A Review on the Challenges and Growth of Cloudlets

Mr R N Panda, Associate Professor ,DSMS
Dr M Dash, FIETE , MIE , Professor
Dr S S Pattanaik, Professor , FM University

Abstract

The 2010 IBM Tech Trends survey predicted that cloud computing will overtake on-premise computing and mobile software application development will emerge as the most in-demand software application development through 2015. Accordingly, it is reasonable to predict that mobile cloud computing, the niche where these two areas merge, will also transpire as a dominant force in both the development and research arenas through 2015 with the convergence of smart phones, tablets, and cloud computing. Therefore the cloud has to be moved closer to the mobile user in the form of cloudlets. Cloudlets do not have to be fixed infrastructure close to the wireless access point ,but can be formed in a dynamic way with any device in the LAN network with available resources .This paper introduces and explores the cloudlets and discusses the challenges and growth prospects of this dynamic field and proposes a new system architecture to overcome them. In this architecture, a mobile user exploits virtual machine (VM) technology to rapidly instantiate customized service software on a nearby cloudlet, and then uses that service over a wireless LAN. The mobile device typically functions as a thin client with respect to the service. A cloudlet is a trusted, resource-rich computer or cluster of computers that is well-connected to the Internet and is available for use by nearby mobile devices.

Keywords: cloud computing, mobile cloud computing, cloudlets

1.Introduction

The vision of “information at my fingertips at any time and place” was only a dream in the mid-1990s. Today, ubiquitous email and Web access is a reality that is experienced by millions of users worldwide through their BlackBerries, iPhones, Windows Mobile, and other mobile devices. Continuing on this road, mobile Web based services and location-aware advertising opportunities have begun to appear. Large investments are being made in anticipation of major profits . C*loudlet* is a new architectural element that arises from the convergence of mobile computing and cloud computing. It represents the middle tier of a 3-tier hierarchy: mobile device --- cloudlet --- cloud. A cloudlet can be viewed as a "data center in a box" whose goal is to "bring the cloud closer". A definition of mobile cloud computing in (Open Gardens, 2010) is expressed as “the availability of cloud computing services in a mobile ecosystem” and refers to an infrastructure where both the data storage and the data processing happen outside of the mobile device from which an application is launched. Having immediate access to email and data and the power to analyze, manipulate it and store it on the fly has become not only possible, but extremely sought after. More and more businesses and individuals are utilizing smartphones and tablet PCs, and mobile service providers are constantly striving to offer faster service, greater mobile abilities and longer battery-life. In order to accomplish this, mobile cloud computing is an essential ingredient.

For the Limits of Cloud Computing , an obvious solution to the resource poverty of mobile devices is to leverage cloud computing. A mobile device could execute a resource-intensive application on a distant high-performance compute server or compute cluster, and support thin client user interactions with the application over the Internet. Unfortunately, as discussed below, long WAN latencies are a fundamental obstacle. Wireless LAN bandwidth is typically two orders of magnitude higher than the wireless Internet bandwidth available to a mobile device. For example, the nominal bandwidths of the fastest currently-available wireless LAN (802.11n) and wireless Internet (HSPDA) technologies are 400 Mbps and 2 Mbps respectively. From the viewpoint of user interaction, the difference in transmission delays at these bandwidths can be very significant: 80

milliseconds instead of 16 seconds for a 4MB JPEG image. This is a huge difference for a deeply immersive application. Even if wireless Internet bandwidth improves by one order of magnitude, wireless LAN bandwidths are also poised to improve by a large amount.

2.Mobile Cloud Computing vs. Standard Cloud Computing

In traditional cloud computing data is stored in the "cloud" of the internet where web-based applications are utilized to access the data and perform various tasks. There are various categories of Cloud Computing systems offered. These include software-as-a-service (SaaS), platform-as-a-service (PaaS), and infrastructure-as-a-service (IaaS), and database-as-a-service (DBaaS). With SaaS, clients can purchase a piece of software that is accessed online, or possibly data that will be stored online. While SaaS can almost generally be applied to all types of software, in the context of cloud computing we take it to represent software that is accessed from a machine residing on the Internet as opposed to software that is downloaded and installed locally. Salesforce.com is a major SaaS provider. The PaaS type systems are especially useful for Web servers. Instead of having to purchase their own infrastructure and hosting that hardware somewhere, users can purchase the rights to host their web site in the cloud on somebody else's infrastructure. Google App Engine is a popular example of a PaaS provider. IaaS providers deliver computer infrastructure, typically at the level of a virtualized environment, as a service including storage and networking. Clients can buy these resources as a fully outsourced service. The amount of services consumed by the client determines the cost. Amazon's EC2 is major provider of IaaS services. There is a relatively new concept to the cloud computing community being termed as Database-as-a-Service (DBaaS). Microsoft recently released a beta version of SQL Server called SQL Server Data Services (SSDS). They intend this to be used in cloud environments so that customers of a specific cloud provider can still use the mainstream database they have come to depend upon without having to purchase full licenses that would have to be installed by the cloud provider. These services are charged by the hour, like the electric company charges their clients per kilowatts an hour of usage. The multitude of service platforms offered allow consumers to tailor their service subscriptions meet their specific needs and ultimately provide both operational and financial benefits to their company. In mobile cloud computing, cloud services converge with Mobile Multimedia Broadcasting (MMB). Mobile cloud computing more-or-less requires an everything-as-a service concept where software, platform, and database services are really all required to be in the cloud to provide small devices with reasonable battery-life quick access to information and the ability to manipulate large quantities of data.

3.Technical Challenges and Limitations of Cloud computing

Presently, MMB services have non-aware production and consumption synchrony, and heterogeneity. MMB services currently use push technology and the desired end result is always to meet the personalized demand of users a quickly as possible. However, when many users are trying to access the same information at the same time and/or a large amount of users need fast access to network storage, or processing capabilities, the ability to provide this information quickly and efficiently becomes a challenge due to the ignorant MMB (Li, Li, Youxia, & Wen, 2010). The battery-life of mobile devices is a constant struggle. Mobile cloud computing has significantly improved the battery-life of mobile devices; however it can be argued that the perfect balance of what is performed and stored in the cloud versus what is performed and stored on the mobile device to offer the most reasonable battery-life is still to be found. Research in this area in underway but more will need to occur. Latency and bandwidth affect the mobile cloud, as well. Wi-Fi improves latency but may decrease bandwidth when many mobile devices are present. The rollout of 4G networks and HTML5 is expected to help with latency and bandwidth as mentioned earlier. With respect to international businesses, there is the issue of lack of international and implementation standards. Devices such as the iPad or a smartphone are resource constrained and therefore a remote display protocol must be able to deliver complex multimedia graphics over wireless links and render the graphics. Data security is a long standing issue and cloud computing is no stranger to its scrutiny. The goal is to ensure that only authorized people have access to data. Traditionally, data governance models have called for individual companies to maintain information and records. With cloud computing, companies must rely on their vendors to ensure the safety of their data and trust that that they are following all the applicable IT governance and rule sets as well as their governing laws (Nkosi & Mekuria, 2010). Cloud computing differs from the traditional computing in the way that data and resources are stored. Even when outsourcing a server to somebody else's data center, the user knows exactly where the data is stored and what resources may be shared. However, cloud computing obscures such low-level details by decoupling the actual data from the physical infrastructure. It is possible that the data could be spread across multiple physical servers that also happen to be storing data for other clients on the same machine.

4.Growth Of Cloudlets

The mobile device functions as a thin client, with all significant computation occurring in the nearby cloudlet. Physical proximity of the cloudlet is essential: the end-to-end response time of applications executing in the cloudlet needs to be fast (few milliseconds) and predictable. If no cloudlet is available nearby, the mobile device can gracefully degrade to a fallback mode that involves a distant cloud or, in the worst case, solely its own resources. Full functionality and performance can return later, when a nearby cloudlet is discovered.

Cloudlet infrastructure is deployed in much the same way as Wireless Fidelity (Wi-Fi) access points. The key challenge is to simplify cloudlet management, and a proposed solution is transient customization of cloudlet infrastructure using hardware VM technology. Pre-use customization and post-use cleanup ensure that cloudlet infrastructure is restored to pristine state after each use, without manual intervention. A VM cleanly encapsulates and separates the transient guest software environment from the cloudlet infrastructure's permanent software environment. Therefore, a VM-based approach is less brittle than alternatives such as process migration or software virtualization. It is also less restrictive and more general than language based virtualization approaches that require applications to be written in a language such as Java or C#.

Although many advances in technology, mobile devices will always be resource poor, as restrictions on weight, size, battery life, and heat dissipation impose limitations on computational resources and make mobile devices more resource constrained than their non-mobile counterparts. Therefore, mobile devices still fall short to execute many rich media and data analysis applications that require heavy computation, and often also have (near) real-time constraints such as augmented reality (AR). One solution to overcome these resource limitations is mobile cloud computing . By leveraging infrastructure such as Amazon's EC2 cloud or Rackspace, computationally expensive tasks can be offloaded to the cloud. However, these clouds are typically far from the mobile user, and the high WAN latency makes this approach insufficient for real-time applications. To cope with this high latency, Satyanarayanan introduced the concept of cloudlets: trusted, resource rich computers in the near vicinity of the mobile user (e.g.near or colocated with the wireless access point). Mobile users can then rapidly instantiate custom virtual machines (VMs) on the cloudlet running the required software in a thin client fashion.

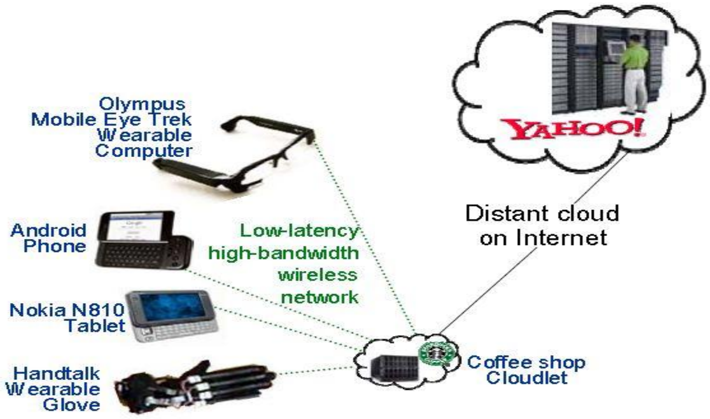

(a) Cloudlet Concept

Satyanarayanan et al. present a new vision for mobile cloud computing , VM-Based Cloudlets . They foresee a new world in which mobile computing seamlessly augments users 'cognitive abilities via intensive capabilities such as speech recognition, natural language processing, computer vision and graphics, machine learning, augmented reality, planning, and decision making.

5. Challenges Of Cloudlets

Although cloudlets may solve the issue of latency, there are still two important drawbacks of the VM based cloudlet approach. First, one remains dependent on service providers to actually deploy such cloudlet infrastructure in LAN networks. To alleviate this constraint, we propose a more dynamic cloudlet concept, where all devices in the LAN network can cooperate in the cloudlet .Next to the cloudlet infrastructure integrated in the mobile network by service providers, or provided by a corporation as a corporate cloudlet, all devices in the home network can share their resources and form a home network cloudlet. On the train, different users can also share resources in an adhoc cloudlet. A second drawback of VM based cloudlets is the coarse granularity of VMs as unit of distribution. Instead of executing the whole application remotely in the VM and using a thin client protocol, better performance can be achieved by dynamically partitioning the application in components . Also, as resources in the cloudlet will still be limited, chances are that even the cloudlet runs out of resources when many users execute their VM simultaneously on the cloudlet infrastructure. With component offloading, a more flexible allocation of the cloudlet resources is possible, so that priority is given for latency-critical parts of the application, while non real-time parts can be offloaded to a more distant cloud.

Conclusions

Wireless and mobile computing technologies provide more possibilities for accessing services conveniently. Mobile devices will be improved in terms of power, CPU, and storage. Mobile cloud computing has emerged as a new paradigm and extension of cloud computing.

In this survey, terminologies and concepts are clarified, and a definition of mobile cloud computing is provided based on an understanding of underlying technologies and applications. Existing work has also been surveyed and two classes of architecture framework are described for mobile cloud computing.

References

[1] R. Balan, J. Flinn, M. Satyanarayanan,S. Sinnamohideen, and H.-I. Yang. The case for cyber foraging. In EW 10: Proceedings of the 10th workshop on ACM SIGOPS European workshop, pages 87–92.ACM, 2002.

[2] M. Butler. Android: Changing the mobile landscape.IEEE Pervasive Computing, 10:4–7, 2011.

[3] B.-G. Chun, S. Ihm, P. Maniatis, M. Naik, and A. Patti. Clonecloud: elastic execution between mobile device and cloud. In Proceedings of the sixth conference on Computer systems, EuroSys ’11, pages 301-314 , ACM,2011.

[4] B.-G. Chun and P. Maniatis. Dynamically partitioning applications between weak devices and clouds. In Proceedings of the 1st ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond, MCS ’10, pages 7:1–7:5, 2010.

[5] E. Cuervo, A. Balasubramanian, D.-k. Cho, A. Wolman, S. Saroiu, R. Chandra, and P. Bahl. Maui: making smartphones last longer with code offload. In MobiSys ’10: Proceedings of the 8th international conference on Mobile systems, applications, and services, pages 49–62. ACM, 2010.

[6] Gartner Group. 2012 press releases. http://www.gartner.com/it/page.jsp?id=1924314.

[7] I. Giurgiu, O. Riva, D. Juric, I. Krivulev, and G. Alonso. Calling the cloud: enabling mobile phones as interfaces to cloud applications. In Proceedings of the ACM/IFIP/USENIX 10th international conference on Middleware, Middleware’09, pages 83–102, 2009.

[8] G. Huerta-Canepa and D. Lee. A virtual cloud computing provider for mobile devices. In Proceedings of the 1st ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond, MCS ’10, pages 6:1–6:5. ACM, 2010.

[9] G. Klein and D. Murray. Parallel tracking and mapping for small ar workspaces. In Proceedings of the 2007 6th IEEE and ACM International Symposium on Mixed and Augmented Reality, ISMAR ’07, pages 1–10. IEEE Computer Society, 2007.

[10] M. D. Kristensen and N. O. Bouvin. Scheduling and development support in the scavenger cyber foraging system. Pervasive and Mobile Computing, 6(6):677–692, 2010. Special Issue PerCom 2010.

[11] D. G. Lowe. Distinctive image features from scale-invariant keypoints. Int. J. Comput. Vision, 60(2).

[12] J. S. Rellermeyer, G. Alonso, and T. Roscoe. R-osgi: distributed applications through software modularization. In Middleware ’07: Proceedings of the International Conference on Middleware, pages 1–20, New York, NY, USA, 2007. Springer-Verlag New York, Inc.

[13] M. Satyanarayanan. Mobile computing: the next decade. In Proceedings of the 1st ACM Workshop on Mobile Cloud Computing & Services: Social Networks and Beyond, MCS ’10, pages 5:1–5:6, 2010.

[14] M. Satyanarayanan, P. Bahl, R. Caceres, and N. Davies. The case for vm-based cloudlets in mobile computing. Pervasive Computing, IEEE, 8(4):14 –23, 2009.

[15] M. Sharifi, S. Kafaie, and O. Kashefi. A survey and taxonomy of cyber foraging of mobile devices. Communications Surveys Tutorials, IEEE, (99):1–12, 2011.

[16] C. Szyperski. Component Software: Beyond Object-Oriented Programming. Addison-Wesley Longman Publishing Co., Inc., 2nd edition, 2002.