Cable Access Redundancy: Opportunities in Virtual Deployments

Amit Singh, Principal Engineer CTAO Cable Access, amsingh@cisco.com
Cisco Systems Inc.

Abstract

The Remote PHY architecture disaggregates the PHY from the physical CCAP (Converged Cable Access Platform). This enables CCAP to be virtualized. Virtual Cable Access Platform (vCAP) is the virtualized instance of CCAP (Converged Cable Access Platform) core software. vCAP has two disaggregated software components, vDOCSIS and vEQAM. These provide DOCSIS and MPEG video services respectively. Together with a RPD (remote PHY device) vCAP provides CCAP services. vCAP is realized via NFV (Network Function Virtualization) virtual machines (VM). vCAP will be deployed in data centers.

One of the key functions of CCAP is high availability (HA). vCAP presents unique opportunities to design and implement redundancy and high availability that are not possible with custom hardware based CCAP.

This paper and accompanying presentation examine the unique aspects of redundancy in a virtualized environment, discuss their pros and cons and weigh them against the relative cost to deploy and maintain them.

INTRODUCTION

Traditional custom hardware implementations of CCAP provide redundancy and high availability via N+1 redundancy for line cards (1 standby line card backing up several active line cards) and 1:1 for supervisor cards.
The typical subscriber scaling of CCAP platforms is about tens of thousands of per instance and thousands of subscribers per line card.

vCAP breaks this model. There are no redundant line cards or supervisors. The instance of a vCAP needs to be redundant to ensure high availability. The disaggregated software centric data center environment provides unique infrastructure, such as, databases and redundant computing environments. These can be leveraged to offer new models of redundancy. For example, databases can be leveraged to store active modem state and redundant compute can be purposed for running back up instances of vCAP functions.

The allocation of storage and compute resources to redundancy is tunable per deployed instance and can be changed over time or altered via application of redundancy policy. The degree to which data center resources are allocated to redundancy services affects both the capital and operating cost.

With reduced scale of subscribers (in the order of tens or few hundreds) per RPD the new availability models become relevant.

The unique flexibility in resource allocation to redundancy results in multiple availability models which offer varying degrees of availability from almost none to availability levels on par or even exceeding as those offered by physical CCAP devices deployed today.

The remainder of the paper discusses these in detail. It is constructed with the following sections. The first section introduces the components of vCAP. The second section introduces and compares new virtual redundancy models. A possible scheme for disaster protection in case the data center is

lost is discussed followed by the mechanics of switchover and revert back.

vCAP COMPONENTS

A vCAP is composed of several components:

1. DOCSIS and eQAM VMs to realize the vCAP instance
2. Remote PHY Device
3. Network
4. Orchestrator
5. Servers

These components are physically distributed, unlike in a traditional CCAP where they are all contained in a single device. The virtual environment is dynamic as functionality could be moved from server to server, rack to rack or even DC to DC. A highly available solution has unique challenges but could also provide unique value over HA solutions in physical systems.

TIERED vCAP AVAILABILITY

If one views a vCAP VM HA from the perspective of the failure domain it is possible to construct multiple tiers of availability functionality. The deployment of these tiers of high availability depends on the number of modems and the type of service levels and service level agreements (SLAs) provided to the end user. The economics of the solution increases as the system becomes more and more available. Further, it is possible to deploy multiple tiers of HA in the network at various service nodes.

The tiers of high availability could be provisioned in a recovery policy with diminishing levels of service in case the primary recovery policies fail. The level of failure recovery could vary with the time of day or annual calendar. Data center alternatives such as vMotion could also be incorporated in the over all high availability policy. vMotion is not discussed here as information on it is available on the Internet.

vCAP high availability tiers are listed below. It is not necessary that the DOCSIS and eQAM VM comprising the vCAP instance deploy the same identical high availability tier.

1. Stateless Non-Redundant VM

This is the lowest tier of redundancy. When a VM goes down a new one is spawned on the same server. Modem state is not saved or cached. All modems will need to re-register. It is beneficial if the VM hosts a small number of modems and the deployed services can withstand the outage time of a virtual machine coming online.

The advantages are, this is the simplest redundancy policy to orchestrate. It is the most flexible and most economical redundancy scheme.

The time taken for the system to recover is the time to stand up the VM, for it to establish communication with the RPD, synchronize configuration with the RPD, and for modems to register.

Synchronizing configuration could be very quick if the RPD and vCAP maintain a checksum of the configuration file. A possible checksum algorithm could be a SHA512.

The RPD and vCAP just need to quickly compare the SHA512 checksum values to validate that the configuration file is the same on either side.

2. Stateless Redundant VM without Pairing

This is the next higher tier of service, where a pool of standby VMs is spawned a priori. The number of standby VMs is a fraction of the primary functioning VMs. The ratio of standby VMs to primary VMs is set by the operator and could change dynamically

depending on the time of day, annual calendar or the current failure rate.

One of the standby VMs takes over when an active VM fails. It will be a little quicker than the previous case as the VM is already up and running. The VM will still need to establish connections with the RPD, synchronize configuration with the RPD and modems will need to re-register.

3. Stateless Redundant VM with Pairing

In this case, a standby VM is spawned for every active VM and pairing information is maintained. Backup VMs could be over subscribed on backup servers. The degree to which backup servers are over subscribed is tunable via policies based on deployment geography, service tier, time of day and annual calendar.
When the primary VM fails the backup VM takes over. The backup VM will pre-establish redundant tunnels with the RPD.
This scenario is quicker than the previous ones, as the backup VM is already up and running, the tunnels are established and the RPD configuration is also synchronized a priori. The time to recover service is the time it takes for modems to re-register.
When the active VM re-spawns a revert back to the active server will need to be affected to recover over subscribed backup server resources if backup VMs are over subscribed on a server.

4. State-full Non-Redundant VM

In this case, all the state of the VM is stored outside the VM in a database. When the VM goes down, it is re-spawned. The VM reloads its state from the database, re-establishes tunnels with the RPD, does configuration synchronization checks with the RPD and re-establishes service to modems. Some modems might need to re-register depending on the time it takes to detect a failure and spawn a VM.
This solution will require access to high-speed databases for storing VM state. The database could itself be replicated to alternate data centers in case it might be desirable to protect against data center outages. This scenario is described in detail in the following sections.
Here server resources are economized as there are no backup servers deployed, however storage for databases will be needed. The database itself would need to be replicated so it is highly available otherwise the scenario mimics the first scenario where the VM recovers statelessly. The time to recover services will be higher.

5. State-full Redundant VM

This is the highest level of availability. A backup VM is spawned for each active VM. The active VM constantly updates its state in an off VM database. The backup VM synchronizes active VM state from the database. The backup VM establishes tunnels with the RPD a priori. Service outage time is minimal and in the same order as that of purpose built CMTS hardware.
Additionally, the database could be replicated to databases in other data centers to protect against data center outages. Further, it is possible to host multiple backup VMs in the same over subscribed backup server, where only a few of them are expected to become active simultaneously. This helps achieve better high availability economics of the entire solution. The number of backup VMs spawned on the backup server can be controlled via policy such as being tunable by the operator dynamically based on time of day and annual calendar. Figure 1 shows the arrangement of active VMs, backup VMs & the databases.

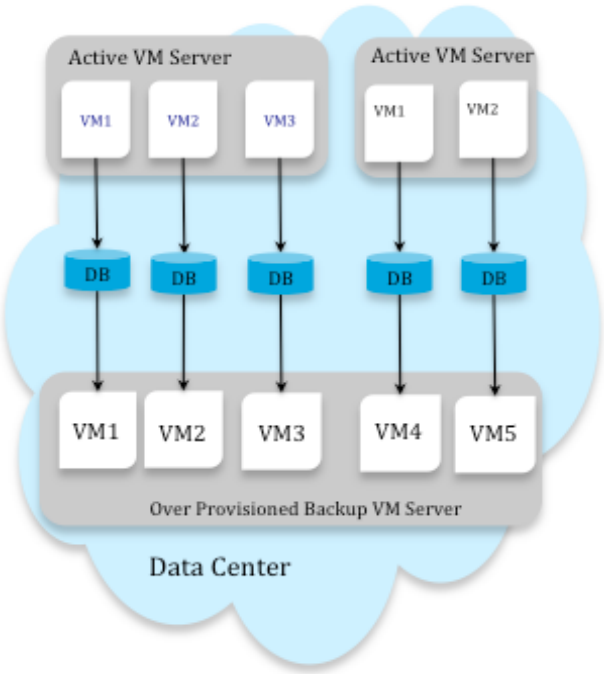

Figure 1: Active & Backup VM Instances

The backup instance does not take much server resources and therefore many virtual standby instances can be hosted on a fraction of the servers used for the active VMs. N+1 can be emulated by creating N instances of vCAP that are on dedicated servers all backed up by a single server. The key statistical assumption is the same as with a traditional CMTS: that N is small enough so that its very unlikely that the backup VM server will have more VMs active than it can support simultaneously.

Further, N is a customer tunable knob, depending on how frequently they experience failures on the vCAP VM. They can also tune N dynamically based on a calendar, for example during important events (such as the Super Bowl) N could be low, to provide higher fault resilience at the cost of more back up servers.

For business services N could be lower than for residential services, thereby providing a higher service level guarantee to business customers.

This flexible and tunable nature of the redundancy mechanism is not offered in the real hardware product. It would require very significant re-architecture of the existing product to offer this level of flexibility and it would cost more in terms of sunk costs of deploying the solution.

If the scale of the VM is small in terms of the number of subscribers it supports a new load of software could be deployed sooner on a limited number of sites with availability cranked up, so if it fails the backup can quickly take over.

The vCAP approach can also simplify the software. Instead of managing N+1 redundancy we mange a collection of 1:1 redundant instances which are over subscribed on the back up server to derive scalable value.

The cost of deploying this flavor of high availability is the highest compared to the rest as it needs backup servers to be deployed and needs databases to manage VM state.

vCAP AVAILABILITY BENEFITS

The other key benefits to arranging HA in this manner are:

1. **Database:** State is stored in an external database. The data center infrastructure inherently provides redundancy of the database. This database redundancy infrastructure can be leveraged as is. This allows for a singe vCAP to have multiple backups. For example, one can keep on backup instance locally to protect from software crashes and at the same time keep another instance in a different part of the data center to protect against hardware failures or at remote data center to cover for the entire data center. Though using a database to store state is not unique, the usage in this case is unique for a couple of reasons, the data base is a transitional

store between VMs that are all running hot and synchronous, with one VM acting as the source of data and the remaining sinks. Also, the database is live and dynamic with real time information such as modem states, service flows, voice sessions, IP video over DOCSIS and subscribers. It can be used to mine and analyze plant and service information.

2. **Backup Tunnels:** Backup tunnels are pre-provisioned so that paths do not have to be calculated on the fly. Intervening switches and networks do not need to be reprogrammed or traffic engineered.

3. **Disaster Recovery:** The approach allows scalable disaster recovery, where the same database is replicated to an alternate data center. Disaster recovery VMs are deployed at even higher over provisioning levels for cost effectiveness. This scenario is depicted in Figure 2

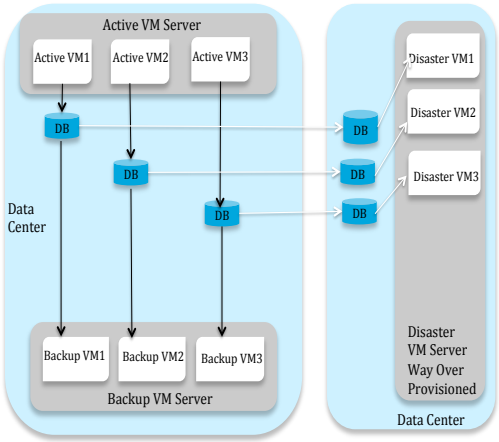

Figure 2: Disaster Recovery via Availability

It is not expected that disaster recovery would switch live modems. RPDs will need to be instructed to connect to VMs in the alternate data centers. All modems will re-register. The network between the disaster recovery data center and the RPDs may not be optimized or traffic engineered to provide an equivalent experience to the primary data center. The service levels of each one of these components are choices for the services provider. They can be tuned via policy, which may include geography or end user service level agreements. Today, there is no cost effective approach to disaster recovery deployed in hubs and head ends with physical equipment.

VM SWITCHOVER MECHANICS

The following sections discuss the mechanics of performing a switchover and revert back in a virtual environment.

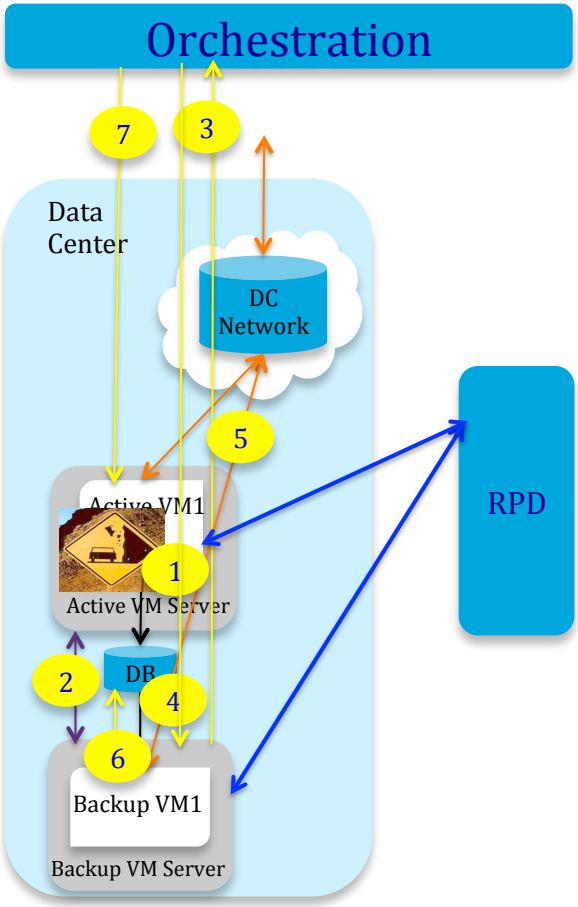

Figure 3: VM Switchover

Figure 3 depicts the typical processing steps for switchover.

1. Active VM1 has an accident.

2. Keep alives to backup VM & orchestration fail; RPD tunnels go down.

3. Orchestration notices keep alive failure OR backup VM notifies orchestration VM1 went down.

4. Orchestration notifies backup VM to take over.

5. Orchestration re-routes backhaul traffic to backup VM.

6. Backup VM becomes DB master & keeps the database updated.

7. Orchestration cleans up active VM1 & spawns new VM1 (waits for 1588 Sync); backup VM1 connects to DB to sync state back.

Figure 4 depicts the typical processing steps for revert back. As the backup server is over provisioned revert back is necessary.

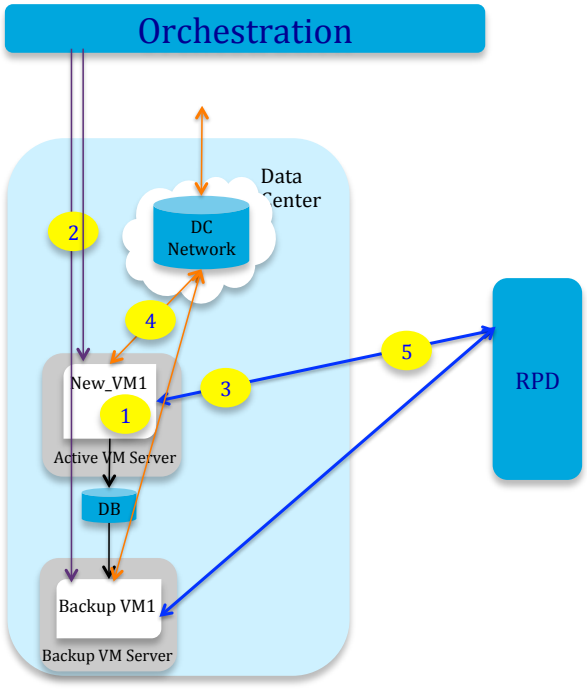

Figure 4: VM Revert back

1. New VM1 initializes & waits for hold over time (~5 minutes). The tunnels/keep alives are up and running. The configuration on New VM1 is applied and synchronized with RPD configuration. NewVM1 connected to the database of the backup VM and is continuously synchronizing real time state from the database.
2. Orchestration notifies NewVM1 & Backup VM NewVM1 is going active.

3. NewVM1 Notifies RPD Its Going Active (RPD tunnel switch is quicker than orchestration re-route). New VM1 becomes the database master and updates state to the database.

4. Orchestration re-routes backhaul traffic to NewVM1. All downstream traffic now goes through NewVM1.

5. RPD starts sending US traffic to NewVM1. The backup VM reverts to backup mode where it synchronizes state from the database.

CONCLUSIONS

The remote PHY architecture enables virtualization of CCAP functions. Virtualization of CCAP functions via a disaggregated vCAP in turn offers new tiers of service, end user SLAs, service availability and tunability via policy. It offers new possibilities in the economics of deployment, service monetization and allows rapid deployment of new services.

REFERENCES

US Docket 03919.0897 (997928): Virtual Cable Modem Termination System Redundancy in a Cable Modem Network Environment, *John Chapman, Jan Medved, Alon Bernstein, Amit Singh, Sep 2015*.