Your Mainframe Tape in the Cloud

Rebecca Alexander
Russell Witt
CA Technologies

Thursday, August 7th
Session Number 16178

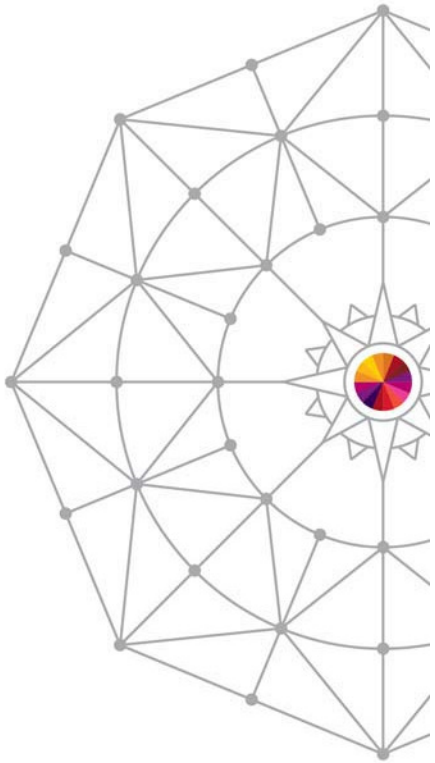

#SHAREorg

SHARE is an independent volunteer-run information technology association that provides education, professional networking and industry influence.

Copyright (c) 2014 by SHARE Inc.
Except where otherwise noted, this work is licensed under http://creativecommons.org/licenses/by-nc-sa/3.0/

Abstract

Need to reduce the cost of managing storage while improving storage utilization? Need a pay-as-you-go storage solution to keep up with business demands? CA Technologies product experts will discuss how you can leverage public, private or hybrid-cloud storage on-demand for your ever changing storage needs. CA Cloud Storage for System z will provide secure, affordable and scalable disaster recovery options to retrieve data remotely and without calling your disaster recovery vendor to manually deliver your tape for restore at your backup site. Join the discussion on benefits of moving to a cloud solution vs. managing and maintaining tape hardware and media on site.

For Informational Purposes Only

terms of this presentation

This **presentation was based on current information and resource allocations as of November 2014 and is subject to change or withdrawal by CA at any time without notice. Notwithstanding anything in this presentation to the contrary, this presentation shall not serve to (i) affect the rights and/or obligations of CA or its licensees under any existing or future written license agreement or services agreement relating to any CA software product; or (ii) amend any product documentation or specifications for any CA software product.**

The development, release and timing of any features or functionality described in this presentation remain at CA's sole discretion. Notwithstanding anything in this presentation to the contrary, upon the general availability of any future CA product release referenced in this presentation, CA will make such release available (i) for sale to new licensees of such product; and (ii) to existing licensees of such product on a when and if-available basis as part of CA maintenance and support, and in the form of a regularly scheduled major product release. Such releases may be made available to current licensees of such product who are current subscribers to CA maintenance and support on a when and if-available basis. In the event of a conflict between the terms of this paragraph and any other information contained in this presentation, the terms of this paragraph shall govern.

Certain information in this presentation may outline CA's general product direction. All information in this presentation is for your informational purposes only and may not be incorporated into any contract. CA assumes no responsibility for the accuracy or completeness of the information. To the extent permitted by applicable law, CA provides this presentation "as is" without warranty of any kind, including without limitation, any implied warranties or merchantability, fitness for a particular purpose, or non-infringement. In no event will CA be liable for any loss or damage, direct or indirect, from the use of this document, including, without limitation, lost profits, lost investment, business interruption, goodwill, or lost data, even if CA is expressly advised in advance of the possibility of such damages. CA confidential and proprietary. No unauthorized copying or distribution permitted.

CA Technologies: The Vision

Be the market-leading provider of a comprehensive platform for managing and securing private, hybrid, and public cloud environments and applications

Cloud Storage – Private / Public

- Private Cloud – A storage solution based within the corporate firewall and under control of the IT department. Usage is by departments, agencies or users within the corporation is done in a similar fashion that would be done if it was a public cloud.
  - But what about a storage device that exists at a remote location with a dedicated network path versus a public network path?
  - Then, instead of a remote location, if is housed at a third-party location but still owned and controlled by the IT dept?
- Public Cloud – A storage solution outside the IT department that is contracted to supply storage on-demand to the various departments, agencies or users within the corporation.

Cloud Storage – Hybrid

- A Hybrid Cloud is when a Private Cloud is supplemented with a Public Cloud. Ideally, this implementation would be done so that it behaves as if a homogeneous solution.
- Most often this is done by using a proprietary commercial storage appliance that serves as both a storage appliance in-house (Private Cloud) as well as a gateway to a Public Cloud storage solution.
- Offers the speed of a Private Cloud implementation for in-house uses; but also the advantages of a public off-site storage repository for DR and/or long-term storage.

The Storage Challenge

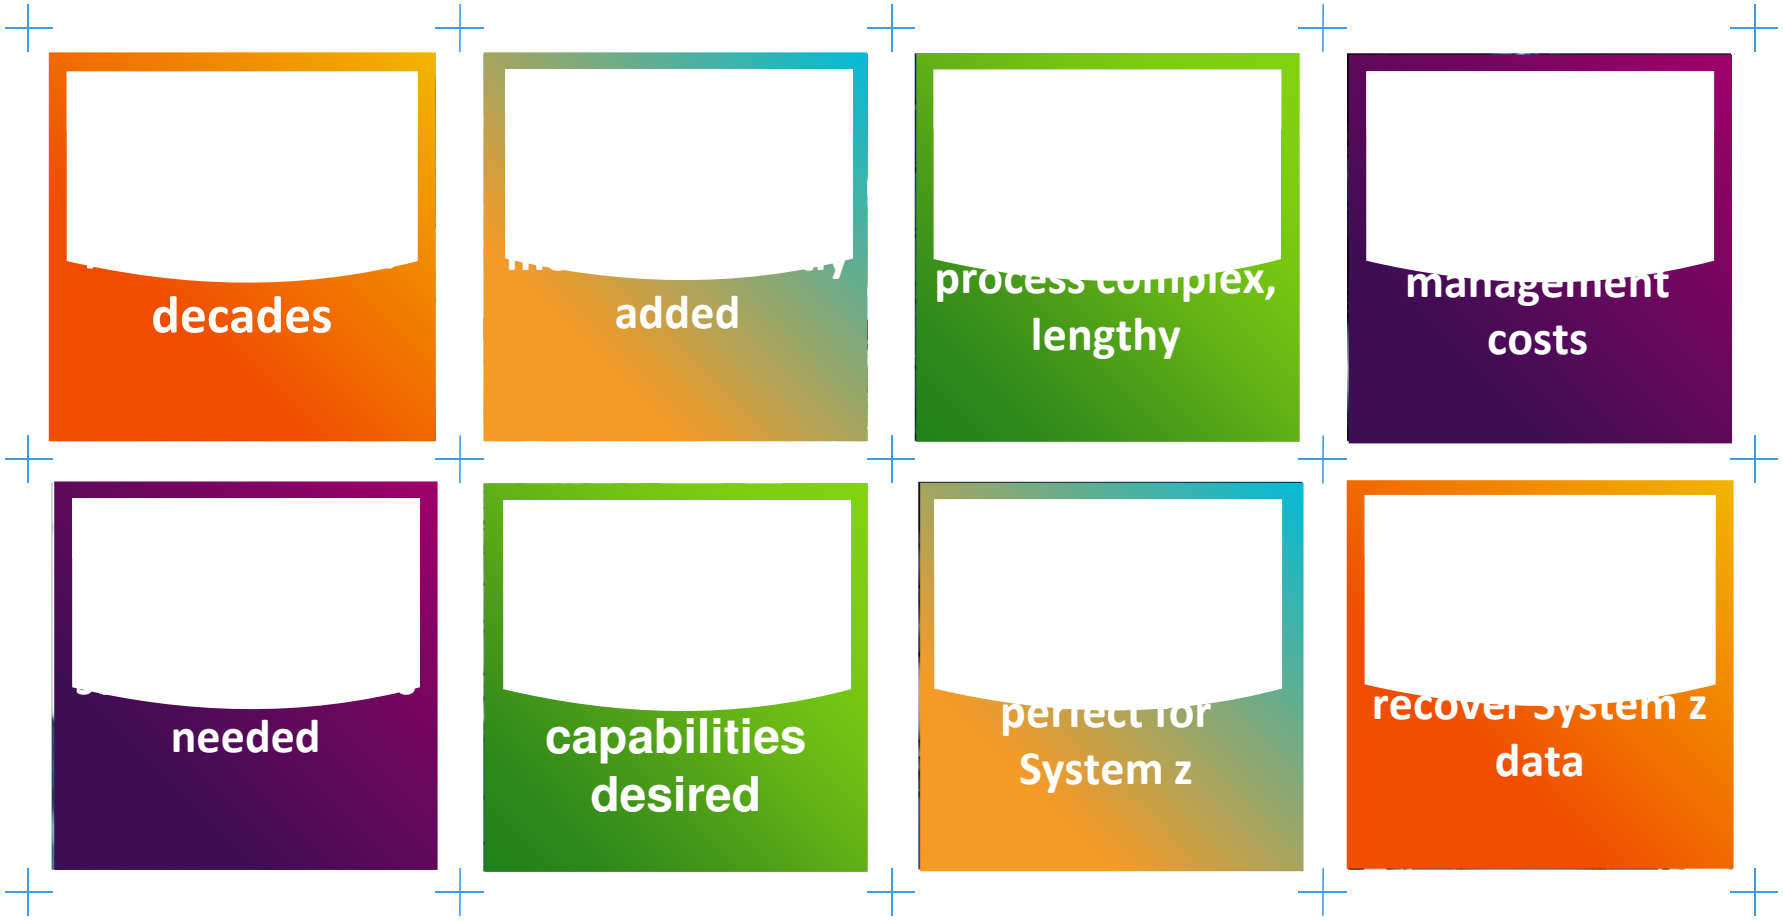

Cloud Storage – What type of MF Tape Data

- z/OS Tape Data can be classified in multiple ways
- How many copies and stored where?
  - One copy in-house (easy to recreate if needed)
  - Multiple copies with at least 1 off-site (DR or critical)
- How long does the data need to be retained?
  - Short Retention – DR backups, incremental backups, data passed between jobs (measured in days/weeks)
  - Medium Retention – log files, transaction files (measured in months/years)
  - Long Retention – legal issues, regulations, historical records (measured in years/decades)

Traditional MF with replication to cold-site

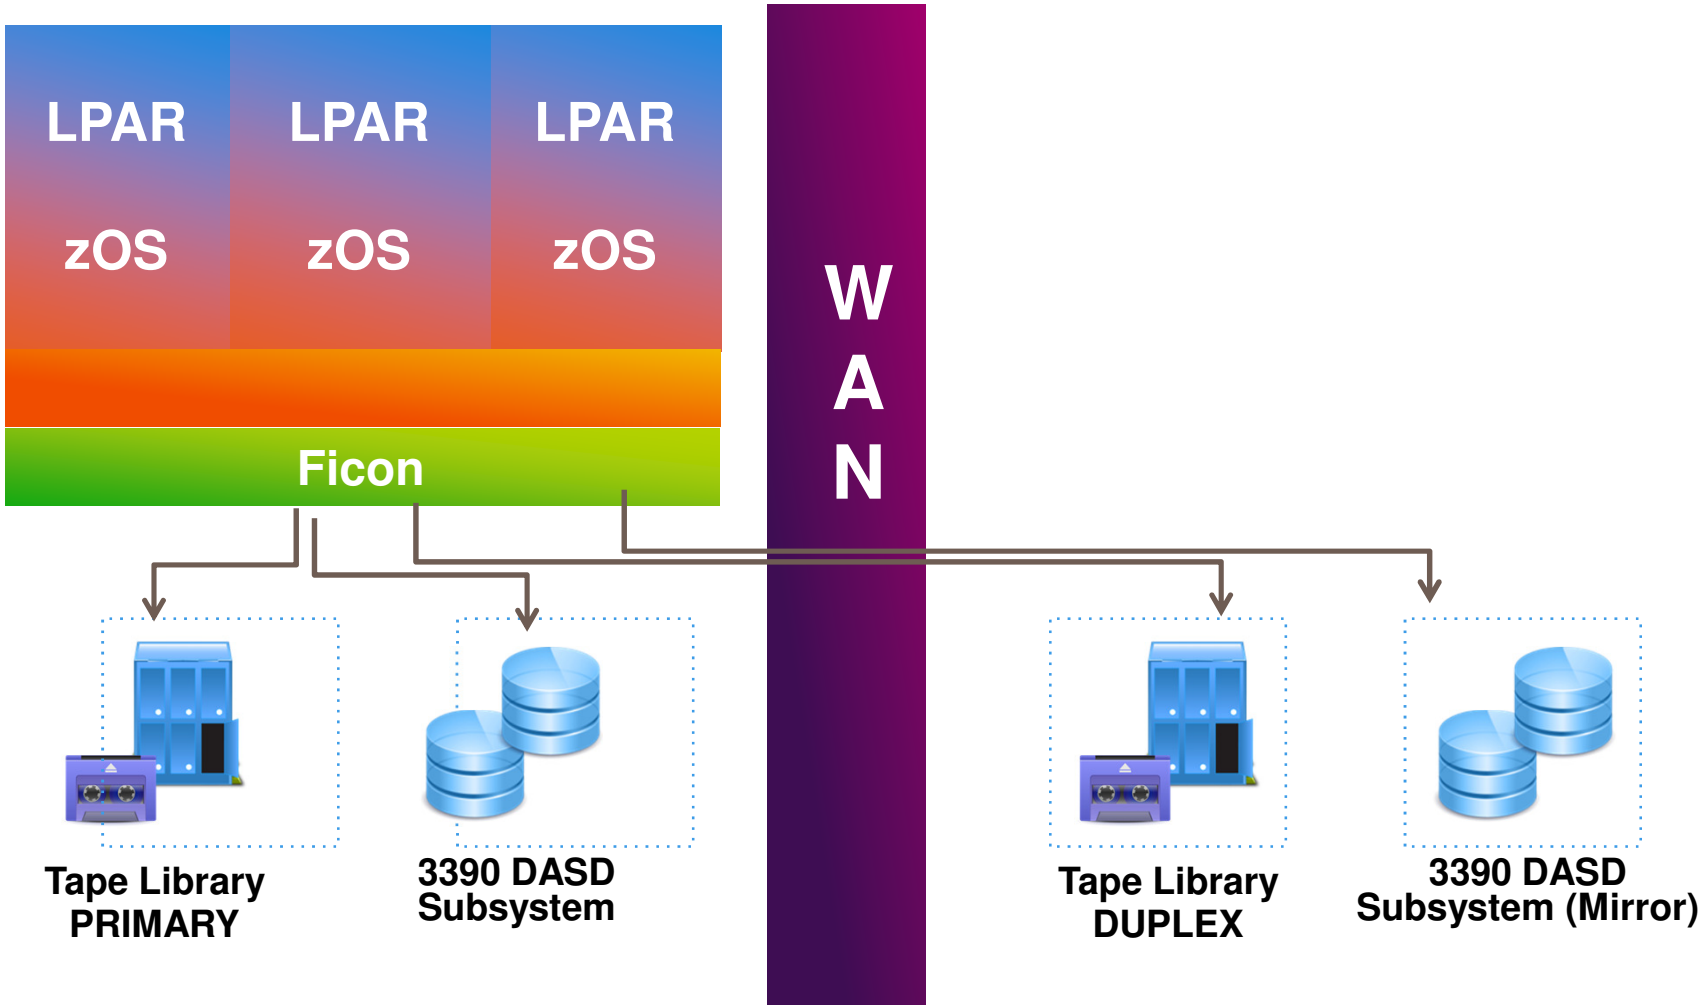

Traditional MF with replication to Hot-site

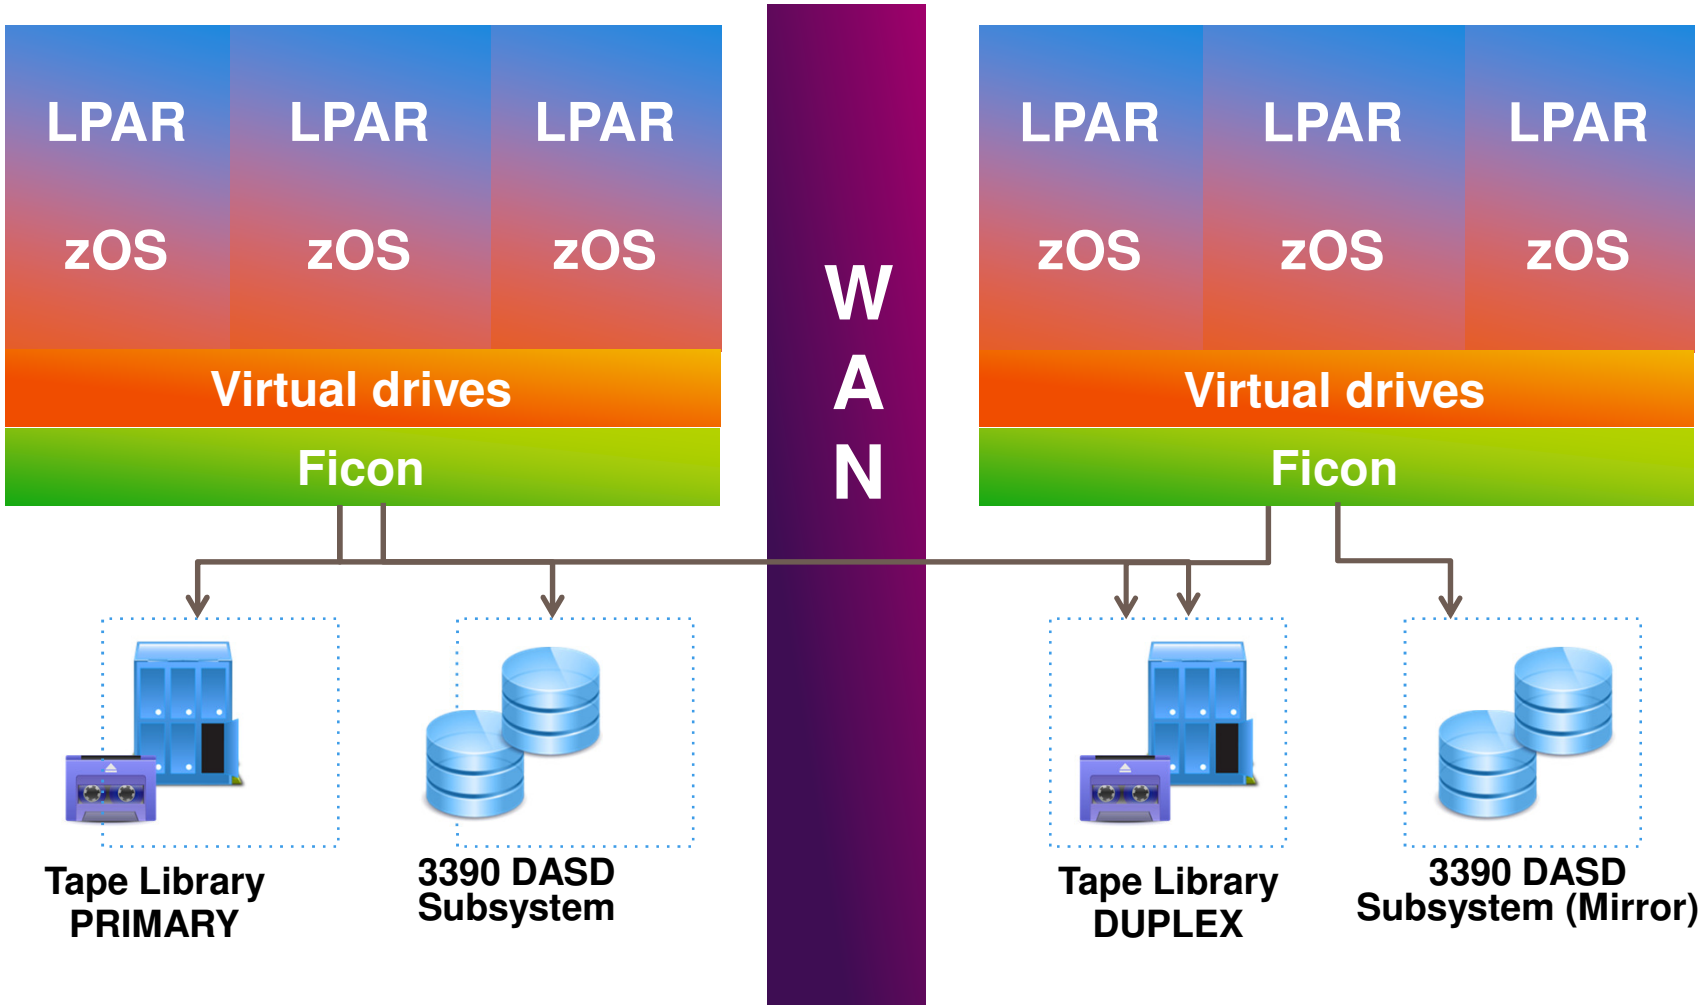

CA Cloud Storage for System z Architecture

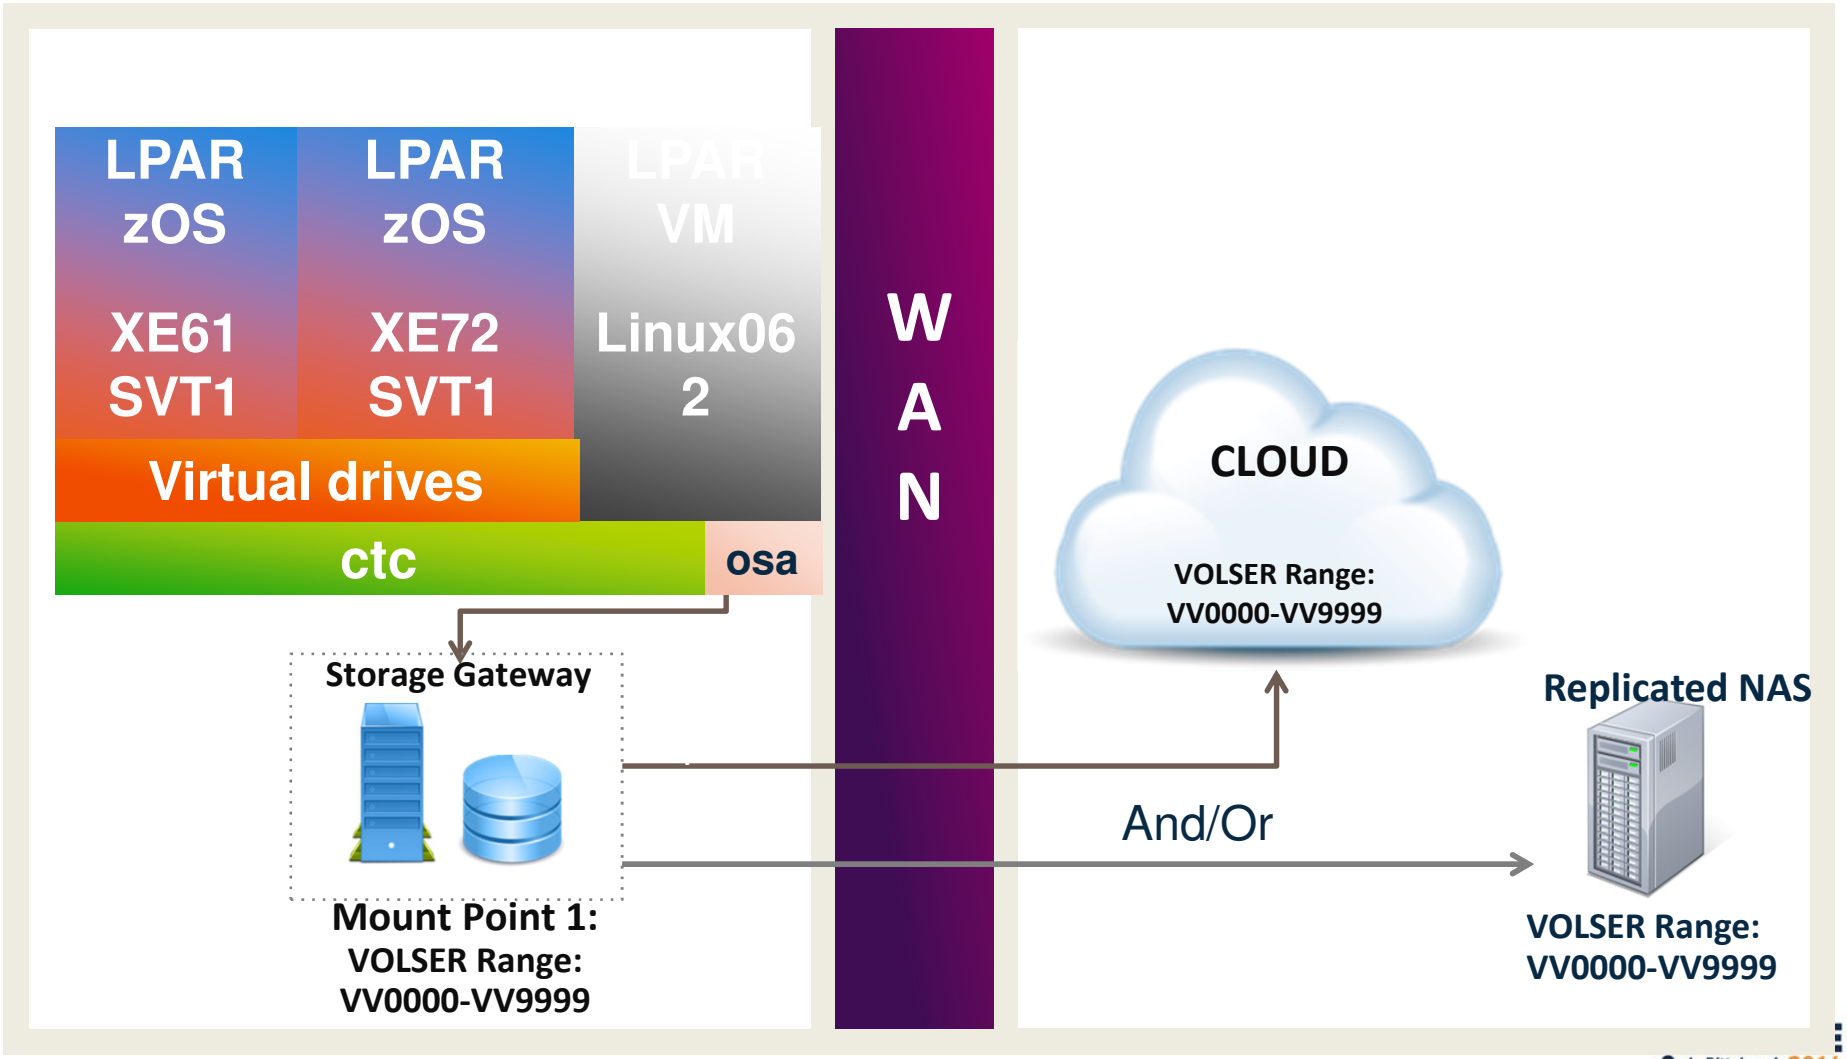

System Requirements - Software

- IBM z/OS - All generally supported releases
- CA Vtape VTS r12.6, with current maintenance
- IBM Linux on System z platform (zLinux)
  - SUSE Linux Enterprise Server 11 SP1 (s390x), kernel level 2.6.32.12-0.7
  - Running under z/VM recommended but not required

System Requirements – Hardware

- A pair of FICON CTC device addresses between the CA Vtape VTS subsystem (SVTS) running on z/OS and the zLinux server
  - Note 4 CTC addresses spread across two CHPIDs is better for throughput and availability.
- A 10-Gbps Open Systems Adapter (OSA) card connected to the Virtual Machine (VM) where the zLinux guest runs.
  - Note that a 1-Gbps OSA card is acceptable but throughput will not be as high as with the recommended 10-Gbps OSA card.
- CA Cloud Storage for System z can be configured to store data on any combination of on-premise NFS devices rather than on mainframe DASD

System Requirements – zVM & Linux on System z

- Minimum of 1 GB of free disk space (for syslog-ng, dumps, tar files, diagnostics, etc.)
- Minimum of 1 GB of memory
- Root authority to perform Network File System (NFS) mounts maintenance and installations

How it Works

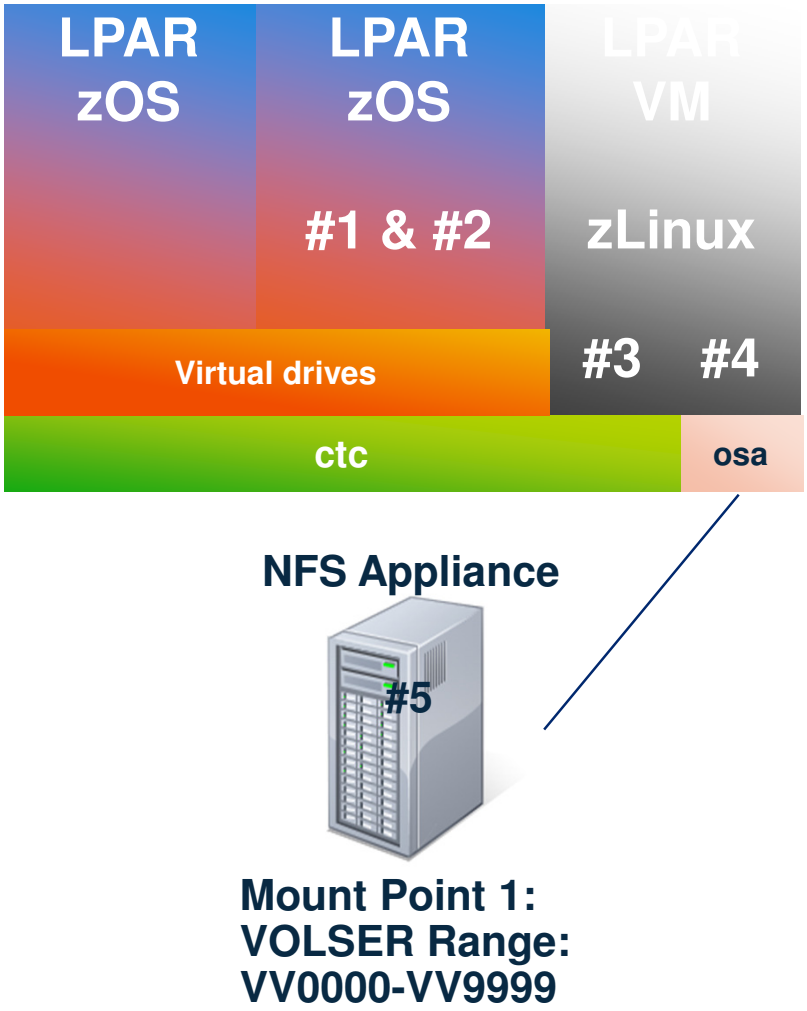

1. Mount is nearly unchanged except TCPIP is used to connect between z/OS and Linux. The Linux listener accepts the connection and forks a new process to do the work on behalf of the virtual drive.
2. The IO Engine is nearly unchanged, SSCH/IO are intercepted and the data is moved to 4MB Buffer. Only change is to use CTC pipe in place of media manager.
3. Linux receives the data over the CTC pipe.
4. Linux writes to the NFS mounted filesystem.
5. The NFS appliance is responsible for compression and encryption.

Note: This is a 100% tapeless implementation with no Backstore capability.

Why do I Need Linux?

- Objective is for the solution to use as little general processor as possible
  - Utilize zIIP when available
  - Utilize IFL when available
  - If specialty engines are not available then general processors are used
- Why not go directly from z/OS to NFS?
  - The z/OS NFS Client adds 20x CPU which is not zIIP eligible
- Ok, then why not use x86 Linux instead of Linux on z?
  - This is under consideration. We know there will be a general CPU increase but it should not be as costly as with z/OS NFS Client

Most Common Configuration

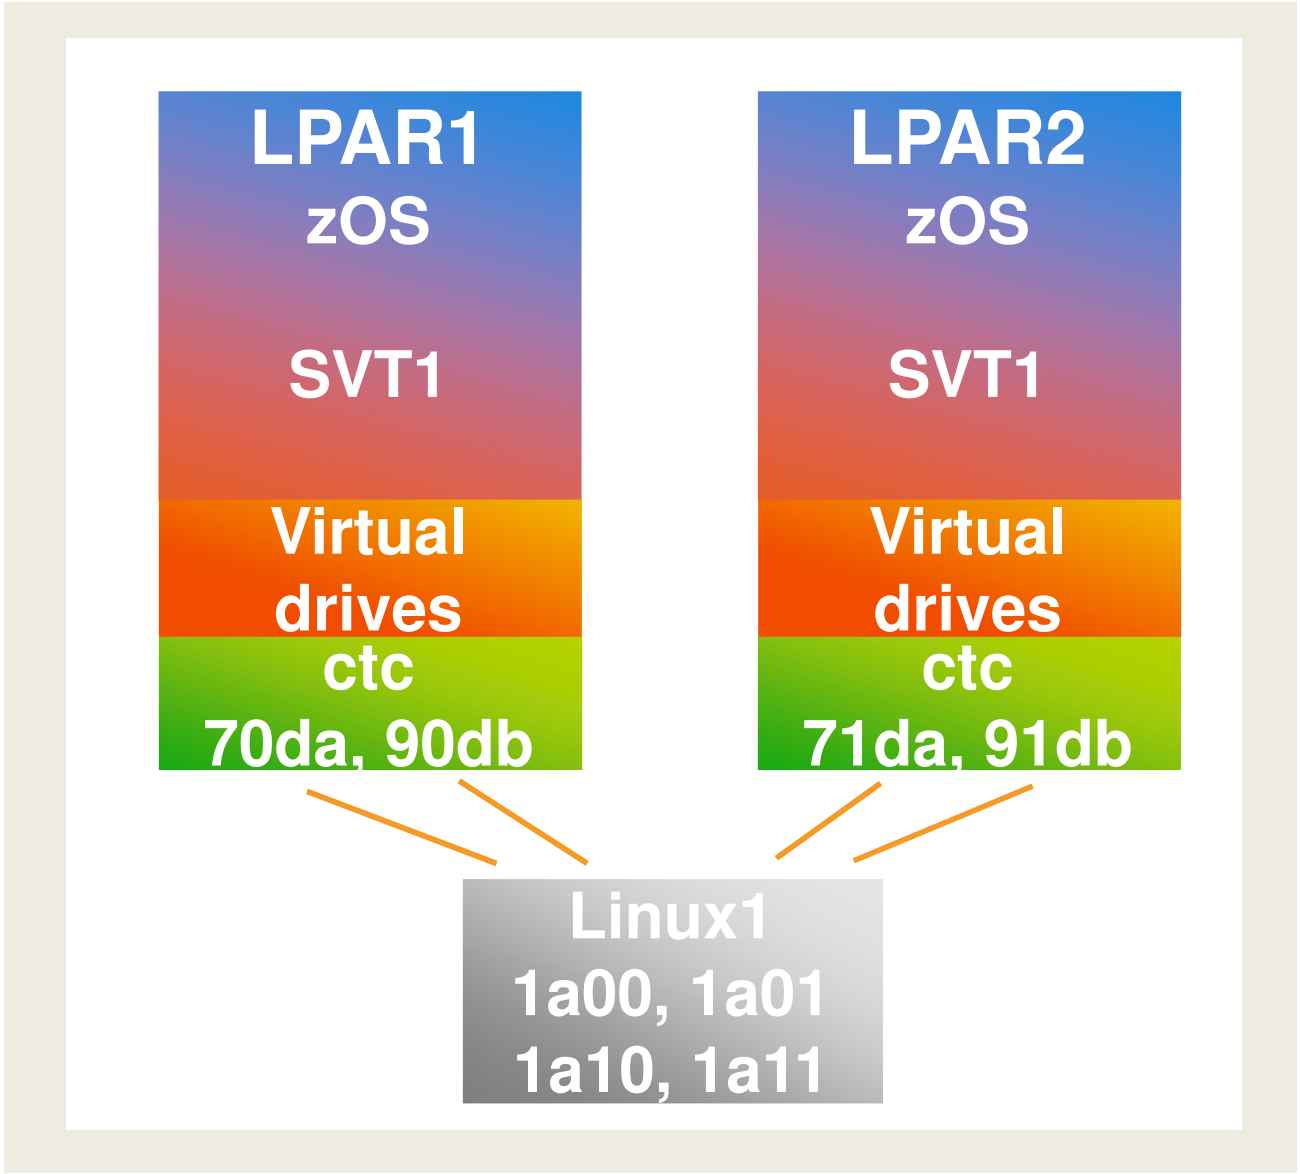

Highly Redundant Environment

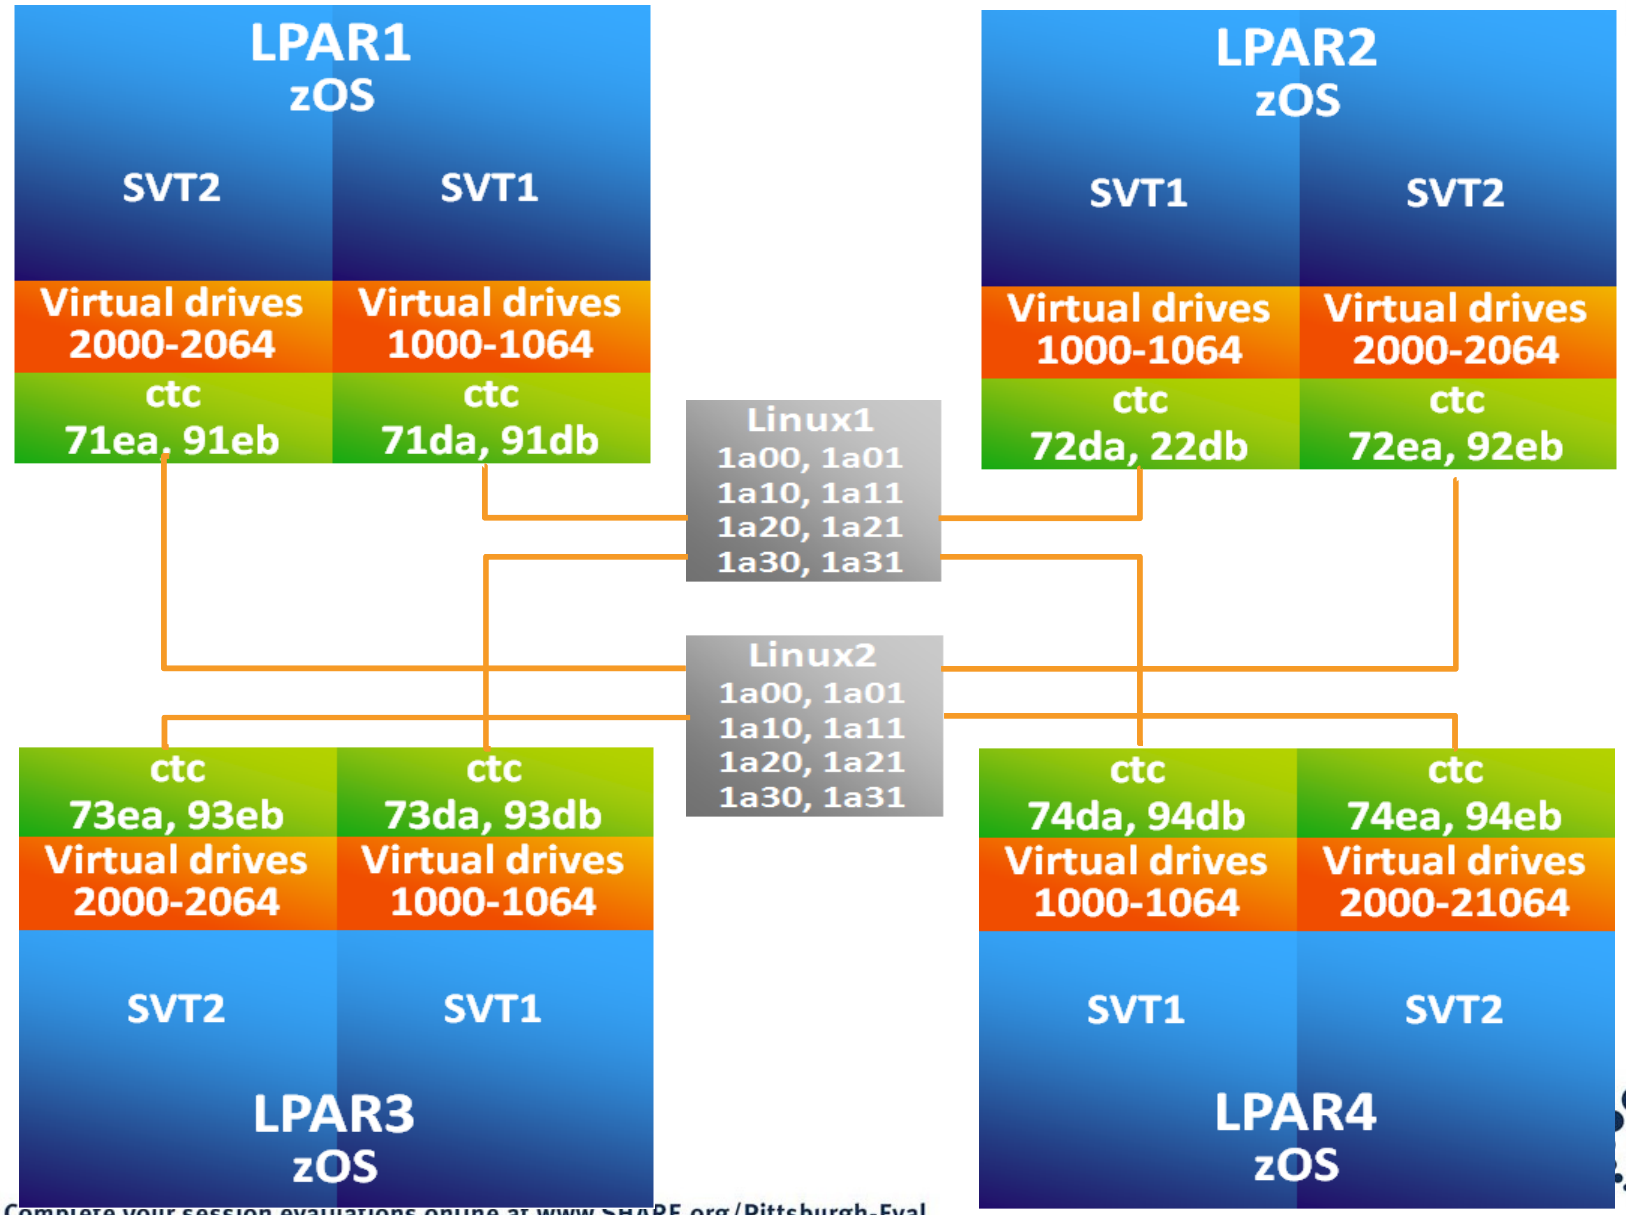

Multiple Mount Point Configuration

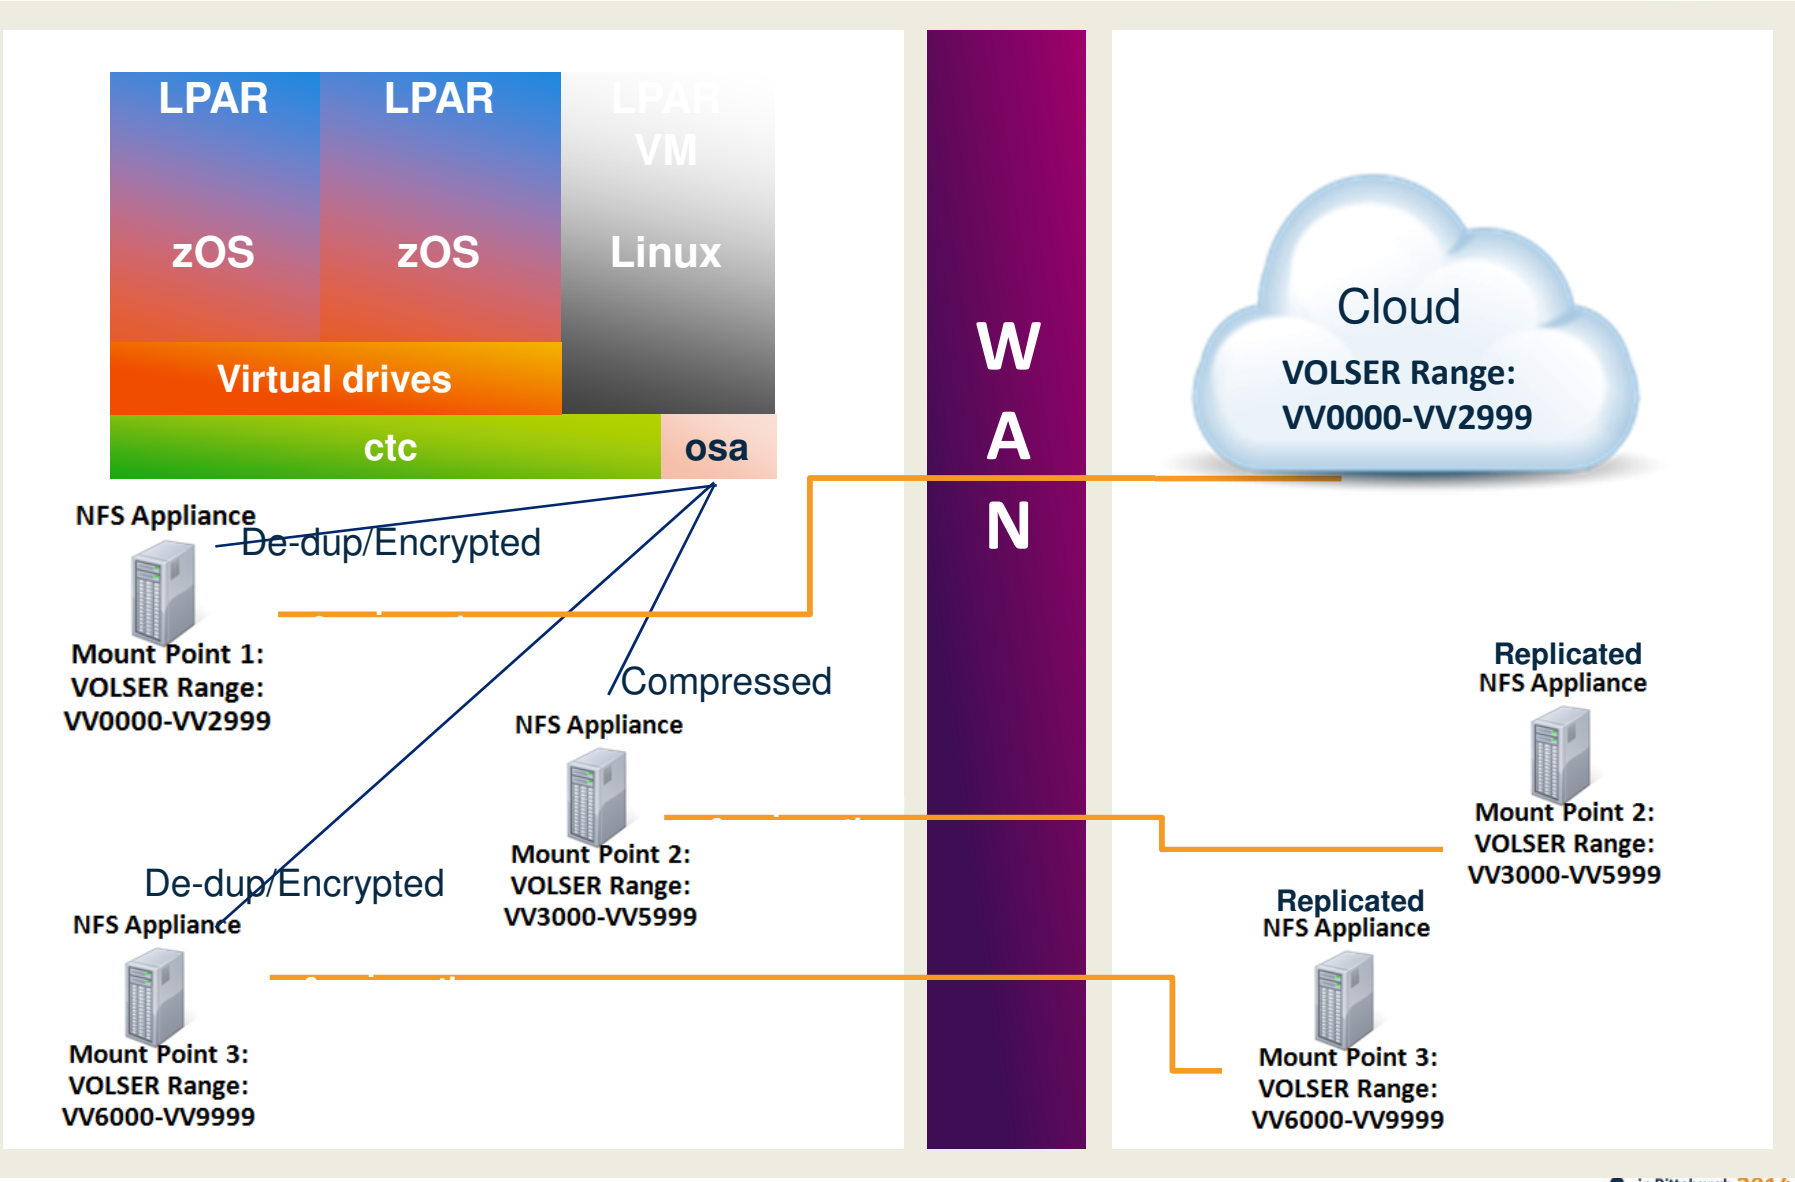

Benchmark Environment

- zEC12 2817-615 utilizing 2 z/OS LPARs and 1 z/VM with 2 Linux guests
- Each z/OS LPAR used 4 CTC addresses on 2 shared CHPIDs
- Both Linux guests shared a 10 Gbe OS4 Express4s configured to use Jumbo frames on a VSWITCH LAN
  - *TCPIP Segmentation Offload was not configured because the OSA must be configured dedicated*
- Both Linux guests wrote to the same 2030 Riverbed Whitewater

Benchmark Configuration

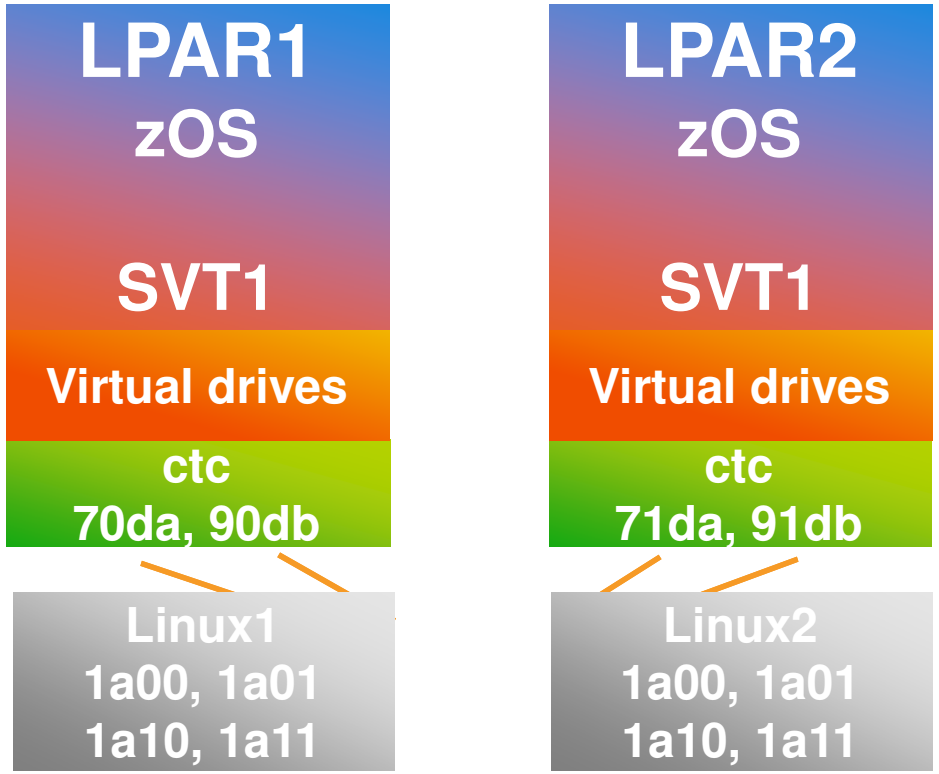

Benchmark Test

- Each z/OS LPAR ran 5 concurrent jobs running IEBDG to write 8GB virtual volumes
  - The 10 jobs ran continuously by resubmitting themselves over a 15+ minute timeframe to span a SMF interval
- CS4z z/OS started task information was collected in a WLM reporting class
- Linux on z IFL resources measured at the z/VM LPAR level
- Throughput reported by the OSA CHPID report
- Collected SMF 7n records and ran RMF reports
  - SYSRPTS(WLMGL(RCLASS(RVTSTC)))
  - REPORTS(CHAN)
  - REPORTS(CPU)

Projected Results

System	CPU Model	General CPU Engine Speed	Specialty Engine Speed	Peak Hourly Data Transfer (GB)	Projected %zIIP processor @ Peak	Projected %GP processor @ Peak	Projected %IFL processor @ Peak (Jumbo=Y, TSO=N)
Baseline	2827-615	947	1514	1830	14.17	5.18	148.5
zEC12	2827-7xx	1514	1514	1830	14.2	3.2	148.5
z196	2817-6xx	768	1202	1830	17.9	6.3	187
z10	2097-6xx	616	889	1830	24.2	7.9	252.8
z9	2094-6xx	468	581	1830	37	10.4	386.8

- The solution is capable of writing 500+ MB/sec which is 1.8 TB/hr
  - This should scale by replicating resources
- Minimizes general processor resource consumption by offloading to specialty engines
- IFL is expected to come down when exploiting TCPIP Segmentation Offload

Disaster Recovery Without Replication

PTAM – Pickup Truck Access Method

Recovery Time Objective (RTO)

- When your RTO goal is measured in hours, you can decrease cost and conserve network bandwidth by replicating to/recovering from the cloud
  - At your DR site you import your cloud settings from your cloud provider into a new Storage Gateway appliance and prioritize the volumes you recover
  - This is where information in the form of a report available at the DR location from the Tape Management System is critical. To know WHICH volumes contain which files is critical to the DR plan. Just like it is now with “PTAM vaulting”.

Recovery Time Objective (RTO)

- When your RTO is measured in minutes then you can have a 2nd Storage Appliance in “warm” mode receiving replication updates the same as when the cloud receives them
  - At your DR site your Storage Appliance already holds a mirror of your primary appliance and recovery can begin immediately
  - You can “pin” critical data in the appliance so that it is available at the recovery site and other data is off-loaded to the cloud; thus your Storage Appliance does not need to hold everything

Cold DR Mode with RTO Measured in Hours

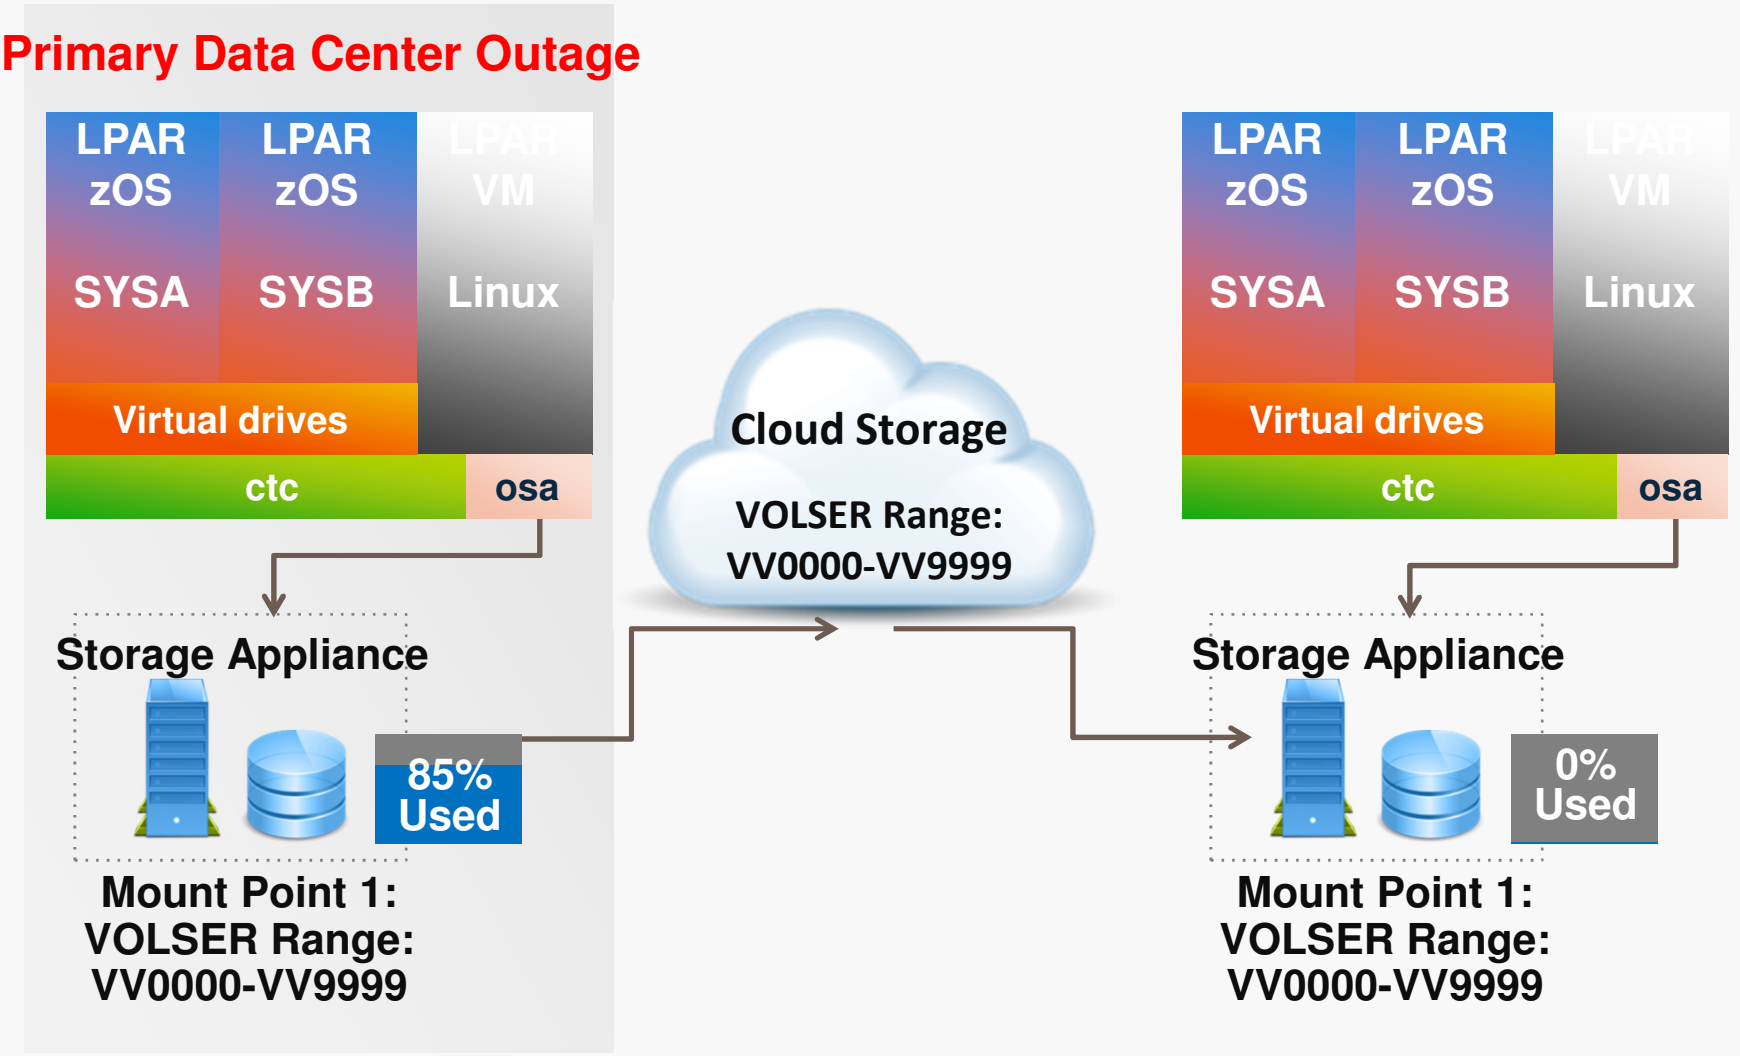

Warm DR Mode with Shortest RTO

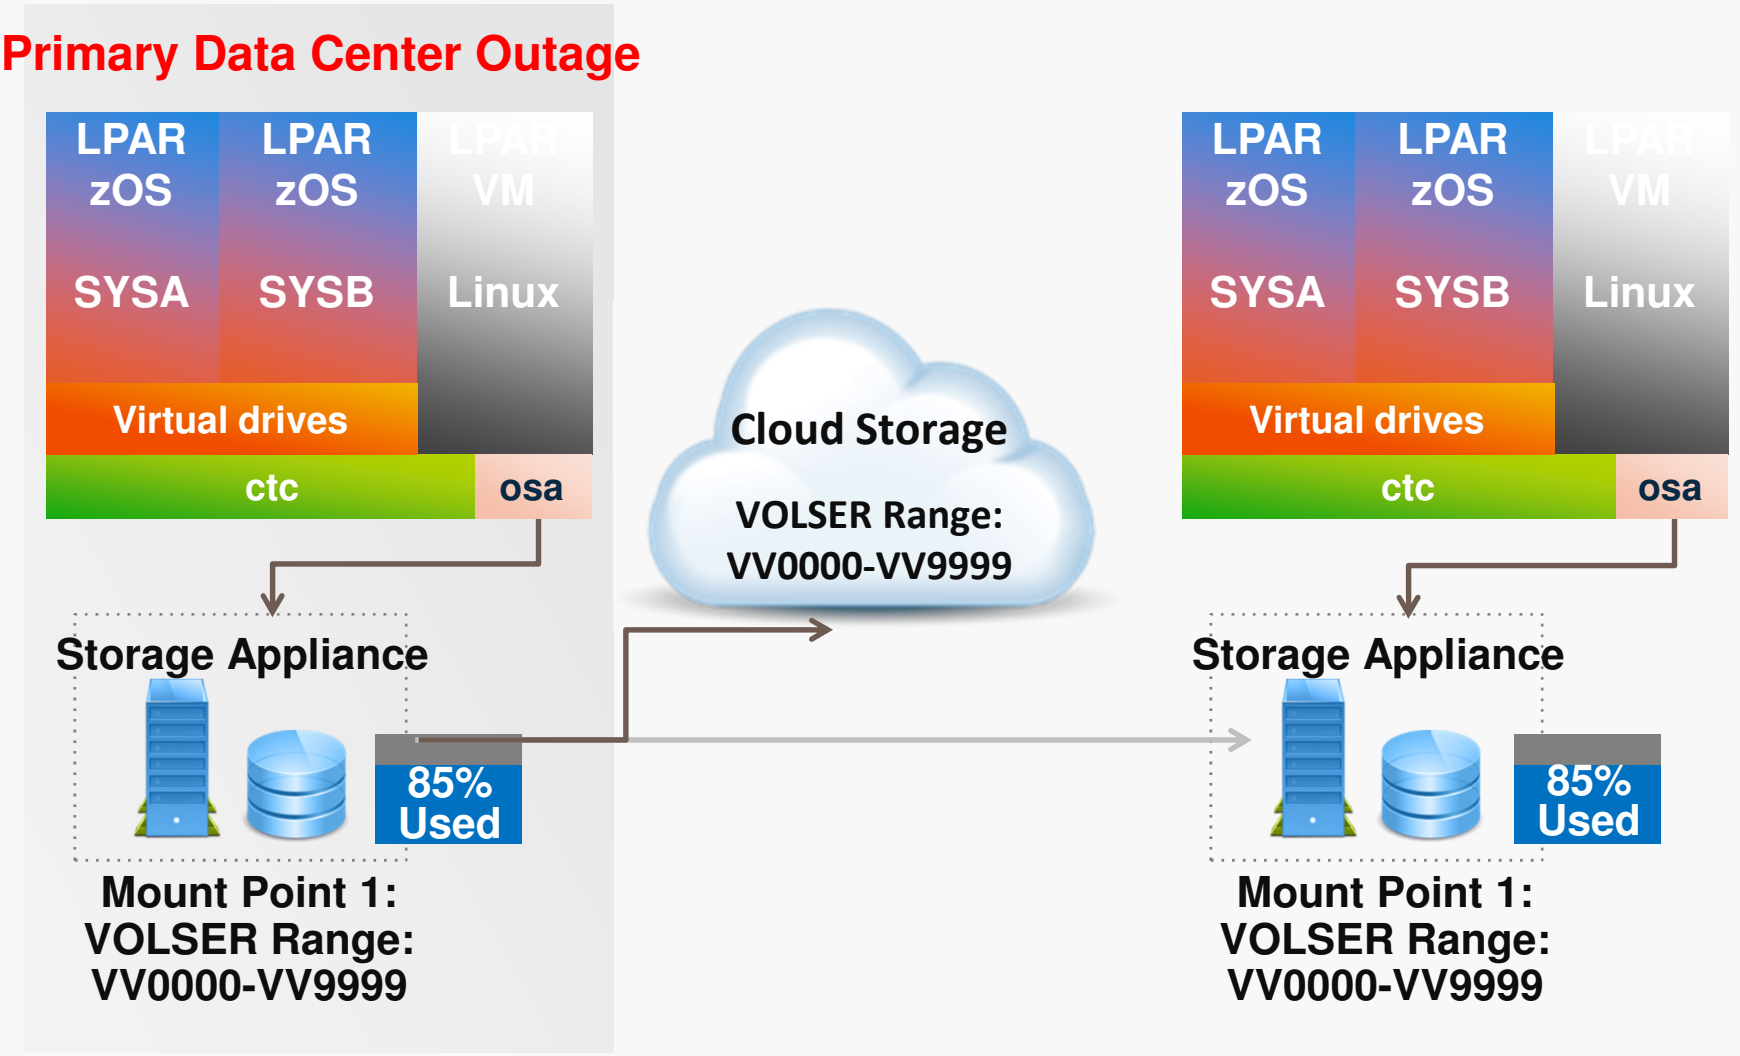

Benefits of CA Cloud Storage for System z

CAPACITY ON DEMAND

Reduce the risk of downtime & impacts to LOB.

Hourly cost of downtime for a typical company can range from $10,000 to $1M

HELPS OPTIMIZE BUS. SERVICES

Recent survey: Respondents said IT downtime can damage staff morale (35%)and can harm customer loyalty (44%)

HELPS MINIMIZE BUS. OUTAGES

Recent survey: delay of 7.5 hours per year during which time data is still being recovered and business operations aren't fully operational

DATA IS SECURE, COMPLIANT

34% of companies would be unable to meet compliance and regulatory commitments due to downtime

GREATER BUSINESS PRODUCTIVITY

When business critical systems are interrupted, companies estimate that the ability to generate revenue is reduced by 29%

CA Cloud Storage Solution Overview

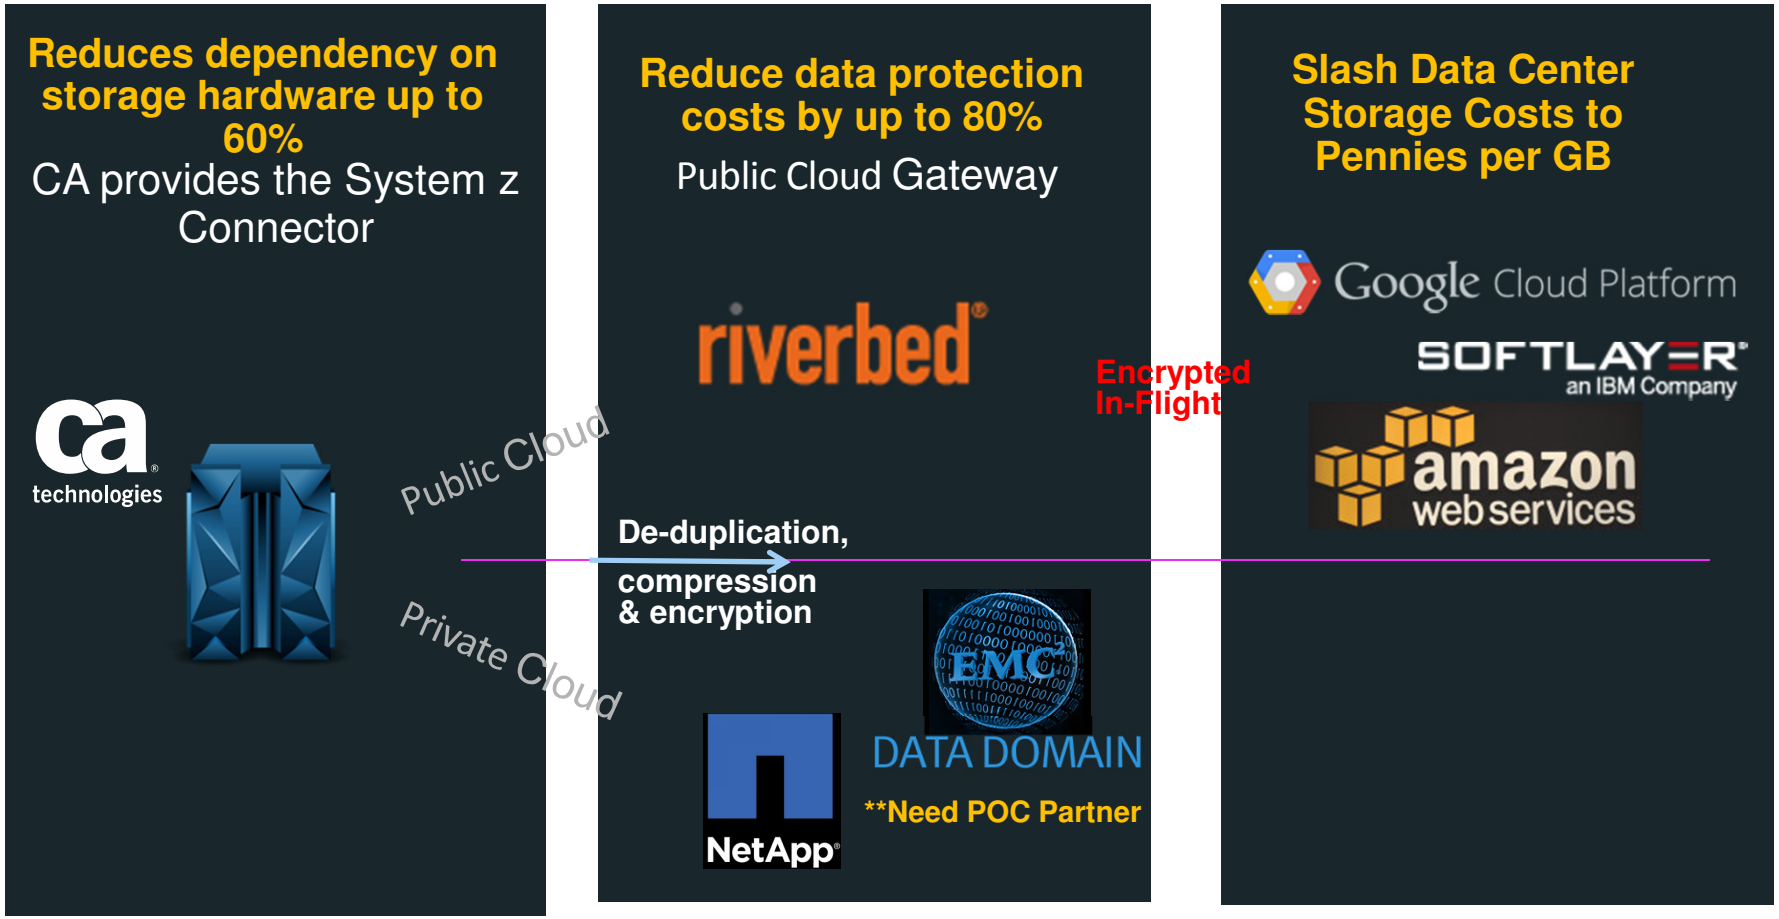

Summary

Develop a Tapeless Cloud enabled Enterprise Storage Virtual Tape Solution that reduces cost and provides a flexible disaster backup and recovery solution

QR Code