


# WHEN TRUE COLLABORATION MEANS GOING GLOBAL



## Matthias Trocme

Sr Solutions Architect  
Dell Technologies  
Matthias.trocme@dell.com

## Mickael Parnet

Technical Sales Representative – Compute Specialist  
Dell Technologies  
Mickael.parnet@dell.com

## Allan Guardini

Sr Presale Specialist – Unstructured Data Division  
Dell Technologies  
Allan\_guardini@Dell.com





The Dell Technologies Proven Professional Certification program validates a wide range of skills and competencies across multiple technologies and products.

From Associate, entry-level courses to Expert-level, experience-based exams, all professionals in or looking to begin a career in IT benefit from industry-leading training and certification paths from one of the world's most trusted technology partners.

Proven Professional certifications include:

- Cloud
- Converged/Hyperconverged Infrastructure
- Data Protection
- Data Science
- Networking
- Security
- Servers
- Storage
- Enterprise Architect

Courses are offered to meet different learning styles and schedules, including self-paced On Demand, remote-based Virtual Instructor-Led and in-person Classrooms.

Whether you are an experienced IT professional or just getting started, Dell Technologies Proven Professional certifications are designed to clearly signal proficiency to colleagues and employers.

# Table of Contents

Introduction .....	4
Collaborative Editing Challenges .....	5
A Proper Information Technology (IT) Infrastructure and Systems.....	5
An Answer to Geographically Dispersed Teams .....	5
Editor Requirements and Needs.....	6
An Editing Workplace.....	6
Standardize the Applications and File Format .....	7
Operations Requirements and Needs.....	7
Security and Operations Requirements and Needs.....	8
Security Matters.....	9
Design and Architecture for Collaborative Editorial Workflows.....	10
Object Storage Global Repository.....	10
Why ECS vs. Native Collaborative Storage used in Media & Entertainment.....	11
Resiliency and Reliability Matters.....	12
Edge Location.....	13
Hardware Innovation to Support Performance .....	13
A Virtual Editorial Workplace.....	14
Graphical Processing Unit Card Nvidia Ampere Architecture: A40 .....	15
Size the Number of vGPU.....	17
Advantage of GPU Sharing.....	17
Size the server or Hyperconverged Infrastructure (HCI) platform .....	18
Choosing a Virtualization Solution .....	19
Bringing Pieces Together .....	19
The Data Governance .....	20
Conclusion.....	20
Bibliography .....	21

## Introduction

Modernizing creativity is always a challenge in architecture and design: how to manage the growing number of media files that need to be accessed in real-time worldwide in a performant, sustainable and secure way?

This article's purpose is to describe a system design that makes collaborative editing possible!

This system will give multiple users worldwide simultaneous access to the same media, and allow geographically dispersed teams to share content, ideas and make working decisions together.

Let us discover it from an end-user experience from the virtual Desktop Render farm, 3D (3 Dimensional), After Effect or Video editing up to a global object storage repository.

Based on our recent work on the transformation of the collaborative platform within a major Media & Entertainment group, we know that this system is built on the combination of multiple innovative Technologies products and solutions, from end-users to edge to core Data Center.

This is a collaborative design, meant to support home office workers spread around the world. Without any compromise on performance in editing content that generates a lot of data.

This article shows how, based on the structural needs, it is possible to build an infrastructure block by block much more agile than the actual monolithic platforms. The system presented in this paper could be relevant to any verticals – not only Media & Entertainment, for any needs in editing large files in a collaborative way with a central repository meeting performances challenge.

How to build such an architecture? Design it? And more importantly what are the sizing guidelines? This document covers a disruptive way to use innovative components combined with 3rd party products to cover all required features needed for such a solution

## Collaborative Editing Challenges

For large Communication Groups specialized in Media and Entertainment the challenges are multiple. But amongst all those challenges, one big challenge is: **collaborative editing**.

- *“The two factors that make collaborative editing possible are shared storage, which gives multiple users simultaneous access to the same media, and fast networks, which allow geographically dispersed teams to share ideas and make decisions together.” (Ref.1)*

This quote from an article published in 2018 was already pointing two crucial factors to making collaborative editing possible:

### A Proper Information Technology (IT) Infrastructure and Systems

To build a technological answer to collaborative editing using a unique solution is not enough. It requires a proper Information Technology (IT) infrastructure, systems, data governance and security policies to ensure the success of any of this type of project.

Over the past years, multiple innovations dramatically changed our world. The communication world is moving fast, and the internal IT systems of media and entertainment often struggle to embrace this digital transformation.

“More agile, more end-user centric, without restriction and with the right performances” are key words to these kinds of IT Infrastructures. Digital transformation and hardware innovations also bring answers to this challenge. The main innovative domains that we will cover are **compute, shared storage, and the workplace of editors**.

### An Answer to Geographically Dispersed Teams

Creating a media, marketing or communication campaign involves many teams and stakeholders working on the same project. Due to the pandemic, the behavior changes so more employees and Freelancers are now working from home with communication group business lines and to do so, **they need to access tools and base contents to perform their day-to-day tasks**.

In this article we will focus on three distinct types of population:

- **The editors** that are working on creating and editing content
- **The Chief Technology Officer (CTO)** and the IT Teams that deliver, manage, and operate the IT systems and infrastructures for editor’s business lines
- **The Chief Security Officer (CSO)** responsible for global and corporate security.

## Editor Requirements and Needs

# Editors concerns

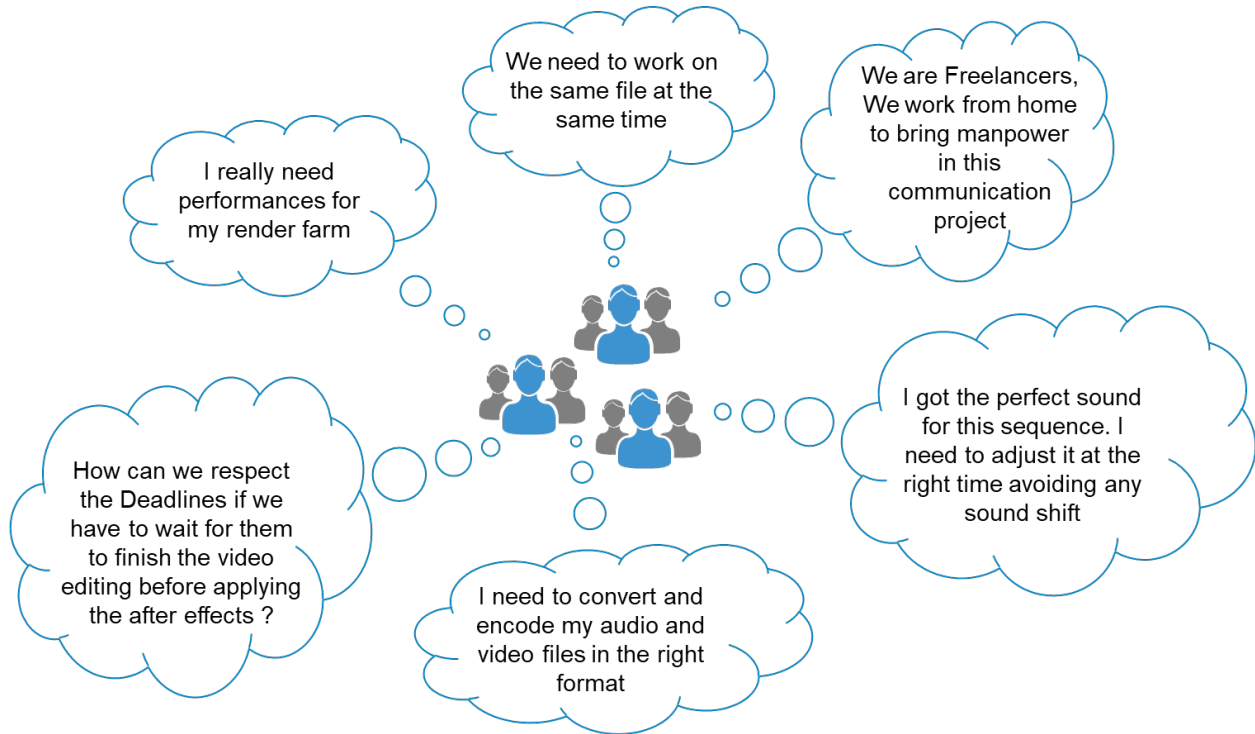


Figure 1 – Editors Concern

## An Editing Workplace

The editing workplace has been transformed over the years, from local edition on a workstation to shared oriented solutions that make collaborative work possible. Many teams are working together on single projects and need to respect the same guidelines and format without having the need to copy the data on each of their workstations.

The first key component that, from an end-user perspective, should be transformed, is its own work environment. In our architecture we analyzed the end-user current workstation in terms of configuration, applications, and usages to build some templates that would be most representative as well as a most common virtual environment.

Building those templates is a critical step. The first thing to do is to identify the use cases such as:





Figure 2 Workplace Templates

An example of today's configuration is a workstation running Windows 10 with powerful Graphical Processing Unit (GPU), remote mouse and keyboard, and all suites of applications dedicated to this world of content editing, such as after effect tools.

One of the main reasons to regroup and standardize those templates is to avoid any operating systems conflicts and to provide a collaborative workplace that will answer the needs of every editor.

### Standardize the Applications and File Format

The process of working on a similar audio or video file brings complexity in terms of the various formats that are nowadays used by editors. With the constant need for quality improvement and the heterogeneity of data sources, the definition of **encoding standards** is one of the challenges of editors.

This standardization of format goes to a process of **transcoding** that converts the files format to the desired one. Splitting the files into multiple small raw files allows more agility in the manipulation of every frame. It also answers the complexity in the usage of those frames containing more information due to the improvement of quality and definition.

But how to handle the master source footage? How is it possible, in terms of latency and number of threads that will need to be running concurrently, to target a single master without causing integrity issues? An appropriate storage and working environment are the key answers to support editors in their creativity.

### Operations Requirements and Needs

2022 Dell Technologies Proven Professional Knowledge Sharing

# CTO & IT Teams concerns

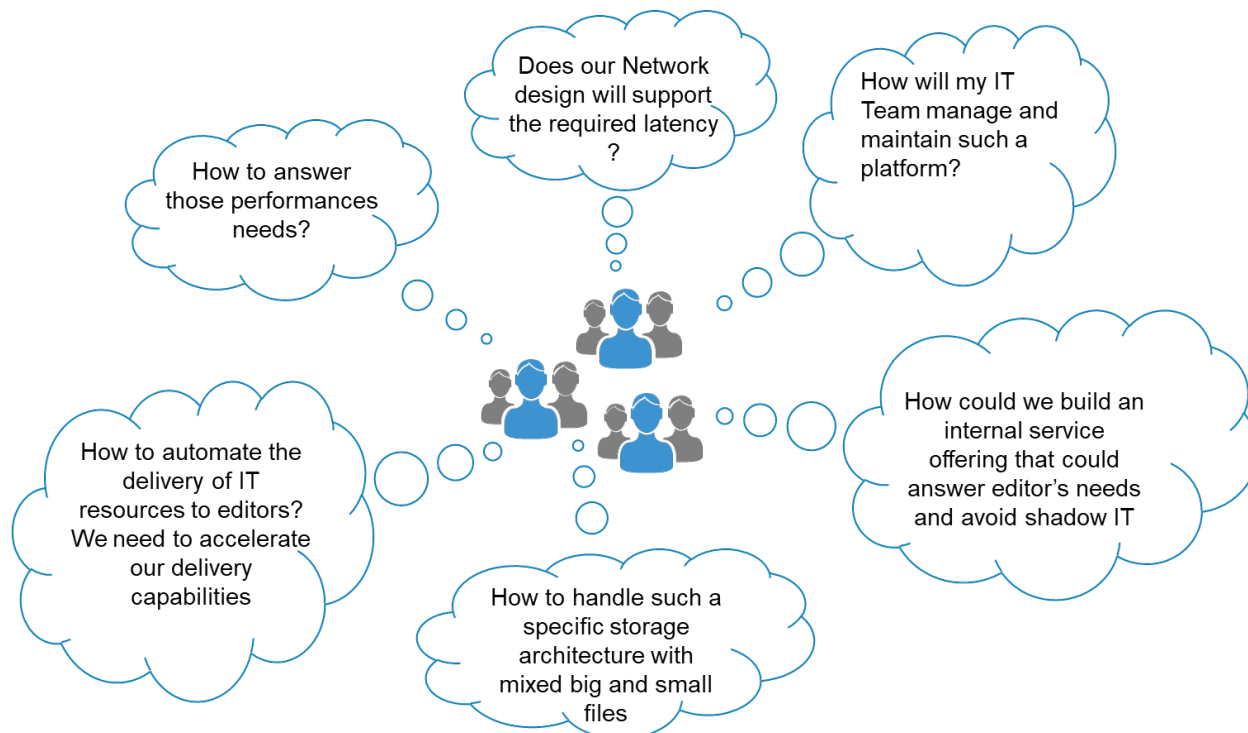


Figure 3 – CTO & IT concerns

## Operations are the Foundation

Content creation is a specific world where IT Teams need to accelerate the delivery of tools that will help content creators streamline their editorial workflows.

Building and maintaining an editorial platform on premise is a challenge. These architectures are different from traditional IT systems in every aspect. The IT Teams are delivering resources to Business Lines users and responsible of the day-to-day operations. A tight spot where any help and advice from infrastructure vendors and solutions editors need to be leveraged to ensure a high quality of service delivery.

Finally, the security of operations is a key topic: on the one hand there is operations, but on the other hand security. How to reconcile an elevated level of security with smooth operations and agility for supporting the business lines?

## Security and Operations Requirements and Needs

Learn more at [www.dell.com/certification](http://www.dell.com/certification)



# CSO & Security Teams concerns

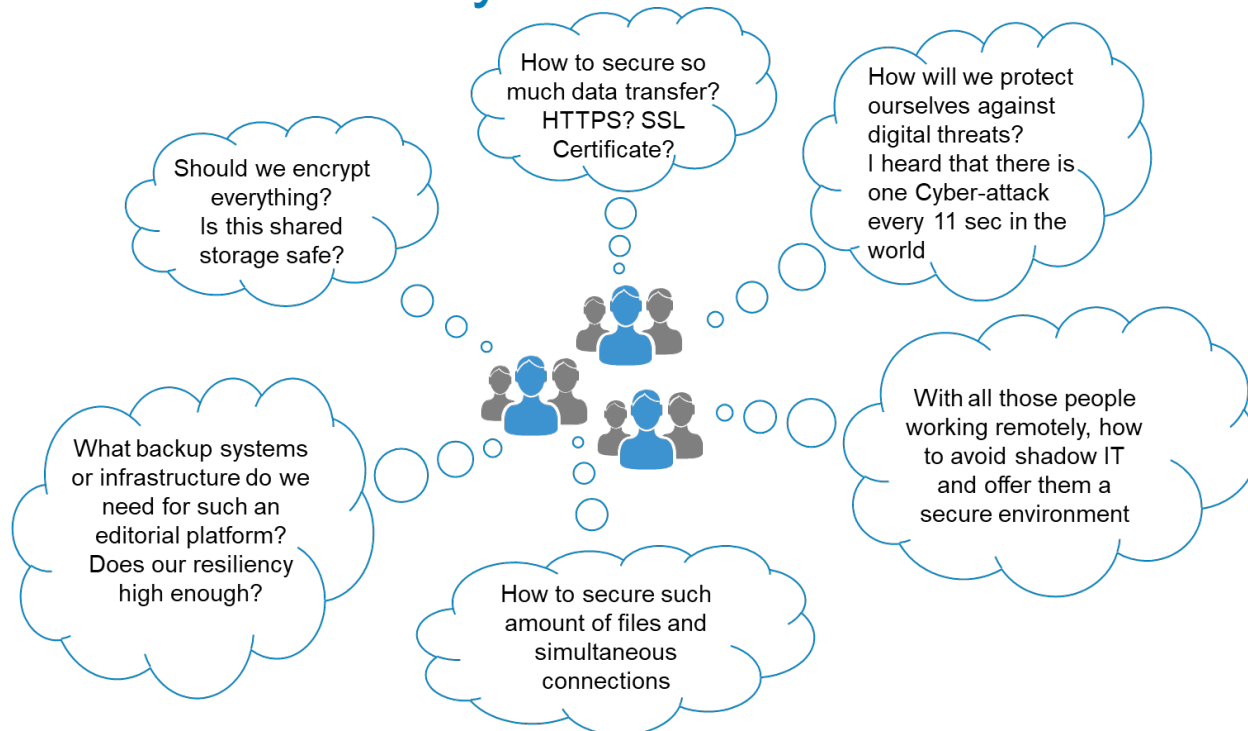


Figure 4 – CSO & Security concerns

## Security Matters

From the editor’s workplace to the core data center, managing the data in a secure way is critical for communication for Media and Entertainment companies. These data are after all what we can call their corporate memory. As the real “data fuel” of Intellectual property, it requires maintaining an elevated level of security at any time, for everyone to avoid data to be damaged, stolen or leaked.

To ensure a correct level of protection of this sensitive information that is often targeted by digital threats, especially within this new remote work world, the design of the storage and of the IT architecture end-to-end needs to be rethought.

Going to public cloud is not an option due to the sensitivity of data which contains critical corporate intellectual property.

Data transfers, multiple remote connections, data governance and end-user’s roles managing access are the challenges of the security teams and the concerns of the CSO. How to define the right level of defense for a geographically spread IT platform and so critical for the business?

## Design and Architecture for Collaborative Editorial Workflows

Now it's time to get an overview of the main components of the solution that will be covered in this document. We will start from the bottom, beginning with the infrastructure in the core data centers moving upwards to the end users' workstations.

The following architecture is based on 3 layers:

- Core data centers with object storage global repository
- Edge locations, that are the Business Lines offices supporting the data cache shared storage gateway
- The virtual editors workforce environment

Even if the modernization and the transformation of each of these layers can be handled separately, building them together is a guaranty of success in a transformation of the collaborative platform.

Analyzing the needs and having an "out of the box" approach leverages the use of latest innovations in every technology and product domain to create a platform built with global governance. Like the micro-services methodology, analyzing every element as an individual one allows us to present a very modular architecture that can scale and support the agility needed by business lines. It also insures a platform with consistent operations and management without any compromise in resiliency and security.

### Object Storage Global Repository

Dell Technologies is offering DELL EMC Elastic Cloud Storage Solution (ECS), an enterprise-ready platform that enables organizations to simplify object storage management, visualize information in new ways, and empower the business to do more with data, all at a lower total-cost-of-ownership. (Ref.2)

This global storage repository perfectly matches our CTO & CSO needs, bringing a solution that can be a global repository on premise, at the end of the value chain of editorial editing.

Indeed, as the data generated are sensitive intellectual properties of the Media & Entertainment companies, keeping them locally at an affordable price is of great value. Re-using those data for future communication campaigns or projects is one of editor's requirements. Combining them together it is representing what we can call the "**Corporate Memory**"

As the foundation of all repositories, the use cases that this technology is addressing are various:

Learn more at [www.dell.com/certification](http://www.dell.com/certification)

# Utilizing ECS object storage in your business

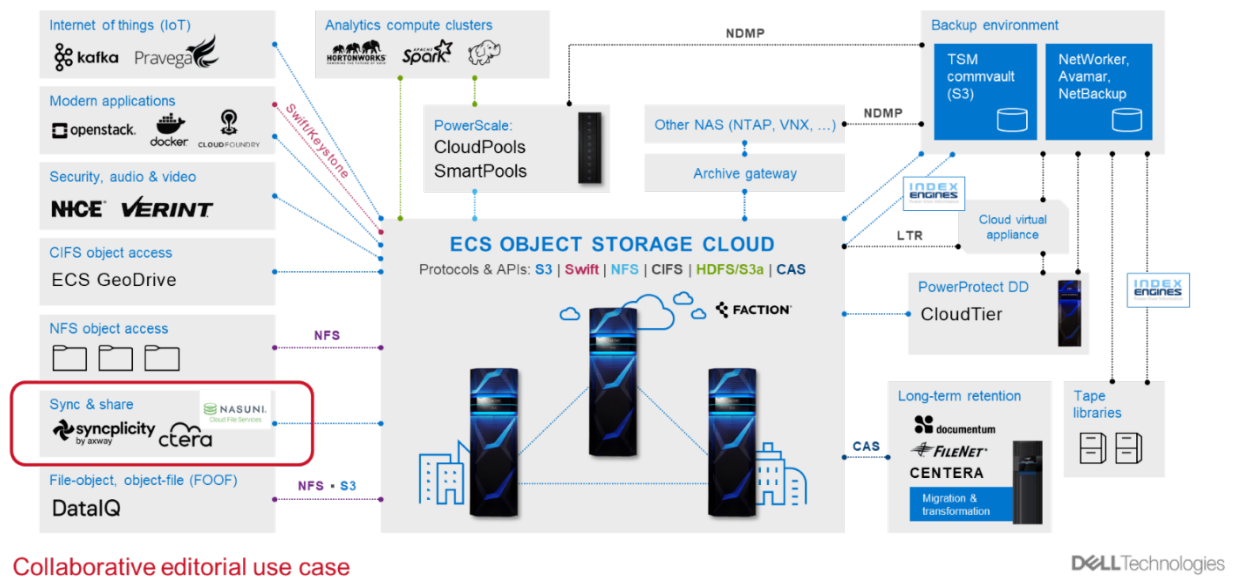


Figure 5 – ECS Use cCases

## Why ECS vs. Native Collaborative Storage used in Media & Entertainment

Unstructured data such as files and objects represent most of the generated data growth over the past years. Choosing an object storage solution is a natural choice in the Media & Entertainment world where specialized solutions exist to address collaborative editorial use cases.

Why is ECS a desirable choice for such a platform? The answer to that question can be summarized in the fact that ECS has been purpose-built to store unstructured data. Offering great scalability and control of operations to operational teams, it offers a global distributed storage infrastructure under a single global namespace with anywhere access to content. Oriented to address operations in a straightforward way, the ECS architecture can evolve over time in a stable way without disruptions as it grows.

Other technologies may have they own value proposition, offering number of “nodes” optimization required to offer cluster protection and geographical distributed access, but it always comes with a cost in terms of resiliency, as well as reliability of the data hosted on such systems.

Making the choice between more affordable or more secure? That is the real question for IT directions to answer to. From our point of view, the security, and the reliability of the data, which

are again the vital sensitive fuel of a communication or media group, is clearly not a subject to take into consideration.

## Resiliency and Reliability Matters

Sensitive data requires specific design and architecture, a global repository built on ECS offers interesting native functionalities as replication and resiliency by design.


Core of the Corporate memory, the security and the data protection issues are handled here by duplicating the architecture on two different sites to prevent any geographical disaster or events. Working in a mirror offers the best protection.

The internal design of each of those two ECS clusters uses an erasure coding schema adapted to this high data protection level.

## Architecture value proposition : Single Site Failure Tolerance – 12+4 EC Scheme

- The Mirror 2 sites architecture (with 5 Nodes)
  - Data can be written on each side even in the case **of one site down**
  - One Node failing has no impact on services – protection still active
  - Two Nodes failing : No impact on services but no more protection
  - Three Nodes failing : Writes disable, Read still possible, no more protection
  
- Erasure coding Protection Level
  - Erasure Coding sheme 12+4, Offering a Data Protection Efficiency of 2,67 with 2 sites in Replication Group

Erasure coding scheme	# Nodes in VDC	Number of chunk fragments per node	Erasure coded data protected against
Default (12 + 4)	5 or less nodes	4	<ul style="list-style-type: none"> <li>• Loss of up to four disks or</li> <li>• Loss of one node</li> </ul>



12 data fragments +  
4 coding fragments

Figure 6 – ECS Architecture value proposition

Here is an example of our architecture choice for the global storage repository using two sites with on both, an ECS cluster with five minimum nodes. The Erasure Coding (EC) on ECS offers an interesting choice in comparison to the standard level of Redundant Array of Independent Disks (RAID) as the disk technology combines multiple physical drives into one unit. It brings protection level on a hardware level to fragment the data across all disks to ensure a disk failure. Altogether this object storage solution offers a multi-site, geographical spread.

## Edge Location

The Business Lines offices are what we call edge locations. In these offices, data reside and are the extension of the IT systems for end-users. Has a “hub” of data, this is where, in an editorial workflow, we will find the solution of cache distributed storage. This local storage will host the data in a current project fueling the editor’s applications.

Spread around the world in the major cities where the communication leaders have their head office, the solutions on these sites need to be:

- Simple in terms of management
- Performant to deliver the best services to the users
- Collaboration oriented

### Overcomes latency to files simultaneous access

Different solutions exist for giving access to large and multiple files at the same time. In our architecture, this is where we begin to work with 3<sup>rd</sup> party hardware or software solutions. In the Media & Entertainment world, Business Lines must take the final decision.

For building our solution, solutions were benched directly by the end users. This article has no meaning in pointing out the Pros and Cons of each of these solutions. We will just list the last three solutions that where shortlisted before the final choice:

- Panzura (Ref.3)
- Ctera (Ref.4)

The final choice brought us to a design using as hardware PowerEdge Servers from DELL and the CTERA Software solution.

The combination of CTERA (Ref.4) software that offers a global file service supporting multiple access on the same files and a powerful hardware design to guarantee performance was the design retained to equip the multiple edge locations.

## Hardware Innovation to Support Performance

“As storage demands expand and processing loads grow, RAID (Redundant Array of Independent Disks) data protection has become a necessary staple for proper enterprise storage management. Dell Technologies’ PowerEdge™ RAID Controller (PERC) provides a RAID solution that is both powerful and easy-to-manage for enterprise storage data protection needs.” (Ref.3)

With the standardization and the innovation brought to the PERC11 systems we can now support hardware Nonvolatile Memory Express (NVMe) for NVMe drives.

This is a **game changing in the design** of our Edge locations design infrastructure: IT Teams will now be using servers capable to use hardware RAID on NVMe drives to protect their data and improve system performance on the “cache shared storage” solutions deployed to extend access their project data.

Here are examples of performance bench realized between RAID with 10 disks Solid-State Drives (SSD) versus 6 NMVe disks in RAID 6 NVMe (Ref.5 Figure 5, Figure 6, Figure 7)

Why is this so important? This is a technological innovation that is:

*Bringing performance to another level*

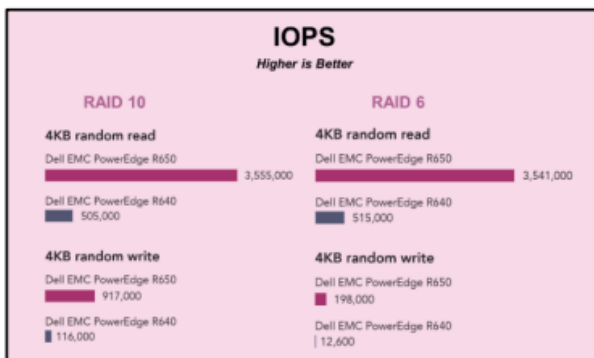


Figure 5: IOPS performance comparisons




Figure 6: Disk bandwidth performance comparisons

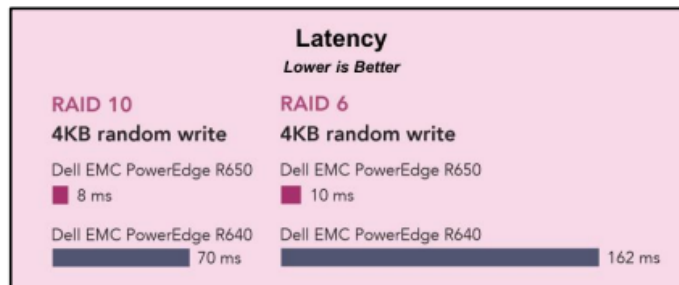


Figure 7: Latency performance comparisons

Figure 7 – Performance of RAID NVMe (Ref.5 Figure 5,6,7)

*Affordable*

Even if the NVMe drives are more expensive than SSD drives, it is required for more performances and less disks. This is an economic equation but the configurations that we designed for Edge locations for local shared storage are cheaper than the traditional RAID 10 SSD configuration that is today a common design in this type of architecture.

### A Virtual Editorial Workplace

Learn more at [www.dell.com/certification](http://www.dell.com/certification)



To modernize and transform the workplace of editors, we investigated the possibility of virtualizing the desktop.

Adopting a virtual desktop infrastructure (VDI) strategy was guided by the benefits of such systems. From an end user perspective, it answers the needs of editors and Freelancer mobility.

From an IT team perspective, it also offers better centralized management. The profiles of today's non virtual workstation of editors are not standard within large organizations. It requires a Helpdesk with dedicated skillsets able to handle the management and the issues of extremely specific highly performant configurations users' processes. As critical for the media business, those users are particularly important business users.

From a CTO & CSO perspective, it brings an opportunity to adopt a better compliant and high security configuration. It is also a choice in terms of investment. More flexible, consistent, and standardized, the provisioning of a new virtual workstation is quicker and the release of resources when a major project ends, much easier.

It is also the ability to accelerate media campaigns that often seduces the business lines when adopting virtual desktop infrastructure. The idea of having a template of a ready to use environment is also important for new users' experience.

Most of the editors' physical workstation configurations use GPU for running professional media and editorial applications. It requires stronger configurations in terms of calculation processor and memory than anyone else in the company.

This architecture also consolidates the CPU and GPU number and usage. Appreciated for its efficiency and positive impact in terms of consumption. This is contributing to optimizing and reducing the environmental impact.

The investment in the adoption of virtual desktop infrastructure is interesting for global companies with worldwide branch offices running 24/7. Indeed, running GPU instance in cloud public is quite expensive and having an on-premises infrastructure that is running all the time in a "follow the sun" approach is more affordable.

From a technical point of view, the performance considerations are, and must be, the first step in such an ambitious project. In the strict respect of every best practice, our experiences as manufacturer are to perform a proof of concept.

### Graphical Processing Unit Card Nvidia Ampere Architecture: A40

Render Farm, 3D, after effect and content editing, requires a top professional GPU Card. Today only one model is adapted to support VDI highly intensive workloads for this kind of use case. This is the Nvidia Ampere architecture A40.

“NVIDIA Virtual GPU (vGPU) enables multiple virtual machines (VMs) to have simultaneous, direct access to a single physical GPU, using the same NVIDIA graphics drivers that are deployed on non-virtualized operating systems. By doing this, NVIDIA vGPU provides VMs with unparalleled graphics performance, compute performance, and application compatibility, together with the cost-effectiveness and scalability brought about by sharing a GPU among multiple workloads.” (REF.8)

## Dell Technologies and NVIDIA Solutions

Recommended for Virtualization

	A100	A30	A40	A16
GPUs/Board (Architecture)	1 Ampere	1 Ampere	1 Ampere	4 Ampere
Workload	AI		VDI	
Virtualization Use Case	Best performance AI, DL training, data analytics, HPC	AI inference, data analytics, HPC	Virtual workstations	Virtual desktops
NVIDIA vGPU Software	AI Enterprise	AI Enterprise	vPC, vApps, vWS, AI Enterprise	vPC, vApps, vWS, AI Enterprise
Dell EMC PowerEdge Infrastructure	R750, R7525, VxRail, R740xd	R750, R7525, VxRail, R740xd	R750, R7525, VxRail, R740xd	R750, R7525, VxRail, R740xd
Max. GPUs/Server	2-4	2-4	2-4	2-4
Memory Size	80GB HBM2	24GB HBM2	48GB GDDR6	64GB GDDR6 (16GB per GPU)
vGPU Profiles (GB) <sup>1</sup>	4, 5, 8, 10, 16, 20, 40, 80	4, 6, 8, 12, 24	1, 2, 3, 4, 6, 8, 12, 16, 24, 48	1, 2, 4, 8, 16
Form Factor	PCIe 4.0 dual slot	PCIe 4.0 dual slot	PCIe 4.0 dual slot	PCIe 4.0 dual slot
Power (W)	400	165	300	250
NVIDIA RTX Technology	-	-	✓	✓
NVLink	✓	✓	✓	-

Figure 8 – Nvidia Ampere Architecture (Ref.9)

Learn more at [www.dell.com/certification](http://www.dell.com/certification)

Available on DELL PowerEdge servers or DELL EMC VxRail hyperconverged solution, the A40 model is offering one of the best performances at that time to address our VDI for editor's virtual workstation needs. Dell and Nvidia are offering an online recommendation available in the references (Ref.10)

### Size the Number of vGPU

After interviewing the editors, we had different templates (Figure 2 Workplace Templates) , the next step was to define the number of virtual machines and number of users for each template.

To design the proof of concept, an audit of each actual physical device is required to define the number of vGPU that will be required. Other things which must be taken into consideration are the number of screens and their resolution.

To audit, in an industrial way, the physical workstations we recommended the utilization of the following tools: **GPU Profiler** (Ref.11) and **GPU-Z** (Ref.12) on customer physical devices with real users. Doing so, it is possible using the friendly recommender (Ref.10) to choose the right model of GPU card and to have good insights on the Frame Buffer usage which is a key element in this typology of use case.

Analyzing the current environment of real editor is, of course representative, only when the GPU Profiler and GPU-Z are running during a real project where the users are working in a "real-life" situation.

### Advantage of GPU Sharing

Another key benefit of vGPU is the capacity of allowing resources to the right user at the right time. As an example, the resources could be significantly increased for dedicated editor during a media campaign that is business critical.




Figure 9 – Best Effort GPU Scheduler (Ref.13)

This mechanism also improves the balance of available graphical calculation processing resources to users. When workloads across virtual machines are not executed at the same time, or are not GPU bound, the performance can exceed the maximum normal performance dedicated to a user by using the “Best Effort GPU Scheduler” mode.

A comprehensive example would be the “Team coffee break.” When all the team members expect one to be at the coffee machine, the remaining user that is still working will have the combination of available resources at his disposal.

On the other hand, a traditional mode exists to “bound” a vm to a dedicated GPU. (Ref.13)

### Size the server or Hyperconverged Infrastructure (HCI) platform

For relevant proof of concept during an editorial workflow, a minimum configuration in terms of calculation processing unit speed and memory is required.

Our recommendation should a customer be interested in using servers (or HCI) is to have a configuration with a minimum of Dual Bi-Processor with CPUs (central processing units) ~3 or 3.2 GHZ with two Nvidia A40 cards. We also recommend avoiding exceeding 8 VMs per physical GPU

The Remote Access memory (RAM) is also important and needs to be sized high enough, as in a vGPU environment, RAM is fully reserved for each started VMs. More guidelines are available through the Nvidia documentation (Ref.9).

Another key point: if you choose an HCI platform using VMware instead of bare metal servers, is that some additional CPU and memory are required for the good operations of the hypervisor.

Learn more at [www.dell.com/certification](http://www.dell.com/certification)

## Choosing a Virtualization Solution

Like our previous choice in terms of the right software for shared storage solutions on Edge locations, the choice of the virtualization solution must involve the end users. But as a component of the IT infrastructure, it is also a corporate choice. The usual choice we mostly see on the market is between Citrix (Ref.6) and VMware Horizon solutions (Ref.7).

In our case we choose the VMware solution even if choosing another VDI solution will not change the hardware design or the other components of our editorial platform.

Again, we strongly recommend a test to run proof of concept.

## Bringing Pieces Together

### The Global Architecture

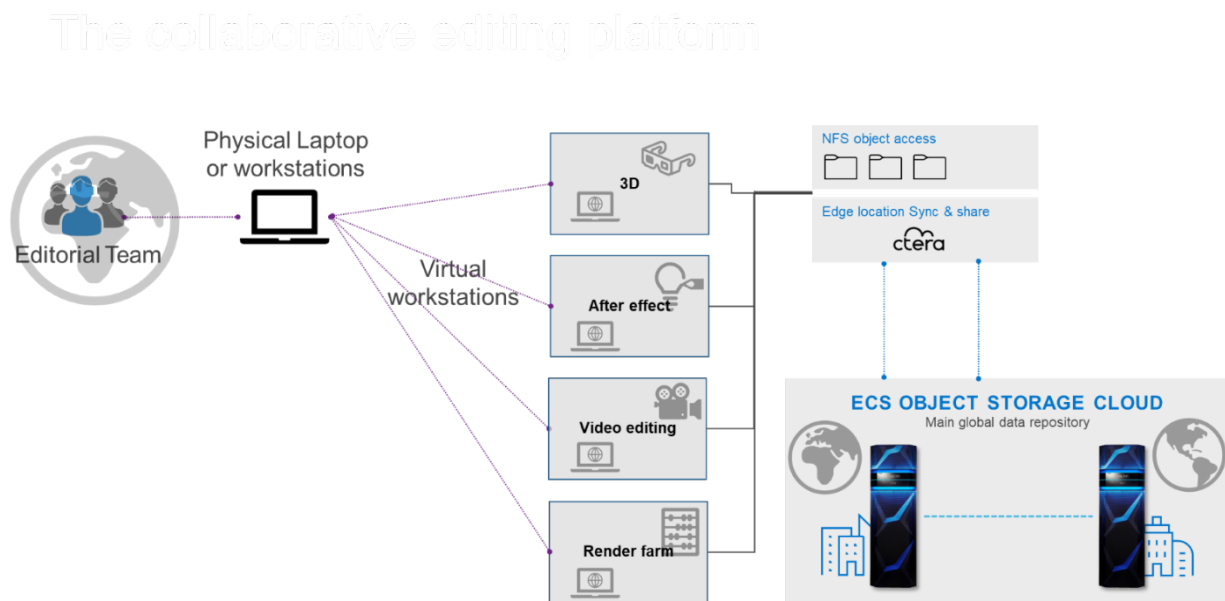


Figure 9 – collaborative editing platform

By combining the pieces together: the global data repository, the edge locations, and the end-user workplace; -we now have the design of our complete platform answering the challenges of a **collaborative editorial infrastructure**.

But what about the data lifecycle?

## The Data Governance

By creating a virtual editorial environment, we offer an agile way to provide workspace with the required performances and applications needed by editors. The applications that they are using generate new data. But it is always easier not to begin a project with a full blank paper, but to have some input to capitalize on.

In the media world, footage (videos or audio) is an important source for fueling the project. Editors need to access and visualize their footage before beginning the editing work. In a collaborating system, the footage needs to be manually organized, metadata will help in the classification process.

Tools exist to help them in this organization, running intelligent algorithms that generate labels and tags which can improve the indexing of the data for business efficiency.

What about security? Security is everyone's matter and concern. To respect the compliance and corporate guidance, the access and end users' actions must be required to be logged in for an audit perspective.

The role base access control of our design is quite flexible. The beginning of the workflow starts with access to the virtual environment. This solution is easy to manage and centralized.

The access to the data is handled on a software level. Current projects are hosted in the Edge locations using the caching shared storage mechanisms.

The "Corporate Memory" or the archive of previous projects, footage media campaigns... etc. are hosted and centralized in the main global repositories where accesses are under a higher control. Using ECS as main global repository offers mechanisms to prevent data deletion or alteration and prevent data loss or malicious behaviors.

## Conclusion

Transforming or building an editorial platform specifically designed for content edition teams is a journey where you need to address multiple challenges such as:

- the geographical multiple locations of users
- the need in term of performance
- the management of the data
- the standardization of format
- the security and administration
- the shared global storage to the core data centers
- performant shared storage to the edge

Learn more at [www.dell.com/certification](http://www.dell.com/certification)



- ...

What if all challenges could be addressed by one seamless platform built in editing collaboration purpose?

Empowered by the guidelines in this document you can begin your journey now!

## Bibliography

1. [how collaborative editing makes creatives more creative isilon ebook](#)
2. [DELL EMC ECS Enterprise Object Storage](#)
3. [Panzura website](#)
4. [Ctera website](#)
5. [Dell Technologies PERC11 Controller Introduces RAID Support for NVMe Drives](#)
6. [Citrix virtual solutions](#)
7. [Vmware Desktop and App virtualization](#)
8. [NVIDIA Virtual vGPU documentation](#)
9. [NVIDIA Ampere A40 Card](#)
10. [vGPU recommender](#)
11. [Download GPU profiler](#)
12. [Download GPU-Z](#)
13. [Tech-Brief-Virtual-GPU-Positioning.pdf](#)-Virtual-GPU-Positioning.pdf

Dell Technologies believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

Disclaimer: The views, processes or methodologies published in this article are those of the authors. They do not necessarily reflect Dell Technologies' views, processes, or methodologies.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." DELL TECHNOLOGIES MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying and distribution of any Dell Technologies software described in this publication requires an applicable software license.

Copyright © 2022 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.