a wide range of devices.

XV. CONCLUSIONS

Mobile device deployment and support in an enterprise environment is a difficult undertaking without the proper systems and infrastructure. The large number of users usually associated with an enterprise deployment can place a significant burden on the support personnel. A cloud-based mobile device management platform can help defray some of the support costs normally associated with a large deployment by leveraging automated problem detection and remediation features in the cloud. The cloud can provide the capacity, resources, and intelligence to provide a minimum number of service interruptions without taxing the limited resources generally available in a mobile device. The cloud-based platform also serves as a central repository for critical performance information, device statistics, asset tracking, and successful problem remediation strategies that can be aggregated and extended to other devices or companies.

Cloud-based mobile device management allows for rapid deployment and provisioning of devices while maintaining the required level of security and governance. Users are able to get on board quickly and begin using their devices in a short amount of time without compromising the integrity of the network or corporate infrastructure. Service providers are able to offer increased performance and reliability while maintaining security and reducing support costs.

WORKS CITED

Chefalas, T., & Mastrianni, S. (2008). Patent No. US 7318226. United States.

PR Newswire. (2011, December 27). 60 Million Reasons to Protect That Mobile

Messmer, E. (2012, June 6). *Gartner: Cloud-based mobile device management (MDM) getting hot*. Retrieved jULY 26, 2012, from Computerworld: http://www.computerworld.com.au/article/426851/gartner_cloud-based_mobile_device_management_mdm_getting_hot/