ACTA TEHNICA CORVINIENSIS – Bulletin of Engineering
Tome VIII [2015] Fascicule 2 [April – June]
ISSN: 2067 – 3809

Octavian-Mihai MACHIDON

ON THE SYNERGY BETWEEN DISTRIBUTED AND RECONFIGURABLE COMPUTING: CHALLENGES AND OPPORTUNITIES

"Transilvania" University of Brasov, Department of Electronics and Computer Science, Politehnicii No. 1, 500024, Brasov, ROMANIA

Abstract: *This paper presents an analysis of the existing trends in distributed systems - particularly cloud computing – regarding the integration of hardware resources using a service orientation approach and identifies solutions to this challenge based on reconfigurable computing devices. The main structural components: instrumentation, computation and network infrastructure are being analyzed, and innovative approaches are proposed regarding a service-oriented integration using reconfigurable hardware. There is a current trend in cloud computing systems for introducing a new layer in the stack architecture model, specifically at its base, namely Hardware as a Service (HaaS) that makes hardware devices accessible through services using the cloud model. For this to be achieved two key points must be addressed: simplified hardware programming - through the development of hardware description services - and enhancing the system's portability by developing a Web service-based access. Last but not least, the paper links all these integration efforts to the most critical issue of the cloud computing systems – security – and proposes solutions based on reconfigurable hardware devices for overcoming them.*
Keywords: *Cloud computing, Reconfigurable hardware, Service-oriented Architectures, Middleware, Virtual instrumentation*

INTRODUCTION

Distributed computing systems are developing and spreading rapidly, and their particular form, cloud computing, is imposing itself as the next evolutionary phase of the Internet. Cloud computing, a revolutionary concept that provides software, infrastructure and storage resources to customers over the Internet in a scalable way, as services, raises new challenges regarding the integration of hardware devices. Having to face this task - developing a new level in the cloud architectural stack: Hardware as a Service (HaaS) - solutions can be found coming from another emerging field of the last decade: reconfigurable computing. These reconfigurable devices are now widely used, with numerous applications in various fields. They provide a high degree of adaptability and scalability, providing flexible solutions for developing versatile systems by minimizing requirements for dedicated hardware and optimizing power consumption.

This paper presents an analysis of the existing trends in distributed systems - particularly cloud computing – regarding the integration of hardware resources using a service orientation approach (SOA - Service Oriented Architectures) and identifies solutions to this challenge based on reconfigurable computing devices. The integration of such hardware resources in the cloud computing infrastructure using web services has the potential to meet the requirements of cloud systems regarding variations in resource demand and workload. Cloud services could gain in configurability and become more independent from the underlying hardware resources; middleware also becomes more flexible, leading to a gain in price / performance ratio. The main structural components: instrumentation, computation and communication are being analyzed, and innovative approaches are proposed regarding a service-oriented integration using reconfigurable hardware.

THE BACKGROUND FOR A SERVICE-BASED HARDWARE INTEGRATION

The service-oriented paradigm is a recent and innovative approach in the efforts to develop new technologies for hardware integration in heterogeneous distributed systems - in this case cloud computing. The main problem that needs to be overcome is the variety and diversity of hardware resources subject to integration - from "pocket" devices to large multi-processor systems.

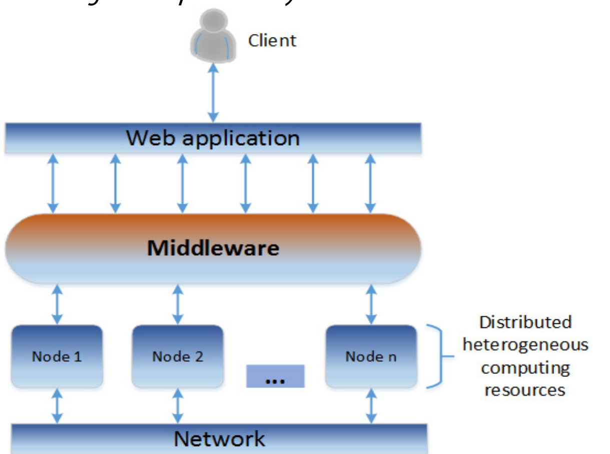

Figure 1. *Integrating middleware in distributed computing systems*

The functional characteristics of these hardware devices, such as energy consumption, storage and computing resources involved, inputs and outputs, the performance constraints (response times, real-time operation, etc.) raise the complexity of the integration effort. Thus, the efforts aim to develop a middleware that allows easy integration of different hardware resources in distributed systems using standard and neutral protocols and technologies (Figure 1).

Service Oriented Architectures and Web Services represent the latest step in the development of middleware technologies. This technology solves the problem of inter-operability and provides the basis for developing large-scale Internet applications. The term "middleware" defines an intermediate layer between the hardware (including its proprietary operating system) and distributed application that accesses it with the aim to mask the complexity of the distributed nature of the application, "hiding" away elements like memory management, network protocols and other functionalities (Geihs 2001).

Middleware technologies that allow the integration of applications are used for different purposes, from interconnecting hardware / software components of desktop or Web applications, to the development of systems that span over the Internet. Traditional technologies are quite limited when it comes to interconnecting heterogeneous software and hardware systems connected via the Internet. Web services and service-based architectures are designed specifically to meet these needs, focusing on interoperability and solving issues raised by the use of different platforms and languages (Figure 2). Thus, SOA is a solution that enables the integration of different technologies.

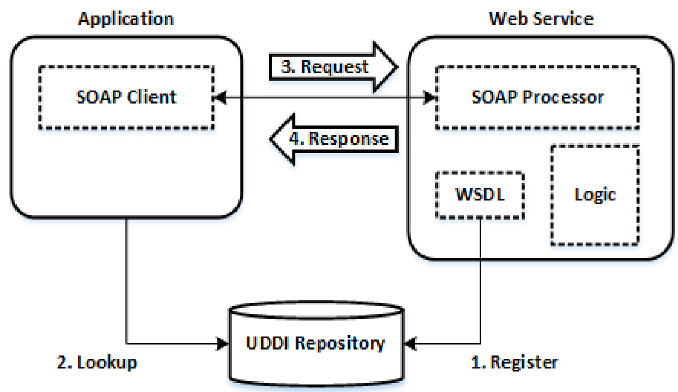

***Figure 2.** Generic functionality of a service-oriented architecture*

Services are independent and autonomous applications and not classes or components closely related to a certain application. Services are designed to be installed on a network, possibly Internet, where they can easily be integrated into the applications where they are needed. Services do not need to know anything about the clients and must accept requests from anywhere, as long as the received messages comply with format recognized by the service and the security requirements are respected. Services can be installed and managed independently one from another and client applications and service owners can modify the interface and functionality of a service at any time.

The scientific community credits service oriented architectures as a viable tool for developing a middleware based on web services, which can achieve a high-level abstraction of proprietary technologies for application developers, thus hiding the physical implementation of hardware devices or other functional aspects (network characteristics, etc.). A major advantage of hardware integration solutions based on Web services is that their architecture, relying on new standards (XML, SOAP, WSDL, UDDI, etc.), allows a unified approach to all hardware resources, despite the fact that each of them requires a specific integration methodology adapted to the particularities of the device.

It is obvious that in addition to the introduction of new technologies, a delicate process represents the ongoing transition from existing traditional architectures to new web service-based integrated ones. Numerous studies (de Deugt et al. 2006, Karnouskos et al. 2007) have been conducted and are ongoing regarding the implementation of Web service interfaces between different hardware devices and distributed corporate systems, especially given the concept of Internet of Things (IoT), which supports the integration of a variety of embedded systems using the Internet (Sommer et al. 2009).

Cloud architectural stack systems traditionally comprise three layers:

- *Infrastructure as a Service (IaaS) provides computing resources, storage and switching, relatively in a less structured way - the operating system is still in the cloud*
- *Platform as a Service (PaaS): provides tools and integrated development environments in a more structured way than IaaS, the operating system running localy.*
- *Software as a Service (SaaS) provides dedicated applications developed using stand-alone software modules, remote accessible (e.g. through APIs-Application Programming Interface)*

Cloud computing architecture enables remote access to the resources physically situated at any location on the globe, in an approach that allows for accurate metering and billing on the "pay-as-you-go" principle (Armbrust et al. 2010). Using cloud computing infrastructure reduces costs, reduces efforts licensing new software tools and increases the flexibility of business processes, virtually eliminating many of the limitations existing in the traditional approach of a computing environment: space, time, power and cost.

Despite the strongpoints listed above, there are areas that cannot benefit from a cloud computing approach, due to technological impediments regarding resources integration (Raj, Schwan 2007); especially in the case where direct access to the hardware device is needed, like in the development of embedded systems. The need for access on the device level does not permit a cloud computing approach, because such access entails that the application is required to run on the same server to which that device is physically connected (Hovestadt et al. 2012); even if this impediment is resolved, a security problem remains, since granting direct remote access makes the *virtual environment no longer isolated from the outside (Szefer, Lee 2011).*

All these issues have generated new research trend (Raj, Schwan 2007) that aims to introduce a new layer in the cloud computing stack architecture, specifically at its base: Hardware as a Service (HaaS), one that would allow the use of distinct hardware devices through services using the cloud computing paradigm.

A PERSPECTIVE ON RECONFIGURABLE HARDWARE ARCHITECTURES

Reconfigurable hardware comprises devices that can change their internal organization "on the fly". This gives a high degree of flexibility in the implementation of circuits, since the hardware resources of these devices are configurable (and usually re-configurable) after production, thus raising the possibility to implement several different circuits using the same device over time.

Reconfigurable hardware devices are constituents for a new discipline, "Reconfigurable computing", that uses such devices (like FPGAs - Field Programmable Gate Array) to implement computing systems. These systems have impressive performance and other advantages like: high processing speed, low power consumption - the circuitry being application-optimized, reduced size, and so on. Reconfigurable hardware devices have a great potential due to high adaptability and scalability, reducing the need for dedicated circuitry, optimizing energy use and minimizing hardware resources required for specific applications.

Cloud computing, being an emerging field with an accelerated growth rate, provides a number of areas where, because of the advantages listed above, reconfigurable hardware can provide significant benefits. Especially considering the increasing amounts of data being moved to such cloud systems, a hardware-based feasible solution is required.

There are also several shortcomings, despite the strengths mentioned above, that impede widespread use of reconfigurable hardware systems; two in particular have been identified by researchers (Vuletic et al. 2004): the lack of unified and standardized programming models and the difficulty of integrating these resources due to their diversity and heterogeneity. It is desirable that application developers for reconfigurable systems can do this without having to bother with low-level details of the underlying hardware. Hardware description languages (HDL) is not an attractive tool for clients who develop applications using such resources. In this context, the integration of reconfigurable devices in the cloud must be accompanied by the development of "hardware description services", including new programming models that provide a high-level development environment, making developing and running applications on reconfigurable hardware attractive for cloud computing systems.

In the broad field of reconfigurable hardware architectures, a special place is occupied by Reconfigurable SoC's (System-on-Chip). This is a solution that integrates reconfigurable hardware with a microprocessor, a synergy exploiting both the flexibility of the software design and the high performance of the hardware implementation. In this integration, the microprocessor is the one having full control over the entire system, being responsible for running the embedded software application and also for the reconfiguration of the programmable logic. These architectures, however, are subject to the same impediments, namely the lack of a unified programming methodology and a standardization of the interfaces (Vuletic et al. 2004). There is a variety of RSoC architectures, each platform having its own specific characteristics and integrated development tools (Mencer et al. 2001). The two directions of research identified as necessary to increase the attractiveness of RSoCs, namely to simplify programming and increase portability, can benefit from a service oriented approach; in the first case by the development of hardware description services (as mentioned above) and in the second case a Web service-based access for ensuring the possibility of programming, re-configuring and communicating with the remote device via the Internet. This creates an intermediary abstraction layer useful both for the programmer and the hardware engineer, thus obscuring the functional and constructive features of the hardware resources.

Service-oriented access to reconfigurable hardware facilitates the development of applications on these platforms, and a variety of fields can benefit from the advantages of this type of hardware implementations: cryptography and security solutions, digital signal processing (DSP), neural networks, control systems, etc. (Rodriguez et al. 2005).

Another application of the synergy between reconfigurable hardware and service oriented architectures is the transition from software to hardware services, namely the possibility of implementing such services on reconfigurable hardware platforms (Smith et al. 2006) ("hardware-accelerated services"). This brings substantial benefits, especially in the cases when various services are launched and executed sequentially over time.

SERVICE-ORIENTED INTEGRATION OF HARDWARE RESOURCES

Processing resources

A decade ago it was common practice in the IT industry to invest in physical equipment and store it in a hosting company's facilities. However, the rapid evolution of the Internet has made this approach obsolete, as it can no longer meet the increasing reliability and availability requirements. Also, the big "static" data centers, growing to a considerable size and complexity, have become completely ineffective in terms of performance and energy consumption due to large variations in workload over time.

The emergence of cloud computing can be seen in these conditions as a natural evolutionary step, but one which also raises new challenges, the main one being linked to integration efforts of distributed computing resources. Thus, the reconfigurable hardware field also receives a boost because of the increasing need of dynamic hardware resources capable to respond in a scalable manner to sudden load changes.

In the cloud computing market, the companies can choose between two scaling solutions for the cloud data centers: horizontal (by increasing the number of stations - computing servers) and vertical (improving the individual performance of computing devices and network components). Both encounter drawbacks: first would greatly increase the power consumption of these data centers (consumption that is very high already, 10-30 times higher than the consumption of the office infrastructure of a company like IBM (Carter 2009)), the second is not feasible in the case of multi-CPU architectures, where the native clock frequency is approximately constant. Under these conditions, the expansion of cloud data centers is a current and urgent problem that needs to be dealt with.

Reconfigurable architectures (such as FPGAs) can play an important role in overcoming the issues described above, since it allows resource scaling in the cloud to a degree that cannot be achieved using conventional processors. New efforts are oriented to finding optimal solutions for integrating these architectures in the cloud through a unified approach that would allow the integration of processing resources, development environments and communication infrastructure. This research direction is the focus of several efforts coming from both the academic and industrial environments (Madhavapeddy 2011).

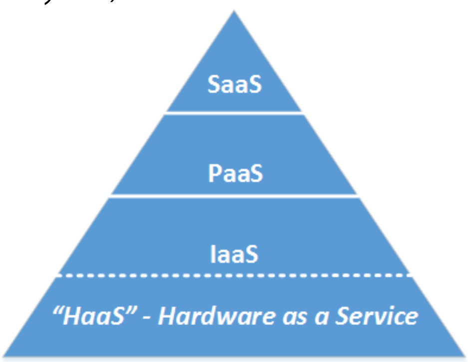

***Figure 3.** The cloud computing architectural stack having as base the emerging "HaaS"layer*

Current research proposes methodologies and platforms for integrating hardware resources as cloud computing services into a new paradigm - Hardware as a Service (Figure 3). In (Stanik et al. 2012) the authors implement service access using hardware resources distributed over different geographical locations but also inter-connected via a virtual bus. In the same manner, (Hovestadt et al. 2012) proposes ways to integrate hardware emulators using services, enabling the simulation and virtualization of a hardware system before the existence of a physical prototype.

In IaaS (Infrastructure as a Service) cloud services, equipment is provided to customers in the form of virtual machines controlled by a software hypervisor. Virtualization is therefore a key issue in a cloud computing system, each virtual machine being composed of hardware resources (CPU, memory and storage devices), and the overall performance depends on how these resources are virtualized and made available to customers dynamically and according to demand. Thus, any optimizations aim at managing hardware resource from the virtualization perspective; current studies (Sefraoui et al. 2014) identify three ways to improve the performance of cloud systems:

- *Real-time "migration" of virtual machines from one physical node to another while maintaining the service functionality*
- *Load balancing by managing the number and operation of servers for improved performance and scaling them depending on the workload*
- *The dynamic reconfiguration of virtual hardware resources during their operation providing thus a real-time scaling of computing and storage resources as needed*

The integration of reconfigurable hardware devices allows overcoming some inherent limitations of the traditional network virtualization solutions with generic microprocessor architectures. Reconfigurable hardware implementations (based on FPGAs) for network virtualization uses the FPGA to implement virtual routers benefiting from the platform's scalability that can easily adapted to possible changes in the network (Vaquero et al. 2011).

Another critical point in the cloud infrastructure vulnerable to heavy traffic situations is the reliability of the Web servers running the cloud services. Solutions to this issue were identified by implementing the Web services protocol stack in hardware - using FPGA architectures (Yu et al. 2011). This approach allows for a hardware accelerated web server to have a higher processing traffic rate, increased reliability and reduced processing time due to the pipelined implementation and direct execution in hardware, without a software operating system.

Cloud service performance can be enhanced by implementing hardware accelerated services ("Hardware Acceleration as a Service" - HAaaS) able to take over the execution of computational intensive tasks that require dedicated hardware resources and deploy them on reconfigurable devices such as FPGAs (Mershad et al. 2013). This way the execution speed increases substantially simultaneously with a decrease in energy consumption; cloud service providers can also increase their earnings by sharing these services with other providers -at their request- and by imposing higher tariffs to customers that require access to such "premium" services.

Instrumentation

Integrating instrumentation into cloud computing systems is a natural phenomenon given the raising need of remote access to a multitude of heterogeneous computing resources, communication infrastructure and measuring equipment/instrumentation. In this regard there are many implementations of instrumentation solutions integrated using service oriented architectures in grid computing systems sharable for academic and research purposes (Cheptsov et al. 2012).

Integrating instrumentation with the cloud computing concept would cause the instrument to "transcend" the physical equipment as cables and connections with a PC would be no longer necessary, and the software does not have to be tied to a specific system. The instrument would be perceived by the user as a Web page accessible using any online device, even a smartphone. Many areas could benefit from such an approach to instrumentation (biomedical, weather, energy, construction, etc.) catalyzing research in this direction. This integration involves on one hand a new philosophy in designing and implementing sensors and measurement hardware resources - which according to the Internet of Things concept must be elements with increased connectivity, always online - and on the other moving related software resources from static local systems to cloud servers for greater accessibility and performance.

In a simplistic approach (displayed in Figure 4), a cloud instrument comprises a measuring device - one or more sensors - place at a certain location, which is connected on-line by various technologies (Wi-Fi, for example). Measured data is converted into digital format and sent via the Internet to cloud servers where it is processed by specific software (control, analysis, metering, data mining, etc.), and the results are provided to users through the Web (Ghercioiu 2011).

Expansion of cloud computing in the field of instrumentation is a new research direction that describes concrete efforts in combining instrumentation and cloud computing both conceptually and technically. This research presents two possible approaches to this integration:

a. Development of models and architectures that allow the integration of instruments and sensors into cloud systems; this direction must take into account the specifics of cloud computing offering IaaS type services.
b. Development of specific instrumentation software cloud services according to the PaaS paradigm.

Currently SaaS cloud services can be used with instrumentation dedicated software, but the problem remains regarding their incompatibility with specific drivers needed for cloud computing architectures. These drivers can be considered an abstraction of the real instruments, which raises the need for a uniform integration approach of the cloud instrumentation.

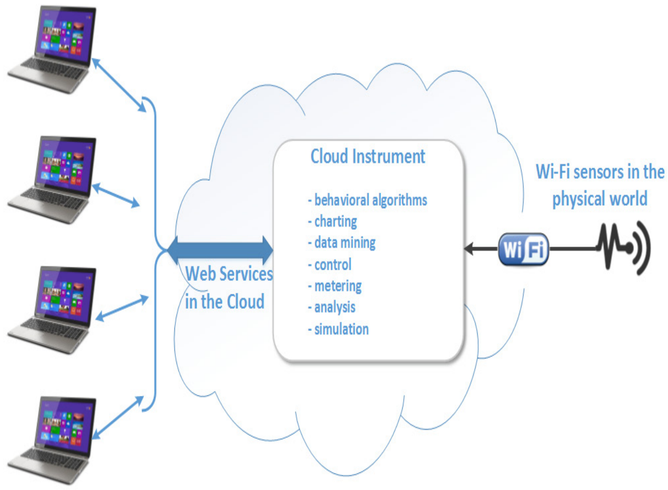

Figure 4. *Generic architecture of a "cloud" instrument*

Instrumentation as a cloud service is a conceptual paradigm that provides users with data acquisition instruments shared through the cloud computing infrastructure. Like any cloud service, instrumentation services present a uniform access and management of these resources, enabling their efficient sharing between multiple users or institutions (academic or industrial), while the instrumentation devices are available physically in different geographical locations.

The development of instrumentation in the cloud must take into account a feature that makes such services somewhat unique - the fact that instrumentation that is a process in close contact with the measured object, which binds these services to real natural elements. This raises a new problem: the difficulty of real-time sharing of instruments for multiple accesses, thus requiring the implementation of resource scheduling services.

Virtual instrumentation is suitable for an approach using reconfigurable hardware devices because each tool performs an operation well defined and bounded, allowing their easy implementation in reconfigurable hardware. The implementation of virtual instrumentation using dynamically reconfigurable hardware platforms is a solution that increases their performance while raising their flexibility because of the reconfiguration potential. It also provides a better management of the complexity of such architectures, because the structure allows the co-existence of reconfigurable integrated scalable IP cores connected to the systems busses. This is particularly important in the case of an accelerated growth in complexity, a situation when reconfigurable instruments benefit from their adaptability and dynamic reconfiguration to meet the requirements. They also allow for an optimized design methodology based on an incremental approach (instrumentation can be updated without physical access by remotely upgrading the bitstream and thus changing - totally or partially - their internal configuration).

Virtual reconfigurable instrumentation is an attractive solution especially in the academic field, having the potential to facilitate student access to real experiences in the field of instrumentation while lowering necessary costs since the same hardware platform can be used to implement a variety of virtual instruments.

Communication

In the view of cloud computing, communication infrastructure undergoes a process of diversification and development, requiring ever increasing resources and also the ability to dynamically respond in real-time to customer requests for services. The constituent elements of the communication infrastructure, namely the network nodes (routers, switches, bandwidth control devices, so-called "load-balancing devices", etc.) must move from their traditional form, with static structure, to a new approach allowing dynamic reconfiguration.

Traditionally, network components are sized according to the maximum needs for which that network is designed. In the case of cloud computing such an approach is totally inefficient because it doesn't allow for an optimal cost management of the working equipment. This is because in a cloud environment there are large variations in the workload, so there are times when equipment *designed to manage massive loads works with a very small volume of traffic. Another problem is the diversity of network services that co-exist in a cloud computing system, which hinder achieving the desired service performance because the "general-purpose" servers, while providing the desired flexibility, lack the needed performance and energy efficiency.*

For these reasons a new approach to network nodes in cloud computing communication systems is needed, since the traditional implementation cannot provide an adequate response to situations where there is a high variation in the types of services required. Four basic requirements needing to be met by network elements have been identified in this research in order for the network to provide the performance required in cloud computing systems (Hayashi, Ueno 2010):

a) Network processing speed
b) Performance scalability
c) Functionality scalability
d) Reliability

Reliability is an important factor since the network infrastructure can act as a social infrastructure, which must guarantee the reliability of services that can be critical to users.

In this context, reconfigurable hardware architectures are a potential solution for a new approach to network nodes in cloud computing systems, an approach that meets the requirements listed above. Research in this area has shown that the dynamic reconfiguration feature integrated into the network nodes enables performance scaling for virtual devices and optimization in the network processing, thus improving the bandwidth and networking resource usage in cloud computing data centers (Hayashi, Ueno 2010).

The big reconfigurable hardware manufacturers have also directed their research towards this synergy of reconfigurable systems and networking, a relevant example are the efforts made by Xilinx to reach a transfer rate of 1 TB/s (currently their FPGAs provide 100G connectivity) by increasing SerDes (serializer - deserializer) resources' speed and by providing fabric support using large width data busses (Brebner 2011).

Another project, this time open-source, covering research in network systems based on reconfigurable hardware acceleration is NetFPGA. This is a platform that has now reached the second version - offering 10G support and being based on a Xilinx Virtex 5 FPGA; also having 4 10GigE SFP interfaces, one PCI Express for connectivity to the host system (Gen2 x8 channels), and on-board SRAM and DRAM memory. Research conducted on this platform highlights the advantages of implementing network nodes in reconfigurable hardware, which enables dynamic re-programming to adapt to different scenarios and needs, which is particularly important given the specifics of cloud computing systems (Rubow et al. 2010).

NetFPGA concept has raised significant interest among researchers, with several groups worldwide working in large research universities and forums (Cambridge, SIGCOMM, SIGMETRICS, CESNET, UNSW, etc.). Experimental implementations have been developed comprising a wide range of applications: congestion control, DPI (Deep Packet Inspection) - FPX, monitoring packages - ICSI, PTP (Precise Time Protocol - for synchronizing between routers).

Networking solutions implemented using FPGA technology bring on one hand the great benefit of "field upgradeability" (one can re-configure the device remotely during run-time without physical accessing it), and on the other support for high performance packet processing and data transfer speeds higher than those obtained with software implementations (Rubow et al. 2010).

SECURITY ISSUES

Cloud computing offers many advantages through its functional paradigm, namely massive computing and storage resources made available to customers as services using public cloud servers. This brings on the issue of dealing with critical data applications that require a high security level. Data security is the main problem blocking the widespread adoption of cloud computing, since customers are reluctant to store critical data in a remote system where they lack full administrative rights and which is physically situated in an unknown location.

Due to the cloud services' features, there is a shared responsibility of data security between two entities: client and provider. What are the specific responsibilities of each depends primarily on the type of cloud services provided (e.g. in the IaaS, the customer has the responsibility to secure everything from the operating system up - data and applications, while in the SaaS they are all managed by the provider) (Ogigau-Neamtiu 2012).

The current cloud service security model requires client data to be encrypted during transit to / from cloud storage resources; however an important vulnerability arises when accessing and processing this data while it is stored physically in the cloud computing infrastructure. Security and data integrity cannot be guaranteed, as data can be accessed by malicious entities without the client knowing. Traditional software-implemented security solutions can provide limited protection and are exposed to attacks from within, exactly the kind of attacks most likely in the case of a cloud computing system. These limitations are generated by the structural and functional characteristics of systems, for example, unified memory spaces for both data and software expose the system to an attack that can change the program memory during operation.

These vulnerabilities identified above must be addressed by cloud service providers by stipulating solid security guarantees for clients in the SLA (Service Level Agreement). In order to be overcome, recent research has identified the need to move from traditional, software-based solutions, to a new perspective where trust is guaranteed by the solutions implemented in hardware. FPGA-based reconfigurable architectures are considered to have a good potential to solve a large part of the existing security vulnerabilities (Eguro, Venkatesan 2012).

A possible direction in which reconfigurable hardware architectures can provide additional security within cloud systems is their *integration as reliable computational modules. Reconfigurable hardware devices (such as FPGAs) can play such a role due to their isolated memory spaces, computing parallelism, dynamic partial reconfiguration and constructive bitstream protection methods. The device operates as a computing element with a high degree of autonomy (a system administrator cannot access the data and operations performed by the FPGA, low-level access to such a resource being impossible). Therefore, cloud customers could use these devices to run critical modules and operations.*

Such integration allows for providing enhanced security guarantees that are absent in the current architecture of cloud servers; however, besides proper integration of these trusted platforms, a new approach on data classification is required. Non-critical data can be processed using traditional cloud infrastructure and sensitive data needing extra security is offloaded to such trusted reconfigurable hardware platforms. There is research to support that FPGAs due to their closed computational environment can be considered as homomorphic encryption emulators (an emerging encryption technique that allows operations to be performed on encrypted data with unencrypted result identical to that obtained when performing the same operations on unencrypted data) (Eguro, Venkatesan 2012). Such encryption has multiple applications for cloud computing and can thus ensure total protection of the critical data (which is always encrypted), maintaining and securing the functionality of the cloud services.

CONCLUSIONS

In this paper we have highlighted the perspectives of hardware integration - especially reconfigurable resources like SoC/FPGA - using service-oriented architectures in cloud computing systems. This integration has the potential to provide technological solutions to the challenges generated by the diversity and heterogeneity of hardware resources.

The main objective is the development of a service-oriented middleware for enabling the easy integration of hardware resources in distributed systems using standardized technologies and protocols, thus achieving a high level abstraction of the technologies and specific features of the devices. Thus, service-oriented integration introduces an abstraction layer between the user and the underlying hardware and software resources, enabling a unified approach.

There is a current trend in cloud computing systems for introducing a new layer in the stack architectural model, specifically at its base, namely Hardware as a Service (HaaS) that makes hardware devices accessible through services using the cloud model. For this to be achieved two key points must be addressed: simplified hardware programming - through the development of hardware description services - and enhancing the system's portability by developing a Web service-based access.

In our perspective an important consequence of this integration is the ability to transcend services from software to hardware - implemented on reconfigurable platforms; bringing on improved performance and increased service flexibility. In cloud systems, reconfigurable hardware can act as an accelerator for services that implement complex computational and resource-consuming tasks - leading to a new sub-paradigm HAaaS: Hardware Acceleration as a Service.

This paper also analyzed issues regarding the integration of virtual reconfigurable instrumentation in the cloud. Uniform access and management of virtual instruments using the cloud model adds greater flexibility in their sharing between users and institutions, an important consideration given that they are usually situated in various geographic locations. Reconfigurable virtual instruments allow a better management of the complexity by enabling adding/removing of new functionalities using the partial reconfiguration feature.

The cloud network infrastructure also makes the subject of a new approach based on reconfigurable hardware solutions that allow the scaling of network resources according to the needs and real time traffic situations, improving the bandwidth and network resource usage in cloud data centers.

We consider that all the key points identified above must also take into consideration the security vulnerabilities of the cloud computing model. We propose that reconfigurable hardware resources integrated into the cloud can act as trusted platforms for running critical computations in secure conditions. Thus, shifting the security solutions from software to hardware is mandatory in order to provide effective counter-measures to the ever-growing security threats.

Acknowledgement

This paper is supported by the Sectoral Operational Programme Human Resources Development (SOP HRD), ID134378 financed from the European Social Fund and by the Romanian Government.

References

[1.] Armbrust, M., Fox, A., Griffith, R., Joseph, A. D., Katz, R., Konwinski, A., & Zaharia, M. (2010). A view of cloud computing. Communications of the ACM, 53(4), 50-58.
[2.] Brebner, G. J. (2011, March). Reconfigurable Computing for High Performance Networking Applications. In ARC (p. 1).
[3.] Carter, J. B. (2009, August). A look inside IBM's green data center research. In Proceedings of the 14th ACM/IEEE international symposium on Low power electronics and design (pp. 153-154). ACM.
[4.] Cheptsov, A., Koller, B., Adami, D., Davoli, F., Mueller, S., Meyer, N., ... & Kranzlmueller, D. (2012). e-Infrastructure for Remote Instrumentation. Computer Standards & Interfaces, 34(6), 476-484.
[5.] de Deugd, S., Carroll, R., Kelly, K. E., Millett, B., & Ricker, J. (2006). SODA: service-oriented device architecture. IEEE Pervasive Computing, 5(3), 94-96.
[6.] Eguro, K., & Venkatesan, R. (2012, August). FPGAs for trusted cloud computing. In Field Programmable Logic and Applications (FPL), 2012 22nd International Conference on (pp. 63-70). IEEE.
[7.] Geihs, K. (2001). Middleware challenges ahead. Computer, 34(6), 24-31. Available: http://www-di.inf.puc-rio.br/~rcerq//semGSD/sugestoes/r6024.pdf

[8.] Ghercioiu, M. (2011). Cloud Instrumentation, The Instrument is" in The Cloud". REE-Revue de l'Electricite et de l'Electronique, (2), 56.

[9.] Hayashi, T., & Ueno, H. (2010). Dynamically Reconfigurable Network Nodes in Cloud Computing Systems. NEC Technical Journal, 5(2), 143.

[10.] Hovestadt, M., Kao, O., & Stanik, A. (2012, December). Hardware as a Service (HaaS): Physical and virtual hardware on demand. In Proceedings of the 2012 IEEE 4th International Conference on Cloud Computing Technology and Science (CloudCom) (pp. 149-154). IEEE Computer Society.

[11.] Karnouskos, S., Baecker, O., De Souza, L. M. S., & Spiess, P. (2007, September). Integration of soa-ready networked embedded devices in enterprise systems via a cross-layered web service infrastructure. In Emerging Technologies and Factory Automation, 2007. ETFA. IEEE Conference on (pp. 293-300). IEEE.

[12.] Madhavapeddy, A., & Singh, S. (2011, May). Reconfigurable data processing for clouds. In Field-Programmable Custom Computing Machines (FCCM), 2011 IEEE 19th Annual International Symposium on (pp. 141-145). IEEE.

[13.] Mencer, O., Platzner, M., Morf, M., & Flynn, M. (2001). Object-oriented domain specific compilers for programming FPGAs. Very Large Scale Integration (VLSI) Systems, IEEE Transactions on, 9(1), 205-210.

[14.] Mershad, K., Kaitoua, A. R., Artail, H., Saghir, M. A., & Hajj, H. (2013, June). A Framework for Multi-cloud Cooperation with Hardware Reconfiguration Support. In Services (SERVICES), 203 IEEE Ninth World Congress on (pp. 52-59). IEEE.

[15.] Ogigau-Neamtiu, F. (2012). Cloud computing security issues. Journal of Defense Resources Management, 3(2).

[16.] Raj, H., & Schwan, K. (2007, June). High performance and scalable I/O virtualization via self-virtualized devices. In Proceedings of the 16th international symposium on High performance distributed computing (pp. 179-188). ACM.

[17.] Rodriguez, D., Sanchez, J. M., & Duran, A. (2005, November). Distributed reconfigurable computing using XML Web services. In Signal Processing Systems Design and Implementation, 2005. IEEE Workshop on (pp. 613-617). IEEE.

[18.] Rubow, E., McGeer, R., Mogul, J., & Vahdat, A. (2010, October). Chimpp: A Click-based programming and simulation environment for reconfigurable networking hardware. In Architectures for Networking and Communications Systems (ANCS), 2010 ACM/IEEE Symposium on (pp. 1-10). IEEE.

[19.] Sefraoui, O., Aissaoui, M., & Eleuldj, M.. Dynamic Reconfigurable Component for Cloud Computing Resources. In International Journal of Computer Applications 88(7):1-5, February 2014. Published by Foundation of Computer Science, New York, USA

[20.] Smith, M., Klose, B., Ewerth, R., Friese, T., Engel, M., & Freisleben, B. (2006, September). Runtime Integration of Reconfigurable Hardware in Service-Oriented Grids. In Web Services, 2006. ICWS'06. International Conference on (pp. 945-948). IEEE.

[21.] Sommer, S., Scholz, A., Buckl, C., Kemper, A., Knoll, A., Heuer, J., & Schmitt, A. (2009). Towards the internet of things: Integration of web services and field level devices. In Proceedings of the International Workshop on the Future Internet of Things and Services—Embedded Web Services for Pervasive Devices.

[22.] Stanik, A., Hovestadt, M., & Kao, O. (2012, April). Hardware as a Service (HaaS): The completion of the cloud stack. In Computing Technology and Information Management (ICCM), 2012 8th International Conference on (Vol. 2, pp. 830-835). IEEE.

[23.] Szefer, J., & Lee, R. B. (2011, June). A case for hardware protection of guest vms from compromised hypervisors in cloud computing. In Distributed Computing Systems Workshops (ICDCSW), 2011 31st International Conference on (pp. 248-252). IEEE.

[24.] Vuletic, M., Pozzi, L., & Ienne, P. (2004, September). Programming transparency and portable hardware interfacing: Towards general-purpose reconfigurable computing. In Application-Specific Systems, Architectures and Processors, 2004. Proceedings. 15th IEEE International Conference on (pp. 339-351). IEEE.

[25.] Yu, J., Zhu, Y., Xia, L.,Qiu, M., Fu, Y., Rong, G. (2011). Grounding High Efficiency Cloud Computing Architecture: HW-SW Co-Design and Implementation of a Stand-alone Web Server on FPGA. Fourth International Conference on the Applications of Digital Information and Web Technologies (ICADIWT), 2011

ACTA Technica CORVINIENSIS
BULLETIN OF ENGINEERING
ISSN:2067-3809