Cloud-Based Operations

A review of: Operational Framework
for Cloud-Based Networks

Author: Byron Spinney

Date: 06 February 2013

Introduction

The intent of this paper is to help the reader understand the concepts of cloud computing, the premise and the strengths and weaknesses of the reviewed paper, Operational Framework for Cloud-Based Network (OFFCBN).

To achieve this intent this paper will have a brief overview of the basics of cloud computing. The paper will then summarize the reviewed paper and discuss perceived strengths and weaknesses in the concepts stated within the paper.

Executive Summary

To obtain the best understanding of the scope, advantages, and risks of the paper it is necessary to clarify that the term “cloud” refers to a complex technology infrastructures of networks, compute engines, data, and applications. Complex technology infrastructures require diligent maintenance, monitoring, and constantly evolving protection to remain viable. This concept is critical to overlay when reading the OFFCBN paper or any other document that utilizes the term “cloud”.

The paper proposes the implementation of a “single Army network” to be nested within and fully utilize the Joint Information Environment(JIE). The concept of this network is to make it “cloud-based” establishing localized clouds (complex network and compute infrastructures) that would be forward deployed. These localized clouds would be flexible enough to interconnect with other local and non-local clouds as required.

The paper discusses utilizing multiple payers of clouds stretching from the enterprise level (data centers) to forward deployed operational theaters. The paper acknowledges potential threats of cyber, electronic, and physical attack. However, there is no mention of increased resource demands to support non-combat focused technical personnel and infrastructure demands (power, space, and cooling.)

There is no setting of expectation of the costs, time, and level of detail required to perform proper requirements analysis to establish a foundation from which to develop a sustainable architecture.

The paper proposes concepts but does not provide any guidance on mitigating risks involved with localized cloud in forward deployed theaters nor does it provide any guidance as to requirements analysis, design, standards, or metrics.

There is strong agreement, with the premise of the paper, that a more tightly integrated network of resources (i.e. sensors, compute, and data) that establishes a highly flexible deployable technology environment can provide significant advantage in the defensibility and lethality of a tactical force.

Cloud-based networking and computing implementations hold significant operational value and competitive strategic advantage. It is imperative to fully understand the technical complexities and associated architectural design considerations before committing operational environments. In no way should complexities and technical challenges keep us from moving toward the future however, a healthy and proper respect for the level of effort in the undertaking is a mandatory requirement.

Overview of Cloud Computing

There is no definitive point from which to reference however, common consensus is that the term "cloud" came from the use of a cloud shaped symbol used in architectural drawings. This symbol was, and is, used to depict complex infrastructures contained within systems diagrams. In other words the symbol of a cloud was used in place of all of the symbols and drawing connections required to depict the complex systems and subsystems within an infrastructure. It is a way to provide a "10,000 foot view" of a complex structure.

The term "cloud" has come to embody the use of compute resources (both hardware and software) which are delivered as a service over a network (in common terminology the network is by default the Internet). The concept is that services are acquired by a user from remote locations. It is notable that the user does not know the details of what service piece comes from which location nor is it important. The user acquires services (applications, server resources, and data) as a utility from a remote provider. The user must entrust the remote service and remote service provider with the data, software, and computation.

Cloud computing has been touted as a revolutionary concept of this century. With Amazon bringing on their first commercial service (Amazon Web Services or AWS) in 2006 and seeing a company called Eucalyptus provide the first platform for deploying private cloud in 2008. Gartner first acknowledge this trend in 2008[1].

However, the reality is that cloud computing had its start in the 1950's and 60's. Large-scale mini and main frame computers stored all of the data and applications, and provided all of the computational power while users connected with terminals. Terminals had no compute power. They provided the human interface (keyboard and screen) and connected via serial connections locally or remotely to the host computers. Les Earnest, a senior research scientist emeritus at Stanford and an early collaborator at the Stanford Artificial Intelligence Laboratory (SAIL) is quoted to say" A bunch of people decided that time-sharing was clearly the way to work with a computer, but nobody could figure out how to make it work for general purpose computing – nobody except John." The John referred to was John McCarthy a professor emeritus at Stanford University. John McCarthy developed the concept of timesharing in the 1950's and early 1960's.

In Douglas Parkhill's book "The Challenge of Computer Utility", published in 1966, most of the characteristics of today's cloud computing elastic provision, provision as a utility, online, illusion of infinite supply, the comparison to the electricity industry, and the use of public, private, government, and community forms, were covered in detail.

---

[1] "Keep an eye on cloud computing", Amy Schurr, Network World, 2008-07-08, citing the Gartner report, "Cloud Computing Confusion Leads to Opportunity". http://www.networkworld.com/newsletters/itlead/2008/070708itlead1.html Retrieved 2013-02-04.

Back in the 1950s scientist Herb Grosch (the author of Grosch's law) postulated that the entire world would operate on dumb terminals powered by about 15 large data centers[2].

The point of the history lesson being that the concept of computation and data as a utility is a concept that has been evolving for over 60 years.

Analysis of the Paper "Operational Framework for Cloud-based Networks"
The paper documents the main benefit of cloud computing as the physical separation of users from data and computation. The paper states that the following advantages are attained:

- Sharing of resources across a large pool of users. – Main point allows potential for centralization of infrastructure (servers, storage, switches, routers, etc).
- Capacity during peak use time increases.
- Enablement of upgrades can be achieved with less effort and time.

Analysis of stated advantages
The first advantage within the paper "Sharing of resources across a large pool of users" is really the only advantage stated within the paper. The advantages which follow in the paper exist only within the first. In other words without the centralization of infrastructure the ability to adjust for capacity demands during peak usage periods or the reduction of time and expense for upgrades would not be relevant within the paper. Centralization of infrastructure resources is the base concept on which the concept cloud computing is formed.

With this said the advantages stated are valid if the initial infrastructure designs are performed and implemented properly.

Advantages not stated in the paper include:

- The potential for increased physical security of critical compute and data resources. This can also be a risk as the criticality of these centralized resource centers makes them a prime target.
- Improved capability to establish, implement and evolve standards for hardware, software, and operations.

The paper states that the caveats regarding cloud computing are focused on network restrictions (bandwidth, service levels, and latency). This is at the heart of the risk of network centric (read cloud) computing. Some of these risks can be addressed through Concept of Operations (CONOPS) and some can be addressed through technology. However it is important to note that the caveats noted in the paper are valid and should be strongly considered in any architectural design.

---

[2] Ryan; Falvey; Merchant (October 2011), "Regulation of the Cloud in India", *Journal of Internet Law* Volume 15 Number 4 (http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1941494)

Risks Not in the Paper
There is no discussion of the requirement for exhaustive analysis of application, network, compute, and storages requirements for every level within the proposed span of architecture. This level of effort will be extremely daunting but must be completed with a very fine level of detail to prevent serious failure. A simplistic example of a potential issue would be an application with timeout sensitivity. Take a single link through a geostationary satellite with the ground station on the equator. The minimal latency would be 240ms. Once ground stations, routers, switches, systems, network interfaces, lengths of cable, and network congestion, jitter and other normal network environmental concerns are inserted in to the pathway latencies can easily exceed 500ms. Many applications are designed for use within a LAN/WAN environment with latencies well under 500ms. Experience has demonstrated that many applications experience fatal timeout issues as latencies become large.

The very nature of cloud-based architectures is to have each cloud contain complex infrastructures. These infrastructures will require high level technical personnel to maintain an acceptable operational status.

The consolidation of infrastructure increases potential for catastrophic problems. Threats to a single location increase and the value of the locations increase.

The failure potential due to technical issues (failed upgrades, patch installations, power loss, and hardware failure) increases as the complexity of the infrastructure increases.

The main concept of cloud-computing is to remove the data, applications, and computing functions from the user to a remote infrastructure, potentially spanning multiple locations. The paper agrees with this premise but then talks about the need to spread the infrastructure from the enterprise level (read traditional data center) down to forward deployed operational areas, reducing the physical separation of users from data and computation. There are many significant risks moving complex infrastructure segments into the forward deployed operational areas. These risks include:

- The need for highly technical staff which are not combat focused staff.
- The need for increased infrastructure support services to support technology demands and personnel requirements.
- Data and sensitive equipment become targets for physical overrun and capture.

Therefore it is imperative to consider reducing footprint as the technology moves farther from the enterprise level data center to the forward deployed units. Foot print includes stored or retained data (volatile and non-volatile memories, SWAP, processing, and dependence) as well as design, implementation, and operational complexities.

There will be some infrastructures that will not transfer well into a “cloud”.

Implementation of cloud-based computing can prove to be a method to contain IT costs while increasing service deliverables. However, architectural shortcuts, design and implementation mistakes, and improper deployment will result in not only lower levels of service but higher IT costs.

Detailed studies will be required to determine requirements for all network links to assure capabilities to deliver on Service Level Agreements. Inability to deliver on SLA commitments will have substantial life and safety as well as security consequences.

Examples of recent historical commercial issues include:

- A total of 568 hours of downtime at 13 well-known cloud services since 2007 had an economic impact of more than $71.7 million dollars[3]
- Amazon AWS service for NetFlix users outage 24 December 2012 caused by issues with the Elastic Balancing service.
- Amazon outage 22 October 2012 for REDDIT, Foursquare, Pintrest, and others caused by latent bug in operational data collection agent.
- Severe storm in Northern Virginia resulted in Amazon outage 29 June 2012
- Amazon outage 20 April 2011 caused by a portion of volumes utilizing the Elastic Block Store service were unable to complete read/write requests (two days to restore service.)

These issues are stated to demonstrate the complexity of cloud infrastructures and services. They also demonstrate the impact of technical issues and the demand for highly skilled technical resources to resolve problems.

Recommendations

Consider utilizing zero clients whenever reasonable to eliminate the majority of network latency issues.

Consider investigating technologies and technology roadmaps coming from the distributed Cloud Research Test Bed run by HP Labs, Intel, and Yahoo. This test be is designed to develop innovations such as cloud-computing-specific chips. HP has stated that they will conduct advanced research in areas such as intelligent infrastructure, dynamic cloud services, and scaling. “A fundamental challenge is that the cloud depends on handling increasing demand by scaling out, doing more work simultaneously in parallel,” he noted. “We know how to do that efficiently for some workloads, but it’s generally quite challenging.”

Also consider potential partnership with IBM. IBM Research launched the Research Compute Cloud, an on demand, globally accessible set of computing resources that

---

3

Essers, Loek. **Infoworld. www.infoworld.com. [Online] June 19, 2012. [Cited: Jan 06, 2013.] http://www.infoworld.com/d/cloud-computing/cloud-failures-cost-more-71-million-2007-195895.**

support business processes. The resources are distributed across the company's eight worldwide labs.

Perform detailed studies required to determine requirements for all network links to assure capabilities to deliver on Service Level Agreements. Inability to deliver on SLA commitments will have substantial life and safety as well as security consequences.

Conclusion

Cloud-based networking and computing implementations hold significant operational value and competitive strategic advantage. It is imperative to fully understand the technical complexities and associated architectural design considerations before committing operational environments. In no way should complexities and technical challenges keep us from moving toward the future however, a healthy and proper respect for the level of effort in the undertaking is a mandatory requirement.