**"Mircea cel Batran" Naval Academy Scientific Bulletin, Volume XIX – 2016 – Issue 2**
**The journal is indexed in: PROQUEST / DOAJ / Crossref / EBSCOhost / INDEX COPERNICUS / DRJI / OAJI / JOURNAL INDEX / I2OR / SCIENCE LIBRARY INDEX / Google Scholar / Academic Keys/ ROAD Open Access / Academic Resources / Scientific Indexing Services / SCIPIO / JIFACTOR**

# AN FPGA-BASED CLOUD STORAGE GATEWAY

**Laurentiu Alexandru DUMITRU**[1]
**Sergiu EFTIMIE**[2]
**Dan FOSTEA**[3]

[1]Eng., Ph.D. (c), Military Technical Academy, 39-49 George Cosbuc Bvd., Bucharest, Romania
[2]Ph.D. (c), Military Technical Academy, 39-49 George Cosbuc Bvd., Bucharest, Romania
[3]Ph.D. (c), Military Technical Academy, 39-49 George Cosbuc Bvd., Bucharest, Romania

***Abstract:*** *Cloud storage solutions are known for their scalability, stability and easy integration. However, many companies choose classical, self-maintained storage because it can be directly controlled, in terms of physical security. With the overall long-term cost in mind, the balance shifts in favor of storage as a service, provided by a cloud infrastructure. In order to meet the security requirements of sensitive data, a gateway that bridges a company's internal storage endpoints with an external resource provider can solve the security issues. This device would be able to interact with existing interfaces and provide a controlled link with remote cloud storage services. The paper proposes such a solution, based on FPGA technology, that will provide seamless access and encryption for data that is stored off-premises.*

***Keywords****: cloud storage, cloud gateway*

## Introduction

Large storage spaces and long term storage have always been technological aspects that require special attention due to their cost, operation and maintenance procedures. Traditionally, long term storage and archives are backed up by magnetic tapes. This method has a low access speed, when compared to other media, but is cost-effective on long term and large data sets. Critical information require fault-tolerant storage which is usually achieved by online RAID arrays, in case of hard disk storage, and a well defined backup plan. Hard-disks are the most popular storage media due to their capacity and speed. In a RAID environment additional disks cover failures and errors, and therefore increase the Total Cost of Ownership – TCO. Additional hardware – enclosures, storage area networks, controllers – and various software licenses increase the overall cost.

With the rapid developments of cloud technologies, storage in cloud infrastructures can now be seen as an autonomous resource. A company can purchase storage as a service, without requiring other components, such a virtual private server to which the storage is attached. At the present moment, the most expensive part is bandwidth. Even if the storage itself is inexpensive, the amount of data transferred, inbound or outbound, is billed by the GB, in most of the cases. For the moment this makes public cloud storage provides inaccessible for large amounts of data, except with the case in which the storage is used solely as a backup space and, therefore, does not require large amounts on bandwidth. Another factor is speed. Usually the typical bandwidth varies from a few hundred MB/s to a few GB/s on premium services. Combined, the high transfer cost, per amount of data, and relatively low transfer speeds, compared to a native disk's speed, put the public cloud storage services among the last options on a system administrator's upgrade list. These shortcomings will be gradually suppressed as more and more companies are using public cloud resources in favor of operating their own infrastructure.

If a company operates its own private cloud, the above cost restrictions do not apply. Instead, the public storage service cost is now translated into TCO, maintenance and spare parts. However, there are cases in which a company wants to store its data off-site, for various reasons. The first obvious problem is data security, both virtual and physical. This is typically mitigated by employing one, or many, encryption techniques. The cloud storage provider is responsible for data integrity, in term of hardware failures. The second major issue is configuration and integration into the current software stack. This step requires high technical experience and compatible software interfaces.

Field Programmable Gate Arrays (FPGA) are reconfigurable integrated circuits that can serve many scenarios. As the name implies, an FPGA can be reprogrammed to match a specific set of physical interconnect, such as SAS or SATA (which are the most used hard-disk interfaces, at the present time). With this in mind, it's possible to imagine an FPGA–based structure that can seamlessly connect to a customer's existing

**DOI: 10.21279/1454-864X-16-I2-059**

hardware, encrypt/decrypt the data and forward it to an external site, such as a cloud storage provider. Such a device could present itself in a form of a hard-disk that provides unlimited storage, a network attached appliance or, even a PCIe card that exposes a storage medium to a standard workstation.

## Related technologies

In terms of remote storage, there are a lot of software protocols that allow transparent remote mounting of file systems. Most of them use the virtual file system (VFS) Application Programming Interface (API) of the Operating System (OS). This requires extra configuration on every machine that accesses the file system. There are lower level mechanisms such as FibreChannel, ATAoE, InfiniBand [1] or iSCSI [2] which can be used to attach storage devices to a server. Depending on various factors, among which TCO, speed and maintenance, a company may choose such a technology. Companies that operate their own private clouds or have a lot of servers usually separate the computing nodes from the storage nodes. Having dedicated, high-density, storage servers that expose resources to various endpoints, is usually a standard approach, even if it implies a higher deployment cost. Storage can be easily replicated from central points and strict access and logging can be deployed.

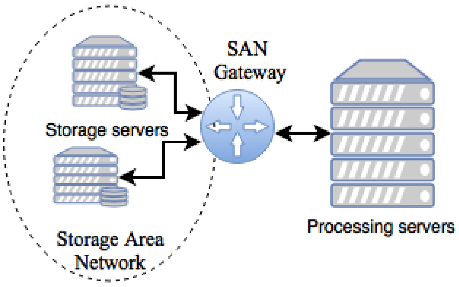

**Figure 1 - Typical SAN**

By default, high speed storage technologies are designed to be used inside the data center and, usually, they use dedicated networks – Storage Area Networks (Figure 1) which add a layer of complexity to the overall architecture. In such a setup, cloud resources are mounted directly from the software stack, bypassing the hardware resources. Some attempts such as a Fibre Channel SAN to Cloud Gateway [3] bridge the cloud directly to the hardware. iSCSI allows a server to acces a storage resource directly by Internet Protocol (IP). This means that a server can add a remote cloud resource that is viewed exactly as a physical disk. However, such and resource (called target) which is mounted on the endpoint can't be added on a normal RAID array due to various technological aspects. Jibbe Mahmoud K., et al. propose in [4] a method in which an iSCSI target can be translated to an SAS environment. By using this method, an iSCSI target could be used as a physical disk.

Remote file systems and protocols such as Network File system, File Transfer Protocol, Common Internet File Systems and others can be used when exported from cloud storage providers [5]. However, when having many dynamic virtual servers inside a cloud that require various storage space with different access lists, managing individual remote mounts can prove to be difficult task and not a desired architecture on a long-term basis. Further more, it is the administrator's job to establish a secured communication channel since most of the protocols were not designed to be used in untrusted environments.

## The FPGA approach

By addressing the two fundamental problems when using cloud storage technologies, transparency and data encryption, an FPGA – based device can bridge the gap between private companies that want to migrate their data into cloud infrastructures, but, at the same time, want to have minimal impact on their current architecture and have maximum security.

FPGA chips are highly flexible and can implement various interfaces to match any environment, given that the appropriate firmware is installed. In simple terms, a device that acts as a gateway can be viewed as a translator to/from the host side to the cloud service endpoint that also employes encryption, access control and logging facilities.

Such a device could easily be implemented using a System On a Chip (SoC) approach, with a minimal set of external elements, such as connectors and non-volatile memories. A generic architecture would contain, as in Figure 2, the host interface (A), the processing plane (B) and the outer interface(s) (C). The host interface, in the physical form, could be anything from ATA, SATA, SAS or USB, depending on the environment in which the device will be deployed. The outer interface will connect provide the communication channel to the cloud providers and the configuration channel. Physically, it can be an Ethernet RJ45, SFP/SFP+ connector or any other connector that provides a connection to a supported transport medium.

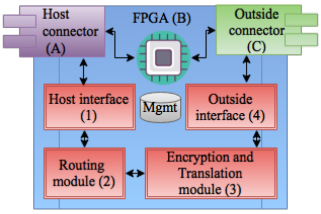

**Figure 2 - System architecture**

DOI: 10.21279/1454-864X-16-I2-059

All the firmware logic is contained in the FPGA. Depending on the chip, one can have integrated hardware media controllers or it could emulate them in software, given that all the connectors and auxiliary electronic parts are in place. Hardware PHY controllers are superior than their software versions but come at a higher cost. The internal architecture of the system is divided in modules perform specific functions according to their nature. The central module is the management module (MM) which monitors the entire system and (re)configures it when needed. It also exposes a communication channel through which a system administrator can set mandatory variables such as cloud storage endpoints. Depending on the hardware setup, a device may or may not have software modules for media access control – modules 1 and 4. If implemented, these modules are responsible with negotiating links, sending and receiving data into a specific format to and from the connected environment and forwarding that data to the next module on the chain. Since a cloud gateway should emulate one or more physical disks, it is very important to map a disk to a storage endpoint. This functionality is done by the routing module (2), which also applies access control checks, rate-limiting and logging, if such features were enabled and configured by the system administrator. After the data exists the routing module it is forwarded to the appropriate encryption and translation module. For each emulated disk, an encryption module implements an algorithm that encrypts and decrypts data before sending and receiving it to and from the cloud, assuring a transparent security layer. Optionally, the encryption service can be bypassed. This module also has the essential role of translation.

The translation service plays an important part into the system since it is the only one responsible from changing the host data format to match the one offered by the cloud service provider. Cloud endpoints could be iSCSI, NFS, FTP or other types of protocols. The translation service wraps the input data, which is usually in the form of raw blocks since the host talks with a disk, into an acceptable format to match the cloud provider's protocol. At this point, some meta-data will exist, in order to reverse the operation when the hosts issues a read request.

All logical modules inside the FPGA use either the AXI4 bus or AXI4-Stream Point-To-Point link. This assures maximum efficiency for data transfers. Other auxiliary components are also present – GPIO ports, timers, FIFOs. When implemented as a PCIe card, an appropriate IP core which bridges the physical PCIe bus to the system is instantiated. The management module can come in the form of a soft processor (such as Microblaze).

Firmware and module updates are done through the management channel which is accessible from the network to which the cloud gateway is connected. If the cloud disk connects directly to the PCIe bus, configuration can be done directly from the host, with the specific toolset. Uploading private encryption keys or setting password for encryption algorithms is also done by the management module. All other modules implement registers through which the MM can verify their operations and can issue commands to them. All of these registers are accessed through an AXI4-Lite bus. Instead of polling, each module can trigger various interrupts to signal specific events. Partial reconfiguration (PR) is used when an encryption module needs to change the certificate, key or even the whole algorithm. The partial bit-stream is sent over the management channel. Vendors provide their own tools and IP cores that implement PR logic but other implementations exist – [6],[7].

**Usage scenarios**

Most of the companies have one or more dedicated storage endpoints to which clients connect. This assures a higher data density in a single spot, facilitating in this way the backup and disaster-recovery procedures. Typically storage services are exported over the network or as dedicated links from endpoints to servers. Performance, fault-tolerance, implementation and operation costs are the key factors that should be assessed when designing the storage logic. Since most of the companies already have such solutions in place, it is important that a cloud gateway device fits without disturbing the current architecture and without the need of major reinvestments.

Network Attached Storage (NAS) devices usually export some form of file system, perhaps as a share, inside the private network. A cloud based NAS is backed by a public storage provider that is totally transparent to the users. Hard disk failures and NAS monitoring become obsolete since these are now covered by the Service Level Agreement (SLA) guaranteed by the provider. This scenario could be used in a case where the administrator does not want to change the current topology, or, when it wants to expose private cloud storage into the internal network.

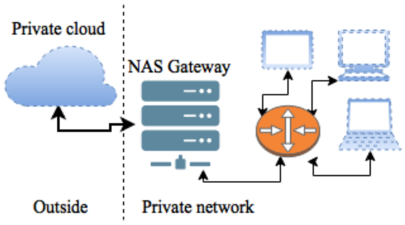

Figure 3 - Typical NAS structure

DOI: 10.21279/1454-864X-16-I2-059

Direct Attached Storage technologies usually connect one or more storage chassis to a server. The storage controller partitions the hard disks and expose certain parts to clients according to the configuration. The client views the resource as a physical disk directly connected to it. A cloud backed DAS gateway would, on the client side, be identical to a standard one, but would also connect to one or more storage providers and bridge the remote resources into the local channels. This method allows, for example, to have private servers that transparently mount remote storage spaces that seem as local disks. Since such a gateway might be built using the PCIe bus to communicate with the host and one or more up-links of various type like SFP or Ethernet to link with the provider. As presented in the previous sections, the FPGA chip can expose several interfaces to the host, thus can present itself as a physical storage device. From this layer up, the server or hypervisor can treat the device as any disk, partition it or use it in raw format.

A cloud disk is the most interesting scenario that brings cloud storage to a new level. It is a device that resembles a 2.5'' or 3.5'' disk in the sense that it has a SATA/SAS connector and a form factor identical to an actual disk, but also has, at the other end, a network connector. Inside there is an FPGA chip that contains all the logic to expose a physical disk through the host interface. All requests are forwarded inside the communication chain, as described in section III, until they reach the exit point and get transferred to the cloud storage provider. Such a disk can be seamlessly integrated into RAID arrays. The method opens a window to a wide range of scenarios. A server could have its whole storage encrypted, off-site and virtually unlimited. A RAID1 matrix could use such a disk in the array and have a transparent backup to a third party with maximum protection and minimum reconfiguration of its current configuration. Any type of RAID could have the hot spares as cloud disks. This would assure a high level of availability. More complicated scenarios offer to possibility of using multiple cloud storage providers in order to assure inter-continental data copies by using already existing storage controllers and technologies.

Regardless of the implementation method, the solution can provide standard encryption algorithms or custom ones provided by the client. The encryption module, discussed in section III, contains a reconfigurable region which can be dynamically programmed to act as a replacement for a per-defined algorithm inside it. The FPGA chip offers a high level of flexibility in order to exactly match the client's security and interfacing needs.

## CONCLUSIONS

Cloud storage services are a viable alternative to in-house storage given their overall lower costs on a long-term view. An increasing number of public and private companies are considering a migration towards cloud services, in any form, but the main concern is data security. Applying in-house security schemes before sending data off-site is the most common scenarios. However, such a setup increases the cost and maintenance requirement up to a point where such complexity overcomes the advantages of using third party providers.

FPGA–based adapters have the flexibility to address the security concerns and interoperability factors. FPGA–based cloud gateways are devices that bridge the host system to one or many storage endpoints provided by third party service vendors. By using existing interfaces, a cloud gateway, shaped as a normal hard disk, can expose a virtually unlimited storage space to a storage controller or directly to a server. Being previously configured with security credentials, the device transparently encrypts and decrypts the data that is transferred to and from the cloud. This way, the data is transparently protected in a similar way in which enterprise hard disk encrypt their content. Another advantage is that such a device can be programmed to support various protocols thus requiring minimal adjustments to a cloud storage provider's infrastructure. Configuration of such a device is straightforward as the operator needs to provide little information, such as the number of emulated physical disks, the addresses of the storage endpoints and security details.

The approach discussed in the previous sections can provide a solid solution to organizations that have not yet migrated toward cloud storage due to high technical complexity or high costs generated by hardware adjustments. Drop-in cloud drives are identified as normal disks, are fully-backed by one or more storage providers and provide transparent data encryption. This method should provide a cost-effective, transparent and secure access to cloud storage without requiring major adjustments to either a client's or a provider's infrastructure.

## BIBLIOGRAPHY

[1] Choi, Jae Woo, et al. "Towards High-Performance SAN with Fast Storage Devices." ACM Transactions on Storage (TOS) 10.2 (2014): 5.

DOI: 10.21279/1454-864X-16-I2-059

[2] Fu, Xianglin, et al. "The architecture and performance evaluation of iSCSI-based United Storage Network merging NAS and SAN." Computing and informatics 21.6 (2012): 547-561.
[3] Sabaa, Amr, Manjunath Aghalaya Gopal Gowda, and Poulo Kuriakose. "Fibre Channel Storage Area Network to Cloud Storage Gateway."U.S. Patent Application No. 13/791,415.
[4] Jibbe, Mahmoud K., et al. "Method and apparatus for enabling communication between iSCSI devices and SAS devices." U.S. Patent No. 8,892,723. 18 Nov. 2014.
[5] Wu, Jiyi, et al. "Cloud storage as the infrastructure of cloud computing."Intelligent Computing and Cognitive Informatics (ICICCI), 2010 International Conference on.IEEE, 2010.
[6] Beckhoff, Christian, Dirk Koch, and Jim Torresen. "Go ahead: a partial reconfiguration framework." *Field-Programmable Custom Computing Machines (FCCM), 2012 IEEE 20th Annual International Symposium on*.IEEE, 2012.
[7]Vipin, Kizheppatt, and Suhaib A. Fahmy. "A high speed open source controller for fpga partial reconfiguration." *Field-Programmable Technology (FPT), 2012 International Conference on*.IEEE, 2012.

**DOI: 10.21279/1454-864X-16-I2-059**