

A full-page background image featuring an astronaut in a detailed space suit, floating in space. The astronaut's helmet is prominent, with a bright light reflecting off the visor. In the background, a large, blue and white Earth is visible against a starry space backdrop. The overall color palette is dominated by blues and whites, creating a futuristic and high-tech atmosphere.

Media & Entertainment HPC and AI reference Guide

Contents

I. Introduction	5
Abstract	5
II. Industry and Applications	6
Applications	6
How do we “render” an image?	7
What is a render farm?	7
Cloud rendering	9
Production management software	10
VFX Craft Groups – Journey of a frame	11
Pre-Production	11
Production	12
Post Production	13
III. How to tune the perfect solution?	14
Solution mock-up	14
Workstations	16
When does it make sense to virtualise?	17
Virtualised workstations with VMware	18
Monitors	19
Calibrate Display	19
Network	20
Isilon Leaf-spine network architecture	21
Render farm	22
Head node	22
Compute nodes	22
Extreme Density Server	24
Storage	25
Option_1: Dell EMC PixStor and Excelero Solution	25
Option_2: Isilon - Home and Buffer in one	27
Isilon Compression	28
PowerScale Dedupe	28
Option_3: Dell EMC Ready Solutions for HPC NFS Storage	29
Cloud and virtualised computing	30
Dell Technologies offers multi-cloud choices	31
Services & software	32
Intel® oneAPI Rendering Toolkit	33

Contents

IV. The Future, Trends & Challenges	36
Challenges in the industry	37
Security and best practices	37
Media and Entertainment – Security considerations and difficulties.....	37
How can facilities work from home if there are heavy restrictions on remote access?.....	39
Surely, they can use cloud	39
V. Glossary and References	41
Glossary	41
Reference	41
Customer stories	42

Contacts & More Info

Contact Dell Technologies team:

Gabriel_Pirastru@Dell.com

HPC_NA_Sales_Team@Dell.com

EMEA_HPC_Team@Dell.com

APJ_HPC_Team@Dell.com

HPC_Latam@Dell.com

Dell Technologies HPC

<https://www.delltechnologies.com/en-us/solutions/high-performance-computing/index.htm>

M&E Dell Workstations

<https://www.delltechnologies.com/en-us/industry/media-entertainment.htm>

M&E Dell Unstructured Data

<https://www.delltechnologies.com/en-us/storage/media-entertainment.htm>

M&E Dell PowerScale

<https://www.delltechnologies.com/en-us/storage/media-entertainment.htm>

M&E Dell Solutions

<https://blog.dellemc.com/en-us/tag/media-entertainment/>

I. Introduction

Abstract

This paper will give the reader a comprehensive understanding of the components of a high-performance computer cluster (render farm) within the M&E space, an outlook of where the industry is heading, an understanding of commonly used applications, and industry best practices.

“The most important thing is that you have to have the visual effects working for you, instead of you working for the visual effects.”

David Heyman
Producer on Gravity that won Best Visual Effects 2014

Computer-generated imagery (abbreviated CGI) and Visual effects (abbreviated VFX) is where computer-generated imagery is created and combined with footage from a camera. It was once a luxury only affordable to productions with big budgets but now with the decrease in hardware costs via utilisation of mass-produced hardware and affordable software, VFX has become mainstream so much so that premium television content also requires high-quality CGI and VFX, making it a vital ingredient in the magical cauldron of movie-making, storytelling, and entertainment. Advancements in modern CGI are only made possible with the aid of sophisticated software algorithms and powerful and robust hardware.

High performance clusters of computers, large pools of storage and high bandwidth networks are the foundation for which a VFX pipeline is built upon. A typical studio will employ specialist engineers in areas of networking, storage, and infrastructure. When it comes to software developers it is not uncommon for developers to have masters and doctorates in their fields of expertise. Production studios blur the lines between several disciplines wearing multiple hats. Are they studios filmmakers, or are they software developers, perhaps they are big data creators or computational crunchers? These days one could say they are all and more.

Artists play an endless game of cat and mouse with technology. As technology advancements allow for faster and more powerful platforms, artists push the creative boundaries producing ever more so complex geometry, lighting and effects, which in turn pushes the technology to its limits thus creating a new set of challenges for technology partners to innovate and overcome.

II. Industry and Applications

“In today’s film and games industry, art production is a communal enterprise. Every CG element that the audiences see (or doesn’t realize that they see) in a theatre, and every character or environment that a player meets in a video game, is the result of the collective effort of many artists and technicians.”

(Renee Dunlop, 2014)

Media and Entertainment encompasses a number of current and emerging industries such as Graphic Design, Movies and Television, Visual Effects, Animated Features and Games. All of which have different workflows and challenges, but all have in common a use for technology to be the tools needed to allow highly skilled creative artists the ability to create amazing art.

Applications

There is a multitude of technologies that are used in a VFX pipeline below is just a small example of some of the applications and terminologies.

Rendering engine is software that is responsible for generating a visual output to a screen or file. A rendering engine takes into account input variables such as light data, camera data, geometrical data such as 3D polygons, etc. Depending on the software rendering engines can be optimised to run on CPU and/or GPUs and can be highly parallelised. Rendering engines typically benefit from multi-core high-frequency clock speeds and large caches.

3D modelling applications provide artists with a comprehensive toolset required to build and create 3D environments and CG Models. Different applications have their strengths and so it is common to find several applications being used throughout the pipeline.

Compositing applications are used to layer various outputs together to form the final movie-quality image. Compositing artists need tools that are versatile enough to handle a multitude of formats and codecs.

Lighting applications utilise real-world physic-based computations to calculate how light from a virtual light source will interact with virtual objects in a scene. These calculations allow artists to create realistic images of shadows/occlusions, spotlights, area lights, light decay, shadow maps, etc.

Examples of applications

Rendering Engines	3D Modeling
3Delight Arnold Maxwell Render RenderMan	Sidefx Houdini Autodesk Maya Autodesk 3Ds Max Cinema 4D Blender
Compositing	Lighting
Nuke Fusion Adobe After Effects HitFilm Pro	Katana Maya Redshift V-Ray

**THE
FOUNDRY.**

NUKE

KATANA



Blackmagic
Fusion



CINEMA 4D



3DS MAX

blender



How do we “render” an image?

Rasterisation

Rasterisation is the most common technique used to create images, objects or 3D scenes, where the GPU draw (render) a 3D scene using two-dimensional polygons, typically triangles called a triangle mesh. After the scene is drawn it is then rasterised by converting the triangle mesh into individual pixels on a 2D screen. A shader then processes the pixels applying textures, colours and lighting effects producing a completed rendered image e.g. OpenGL, DirectX.

Ray Tracing

Ray tracing is a realistic lighting technique that is used to light objects in frame by emulating the way light paths from virtual light source reflects and refracts of objects along an image plane, just like light would interact with objects in the real world e.g. Intel® Embree Ray Tracing.

Benefits of Ray Tracing

Better image quality and faster feedback also typically easier to implement. V-Ray on Intel® Xeon® Scalable Processors takes advantage of hardware infrastructure, including open-source tools, so there’s no worrying about compatibility, memory limits or updating drivers.

CPUs vs GPUs

CPUs (especially Xeons) make up 99% of the render farms today, and GPUs are playing a big role in the artist workstations in the VFX pipeline for real-time visualisation, obviously orchestrated by Xeons at the heart of the client and remote workstations. GPUs are starting to play a role in render farms and will be important to consider going forward.

What is a render farm?

You will be pleasantly surprised that within the M&E space no animals are harmed when it comes to the render farm. But before I begin it would be prudent to first define a few terms that will be used throughout this document. A more comprehensive lexicon is located in the glossary section. The following terms are used loosely in the industry and vary from site to site.

Render job/sub-job/task/instance – Is an encapsulated package of commands to be executed by the worker. A Job can contain one sub-job with one task, or a job can contain many sub-jobs with many tasks. A job with many sub-jobs is sometimes referred to as a chunk.

Chunk - Chunking jobs is highly beneficial when application start-up and scene load times are significant compared to processing times. Chunking loads the scene once and renders many frames saving time on application start-up and scene load.

Server/supervisor/master – A central server that is responsible for the submission, dispatch, monitoring and coordination of jobs.

Client/worker/slave – A node designated to be managed by the supervisor for executing a package of work.

Dependencies – Jobs that require other jobs to start or finish before they can start processing are dependency constraint.

Render wranglers – The department that is responsible for the render farm in its entirety and is the first level of render support for Artists. Ensuring render jobs are scheduled correctly, failed to render jobs are debugged, faulty nodes are dealt with and several other various tasks.

A render farm within the M&E space refers to a networked cluster of computer nodes typically arranged in a master/worker model to coordinate and execute predefined tasks in an efficient and highly parallelised system. Nodes can consist of a mix of hardware configurations, workstations, dedicated servers and cloud-based compute nodes, located either on-premise, data centre or in the cloud.

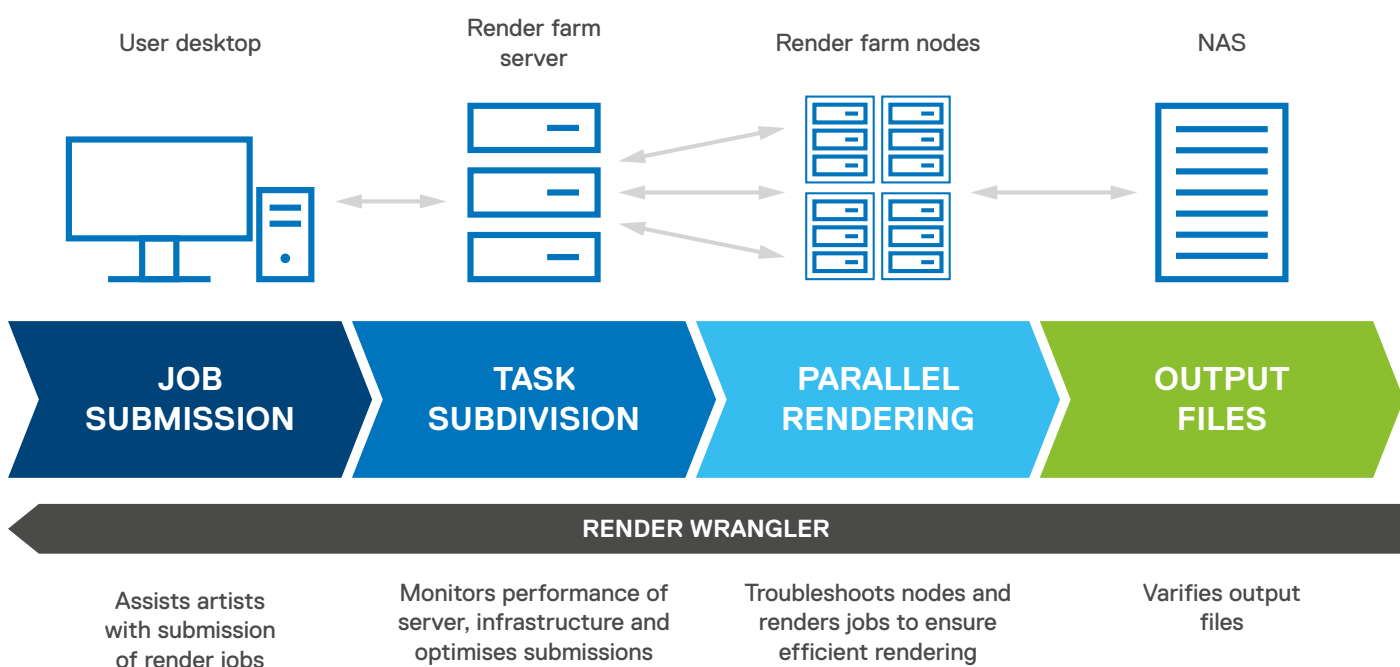
The master consists of a scheduler and queue manager whose responsibility is to accept jobs and constantly evaluate and groom the backlog based on various criteria’s, from simple FIFO model to more complex models that factor in production priorities, farm partitioning via the show and/or user quotas and availability of appropriate render resources. Job stats collection is either centralised or handled via a separate process and typically records a jobs

runtime, memory consumption, submission time, user info, runtime logs and error logs. More advanced render farms monitor job health and flag to the render wranglers when a job or render node does not look healthy.

The scheduler is the brains of a render farm and is the real differentiating component between render farms commonly found in studios. A robust and versatile policy-based scheduler that can manipulate the order that jobs are dispatched is advantageous in ensuring that the farm is processing job in the correct order that makes sense for production to meet their schedule and deadlines. Job priorities tend to change often so having the ability to shift resources from one job or shot to another as quickly and efficiently as possible with minimal loss in compute time spent is vital.

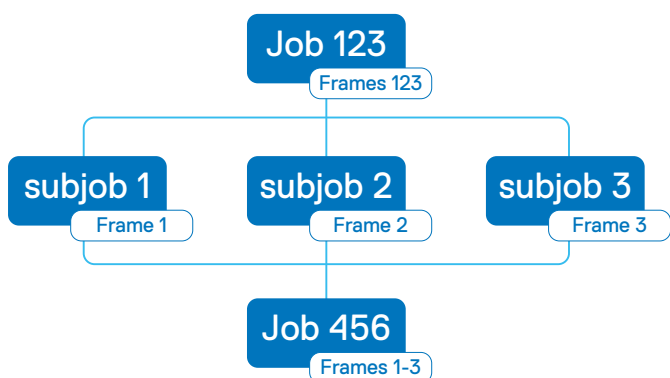
The render farm manager and scheduler have to also track license requirements of render jobs. A scheduler will typically query a license server via an API or command line example FLEXlm or RLM and dispatch render jobs accordingly.

The job package is dispatched to the worker to be executed. It contains all the information necessary to set up the right environment variables, defined application and input paths, etc. A basic render starts with establishing the appropriate runtime environment, loading the render application, opening the scene to render, pulling in the prerequisite input files from a shared file server and then rendering the scene file where finally once the render is complete the output is saved back on the shared file server. The render farm server keeps track of a job's progress and any errors via standard I/O streams.



Sometimes a render may require pre and post-processing in the form of dependent job. Job Graphs are an effective way to visualise a render job's dependency tree. Jobs can be chained to start after any number of upstream jobs and also have any number of children jobs chained after it.

Example of a render job graph



In the early days studios would write their proprietary render farm manager that was highly integrated within the studio's pipeline. In some cases, advanced custom schedulers that went well beyond FIFO/FIFO models were developed to allow production great control over the order of which render jobs were processed. Render wranglers would monitor the jobs and ensure no job was stuck, failed and ensure that the output files are consistent and expected.

Below is list of a few off the shelf render farm software

- Tractor 2 - renderman.pixar.com/tractor
- Deadline - www.awsthinkbox.com/deadline
- Qube - www.pipelinefx.com/
- Flamenco - www.flamenco.io/
- Open Cue - www.opencue.io/

Cloud rendering

Cloud is on the verge of becoming a mainstream solution where more companies are offering cloud rendering as a service than ever before. Cloud vendors have worked hard to ensure adequate available capacity and that the cost is relatively comparable to the total cost of ownership of traditional on-premise render farms. Industry compliance bodies such as the MPA and TPN are working the industry to ensure all security guidelines and certifications are in place. The consensus is that the movie industry is embracing, albeit slowly, cloud services and it is only a matter of time that cloud rendering, backup, remote workstations and other parts of the pipeline all run in the cloud.

Cloud render farms have recently gained a lot of momentum within the media and entertainment industry. As film budgets get smaller, fierce bidding among competing studios reduces profit margins and increased appetite for rich blockbuster quality VFX, studios are looking for ways to decrease costs and are looking at reducing the total cost of technology.

One would think a cloud on-demand, pay for what you use model would tick all the right boxes, unfortunately, it is not that simple. While the flexibility provided by pay for what you use model is very enticing a VFX pipeline can be very complex with years of legacy software, complex workflows and reliant on custom scripts that may not be suited to run outside of the studio ecosystem.

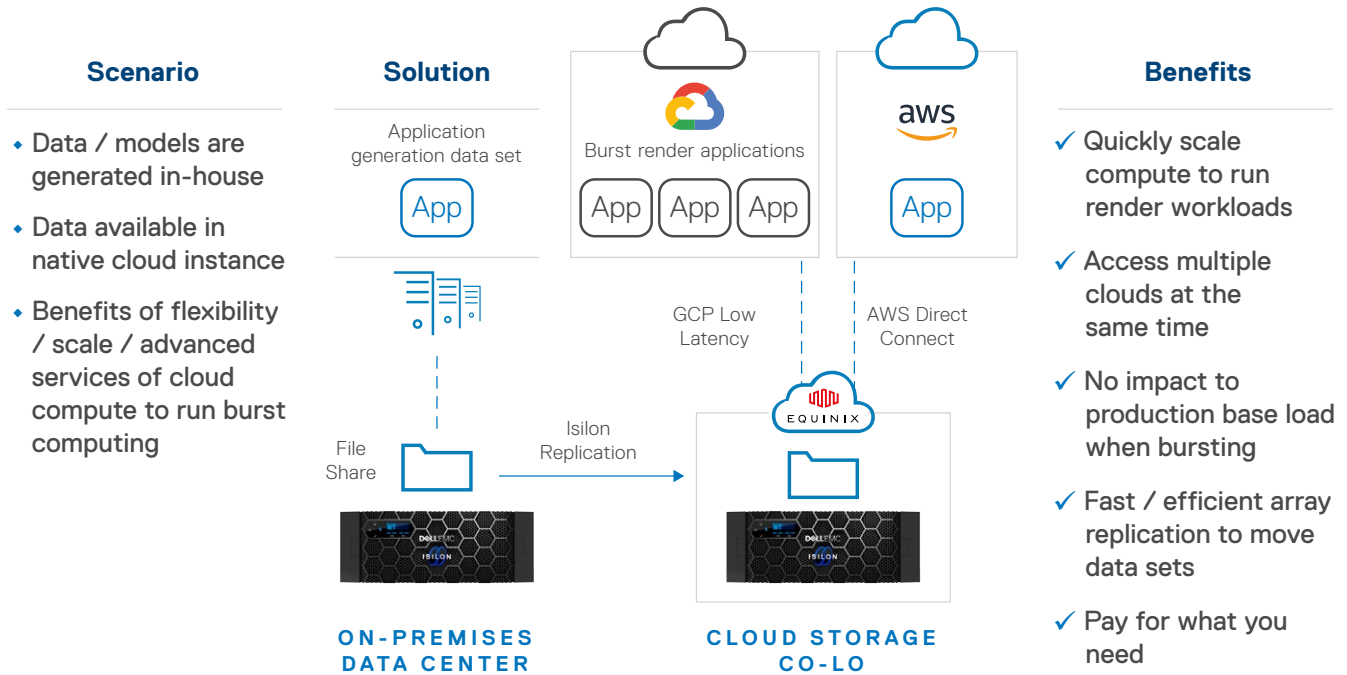
Cloud has been proven to work within the M&E space and the idea that a studio can scale on demand and then drop when not needed is very enticing and suits the workflow of production. The holy grail would be to seamlessly scale up and down as needed without any noticeable performance degradation and costs to be on par or cheaper than on-premise solutions.

In some regions of the world there is simply not enough compute power to allow a moderately sized studio to scale, if you think about several active shots with a few passes each 100 plus frames and all needing to be rendered and certainly not within the ease of setup compared to a local render farm.

There are several considerations that need to be thought out before a studio can jump into the cloud.

1. Cloud readiness: Studio Pipelines needs to adapt to cloud workflows
2. Billing Models: Software licenses need to be updated to suit cloud-based infrastructure. Cloud pricing needs to be better than the total cost of ownership of on-premise infrastructure
3. Complexity Scale of hardware needs to be instant, balanced and adaptable to suit studio needs
4. Cloud security: The Industry needs to set acceptable Security guidelines, cloud best practices and certify "approved" cloud vendors
5. Connectivity, performance and costs: Edge caching technology vs X number of Direct Connects < 8 milliseconds
6. Vendor flexibility

Use Case: Burst render in cloud



Production management software

typically refers to the management of scenes, shots and assets. It is a centralised file management platform that tracks the progress of shots, file and image review notes and iterations of changes. It also provides the production with reports on a shots progress and in some cases can also integrate with a render farm manager to track shots compute and storage consumption.

Examples of Production Management software

- Artella - www.artella.com/
- Ftrack - www.ftrack.com/en/
- Prism Pipeline - prism-pipeline.com/
- Shotgun - www.shotgunsoftware.com/

Breakdown of a Typical Production Pipeline and VFX Craft Groups

In its basic/linear form a typical 3D Production pipeline has three main stages

Content Creation Workflows




Sizing Questions	
• Number of concurrent streams per editor	• Number of editorial seats
• Codec, resolution & bitrate	• DPX-based or clip-based
• Mac or PC platform	• Throughput requirement / stream


Sizing Questions
• Render Nodes dedicated to 3D & 2D rendering
• CPU vs. GPU rendering / Pre-Vis
• Avere or other caching appliance in use
• NFS, SMB, both?

3D Production Pipeline

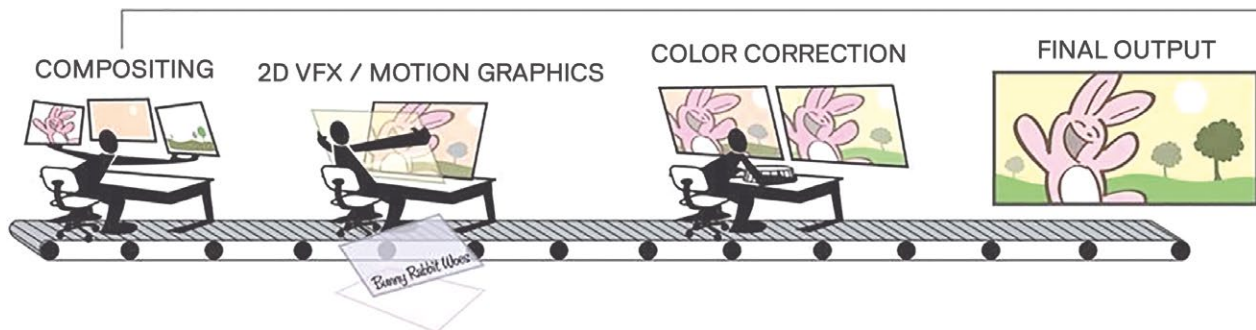
Pre-Production



Production



Post-Production



Pre-Production

Pre-production is where it all starts, it is where ideas are drawn, storyboards are created, and the pre-shoot plan is drafted. Depending on the show this stage can last from a few weeks to over a year.

Research & Development

R&D is typically the first phase in the VFX pipeline. It is where highly skilled and sometimes specialised developers create the customised software and techniques needed by artists to create the desired look with ease and consistency. R&D department tend to be highly technical software developers relying on solid foundational computer science and cutting-edge research. Siggraph technical papers are a great resource for where the industry is heading from a software and technique point of view. s2019.siggraph.org/conference/programs-events/technical-papers/

Art Department

The Art department is responsible for translating the written script and directors' vision into still imagery used to seed and guide digital artists. The collection of images forms a storyboard that gives direction and confirms the look, feels and style of the movie. Technology in this department has evolved from traditional paper and pencils to software-

based drawing suites like Adobe photoshop or similar, coupled with interactive pen displays, tablet styluses, colour accurate monitors, large amounts of storage and workstations with lots of memory.

Pre-visualisation

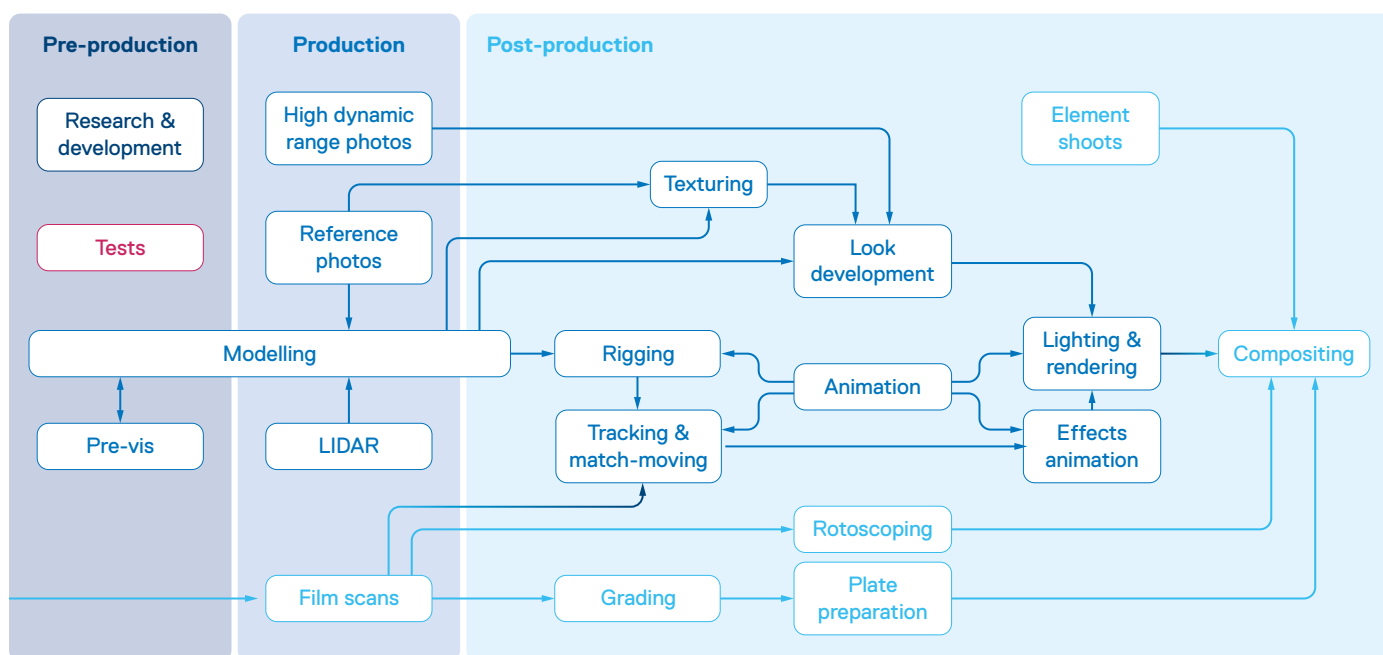
To better prepare for principle photography pre-viz gives the director and production additional context by creating a low-quality version of each VFX shot from the script and storyboard. By this stage software, techniques and workflows are refined to achieve the directors desired look.

Production

Production is where the scenes are virtually created, images are rendered, and creative adjustments happen via an iterative workflow.

The above image is considered basic in that it shows a very linear workflow where work is passed down the line but in reality, modern 3D pipelines processes are complex, highly parallel, iterative and have non-linear workflows as per the image below

A Typical VFX Pipeline



Matte painting

Everything from realistic backdrops to Alien-like environments, Matte painting artist create backdrops using digital or traditional painting methods. Matte paintings are still used today as they provide unlimited creative control and most of the times it is more practical than travelling to a location for filming. Like the Art Department Adobe Photoshop or similar, is typically used to create extensive matte paintings. Requiring interactive pen displays, tablet styluses, colour accurate monitors, large amounts of storage and workstations with lots of memory.

Asset Department

Is a department of riggers, texture painters, modelling artists and sometimes shader writers that work together to create digital assets or props of the real world and or fictional objects. The digitally created assets are then used further down the pipeline by other departments as the scenes are being built.

Animation

Animation artists bring to life digital characters or 3D models by manipulating the digital asset. They work with the rigging department to ensure that movement points, set up by the rigging artists, can be adjusted to achieve the desired motion as directed by the director.

Post-Production

Matchmove (also referred to as Motion Tracking)

It is through the analysis of real-world camera movement and matching it to virtual cameras that allows downstream artists to insert 3D digital objects convincingly into real-world footage.

Rotoscoping

Is a painstaking task of tracing and masking parts of a frame for extraction from a live-action shot. Green/Blue screens are used as backdrops to remove a designated colour from the scene also known as colour keying.

FX Simulation

An FX artist is responsible for creating water, smoke, fire, explosions and many more effects through dynamic, procedural simulations and fluid, particle systems. Wherever you see fire or water, cloth on a character, hair on a creature, a dust storm or fire, you likely wouldn't notice it was digitally constructed, but somewhere an FX artist digitally constructed the element to be as real as the real thing!

Lighting

In live-action lighting, artists apply digital lights to a scene to match the lighting in the real-world footage. This allows for a convincing blending of real digital assets with live-action footage. Lighting artists in a full CG movie apply their craft via digital lighting rigs in line with director notes to bring environments, characters and key elements in a scene to light akin to a real light in a traditional movie set. Computer software calculates the physics and the lighting artist adjusts to achieve the desired look.

Compositing

Compositing Artists artistically layer all the output elements from all the upstream departments. Compositors blend seamlessly all the live-action, matte paintings, lighting, animation, simulated effects into the final shot. Unfortunately, a compositor is at the last link in the chain and so sometimes have to work within very tight deadlines.

Production

Consists of several roles that do not fall into any of the artistic craft groups. Roles such as management teams e.g. production management, artist management, production accountants, producers, schedule managers. They manage artists and ensure that shows are on track and in the budget. Production primarily require laptops that have long battery life, are light weight and have backlit keyboards so that they can type notes in the dark during review sessions.

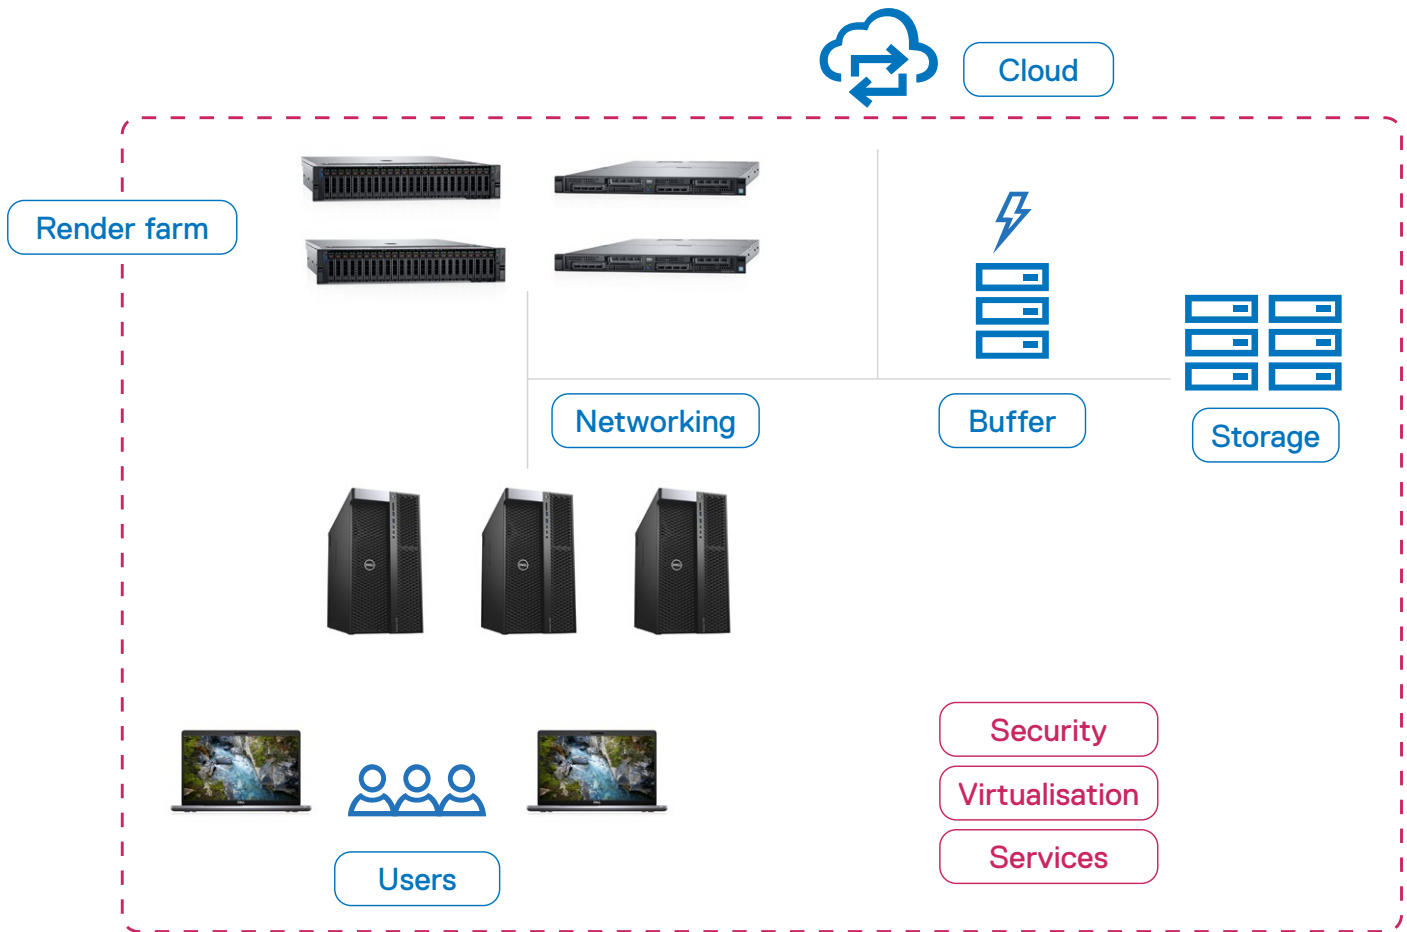
How is Technology used in departments?



Most departments tend to have similar hardware requirements with configuration variations in the amount and/or speed of RAM, fast scratch disk and CPU. Variations tend to be generational based due to purchase cycles. Ideally hardware spec variations are kept at a minimal allowing for departments to pick up and move to different parts of the facility without having to also move their workstations. Department moves are common in studios mainly because departments grow, or various departments need to sit in close proximity to aid in creative collaboration. Workstation design is very important, with emphasis on weight, carry handles and ease of accessibility to internal user serviceable parts highly desirable.

III. How to tune the perfect solution?

Solution mock-up



The solution takes in as example a team of 8 creatives working in a special effect studio.

Users	Option_1: Precision T7820 Intel® Xeon® D/W/E series Gold 6234 3.3GHz, (4.0GHz Turbo, 8C, 10.4GT/s 3UPI, 24.75MB Cache, HT (130W) DDR4-2933) Windows 10 Pro 32GB 4x8GB DDR4 2933MHz RDIMM ECC Memory GPU Boot-drive: 1TB PCIe NVMe Class 50 Solid State Drive Storage: HDD 1TB Quietness
Option_1: Workstation (x8)	
Option_2: R740 (x2)	Option_2: Virtualised workstations Wyse 5000 Series Thin Client with PCoIP R740xd Intel® Xeon® Gold 6242 2.8G, 16C/32T, 10.4GT/s, 22M Cache, Turbo, HT (150W) DDR4-2933 GPU
Monitor (x8)	Dell UltraSharp 32 Inch 4K UHD (2160p) 3840 x 2160 (DisplayPort: 60 Hz, HDMI: 30 Hz) 1.07 billion colors Ultra HD 4K Monitor with PremierColor - UP3216Q
Admin Network	Dell EMC PowerSwitch S serie 10 GbE switch

Compute Network	Dell EMC PowerSwitch S serie 50 GbE switch
Render Farm (x2 HA) Head node	<p>PowerEdge R740xd (2 sockets) Intel® Xeon® Silver 4208 2.1G, 8C/16T, 9.6GT/s, 11M Cache, Turbo 2.0 to 3.0, HT (85W) DDR4-2400 1TB HDD 960GB SSD ME024 storage array for high availability</p> <p>Optional: Two GPUs supported with NVMe configurations for virtualisation VMware® ESXi® for virtualised computing</p>
Render Farm Compute Nodes (x10)	<p>PowerEdge C6420 (2 sockets) - 4 servers 2 U (C6400 enclosure) Intel® Xeon® Gold 6254, 18C, 3.1 Ghz, DRR4 Memory: 12x32GB</p>
Storage – Option 1	<p>Home Storage PixStor Dell Solution</p> <p>Buffer Storage Excelero Dell Solution</p>
Storage – Option 2	<p>Home and Buffer in one</p> <p>F800 (All Flash) H500 Isilon (Hybrid) ECS (Archiving)</p>
Storage – Option 3	<p>Dell NFS Ready Solution High-availability 84 - 3.5" NL SAS drives, up to 12TB. Supports up to 1008TB (raw space) 8 LUNs, linear 8+2 RAID 6, chunk size 128KiB. 4 Global HDD spares. Peak performance, with read at ~ 7 GB/s, write performance ~ 5 GB/s</p>
Cluster Manager	<p>Open Source Option: OpenHPC - XCAT Cluster Manager</p> <p>Enterprise Option: BrightComputing</p>
Software	Intel® OpenAPI Rendering Toolkit
Security and Services	<p>Dell Endpoint Security Dell Services Dell Services for HPC</p>

Workstations

Dell Technologies offers a wide array of workstations that can be in the mobile format. They are divided them into three categories, 3000, 5000, 7000 from entry to high-end.

The workstation is where the creative process starts, this is the main tool of creators and designers.

How to choose the right workstation?

Core count/ Frequency

Check how your application behaves to best match it with the relevant CPU. Workstation are available with the latest Intel® Xeon® processors. As a rule of thumb, it is recommended to aim for CPUs with high frequencies (from 3GHZ) and at least 6 cores for applications such as Autodesk. Large caches and memory channels per CPU type do make a difference in performance.

For choosing the best CPU we recommend **checking the application benchmarks** such as with Autodesk CPU benchmarks

Memory

This might be the main bottleneck in the creative process. You absolutely don't want your frame to be sitting in the disk – which would dramatically decrease performance. To avoid that, choose a high memory count (i.e. >=32GB).

Solid State Drives

Previously in the expensive end of the spectrum, SSDs have now become cheaper and speed-up both applications and file access significantly.

Graphic cards

GPUs have proven to accelerate M&E applications, different vendors now offer them and checking performance/ price is important.

GPUs must be certified by software vendors, it is therefore advised to pick certified cards, such as the hardware certified by Autodesk.

Liquid Cooling

To get best performance out of your workstation, cooling is key. Dell has recently integrated liquid-cooling capabilities in its workstations which allow you to run more powerful CPUs longer.

Licences & Certification

In case you run into a problem with the software you want to be sure that the hardware has been certified in the 1st place. The good news is that Dell workstations have been certified by all major applications in the M&E space. You can find the full list of certification on this page.




Noise

A noisy workstation might impede work of the creatives, therefore it is recommended to have a good understanding of the noise level of a device running overclocked CPU with GPU enabled to make sure that getting best performance doesn't come to the cost of inconvenient noise. The T series of Dell are renowned for their quietness.

www.delltechnologies.com/nl-nl/industry/media-entertainment.htm

Dell Mobile Workstations

○ N/A ○ Good ● Better ● Best

	Performance Options	Professional Graphics Performance	Memory Capacity	SSD Storage	Display Resolution & Colour	Mobility	Docking
 3000 Series	●	○	●	●	○	●	●
 5000 Series	●	●	●	●	●	●	●
 7000 Series	●	●	●	●	●	○	●

www.dell.com/en-us/work/shop/workstations-isv-certified/sf/precision-laptops

Precision Fixed Workstation

Compare the full breadth of our ISV-certified fixed workstation portfolio.



3000 Series

Affordable tower workstations for professional performance.



5000 Series

Compact tower workstation for mission-critical performance and reliability.



7000 Series

Powerful tower and rack workstation for scalability.

Single Processors	○	●	●
Dual Processors	○	○	●
Liquid Processors	○	○	●
Performance Options	○	●	●
Storage Capabilities	○	●	●
Professional Graphic Performance	○	●	●

○ N/A ○ Good ● Better ● Best

www.dell.com/ag/business/p/precision-desktops

When does it make sense to virtualise?

Instead of owning large and expensive workstations/ laptops, another option is to optimise resources by using virtualised servers for all users. Choosing this infrastructure would mean having wyse/PCoIP with monitors and run computational tasks on the server in the backend. Therefore four to six high-end workstations can be replaced by one or two powerful servers.

Dell Technologies supports desktop virtualisation software from Citrix, Microsoft, and VMware. What makes our data centre hardware differentiated is that we make it a lot easier to connect that hardware in ways that's useful to

our customers. Dell makes use of reference architecture to build out sample environments. Dell has several certified partners: Citrix, MSFT, and VMware.

[Wyse 5000 Series Thin Client with PCoIP](#)

Users can either have 1:1 connection (from 1 remote workstation to 1 server) or share a pool of resources. When working in a creative team, this can be particularly interesting since some of tasks (e.g. drawing and sketching) don't require high level of computation.

Key benefits include added network bandwidth and better resource management. By sharing a pool of resources, users can get access to more when others need less.

The only challenge of virtualising the environment is to be able to dynamically manage resources, this is being done by a hypervisor, available Dell tested and supported options include:

- [Teradici PCoIP Software Client](#)
- [Amulet Hotkey – a Dell OEM and Global Alliance Partner](#)

How do you size it? Just multiply the necessary cores/ storage/ GPU needed by the amount of people and fit it into a server.

If the servers are located miles away, will there be a delay? Studies have shown that from 100milli-second (10,000th of a second) of latency one would feel a delay. Which is unlikely to happen even if servers are based in another continent.

[Virtualised workstations with VMware](#)

VMware Horizon is a modern platform for secure delivery of virtual desktops and apps across the hybrid cloud, from the market leader in [software-defined data center](#) and [digital workspaces](#). By leveraging unique integration with trusted VMware technology, Horizon helps IT efficiently deploy and scale virtual desktops and apps from a single control plane with rapid provisioning, automation, and simplified management to extend the best digital workspace experience to end users.

[VMware Horizon Page](#)

[VxRail HCI platform](#)

The solution configuration based on the VxRail HCI platform is recommended for running professional graphics applications such as Autodesk.

[Solution configuration for virtual workstations on VxRail](#)

The compute layer consists of VxRail V570F appliances. These are VDI-optimised and offer the highest processor speeds and graphics capability. The storage layer is powered by VMware vSAN software-defined storage technology. The network layer consists of Dell EMC PowerSwitch S5248 (25GbE ToR switches). Virtual workstations run on a VMware ESXi hypervisor. VMware Horizon 7 is the virtual desktop brokering software. NVIDIA Quadro vDWS virtualises the GPU hardware to provide virtual GPUs for each virtual-workstation.

For more information, see the [VMware Horizon VDI solution on Dell EMC VxRail Design Guide](#).

Monitors

Choosing The Right Display Monitors is key for productivity (see [Forrester Study](#)), you can choose the right monitor this by considering following factors;

- **Size:** ability to multitask and reduce the time required to gather information from multiple sources, applications, and/or less
- **Colour level:** for graphic works every detail matters. Colour accuracy and the ability to calibrate apply different Lookup tables (LUTs) are Key features.
- **Screen resolution:** effect on content creation quality through the ability to view extremely sharp images with a high level of detail
- **Overall comfort:** 84% of respondents indicate that being able to view a great deal of work on a large display monitor is important for their work

Dell UltraSharp PremierColor monitors – True colour resolution are the products that satisfy highest demands.

UltraSharp monitors offer nearly 100% sRGB color coverage, while UltraSharp monitors with PremierColor – designed for color critical tasks – offer virtually 100% AdobeRGB and sRGB coverage, as well as newer colour spaces like DCI-P3, REC 709, REC 2020 for video formats. With up to 1.07 billion* colours that's 64 times more color depth than standard monitors can provide. Customise colour parameters on PremierColor monitors to your preference using the Custom Colour mode, the SDK2* or the optional X-rite iDisplay Pro colorimeter. Your UltraSharp monitor is calibrated on sRGB to an accuracy of Delta-E less than four, and AdobeRGB to a Delta-E less than two for PremierColor models.

www.dell.com/en-us/work/shop/dell-ultrasharp-monitors/ab/dell-ultrasharp-monitors



Curved Monitors



PremierColor Monitors



InfinityEdge Monitors



Touch Monitors



4K & 8K Monitors

Calibrate Display

To get best colour resolution for your display, Dell helps you calibrate your device;

- Dell offers [X-Rite i1 Display Pro calibration package](#)
- Dell delivers a pdf manual indicates how to use the Dell Ultrasharp Color Calibration Solution along with the X-Rite package (shipped in a CD) to easily set-up your product

[Check the Dell Ultrasharp Color Calibration Manual here](#)

Dell UltraSharp Color Calibration Software

The link to the to the current version of the software can be found here: [Dell Ultra-Sharp Color Calibration Software \(DUCCS\)](#). You can also find the software download by going to [Dell's Drivers & Downloads site](#) and searching for DUCCS in the search bar at the top of the page. Then select Drivers & Downloads under the Content section of the next page to find the latest version of the Dell Ultra-Sharp Color Calibration Software.

Accessories

From keyboard to connected docks, Dell Technologies can provide you with all components that you need to get started with your work. [Check-out all our accessories on this page.](#)

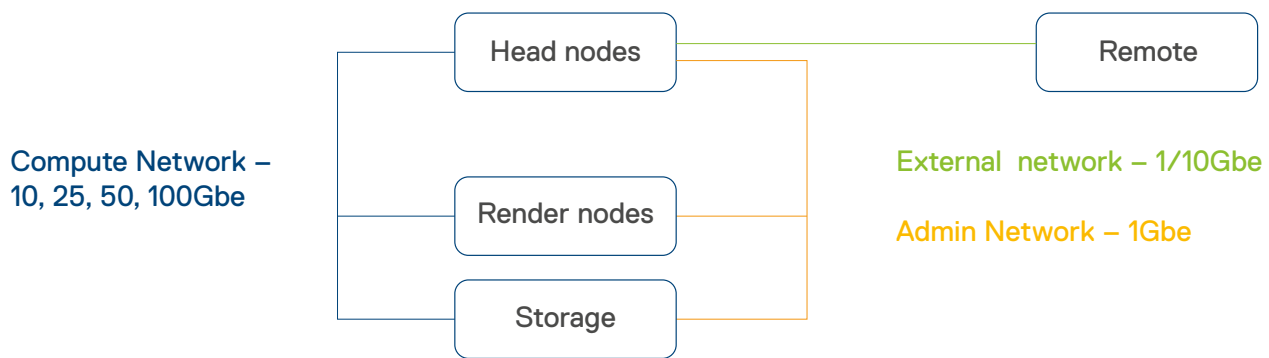
www.dell.com/en-us/work/shop/monitors-monitor-accessories/ac/4009

Network

There are mainly three types of network; admin, compute and external. Whilst admin and external can work with 1Gb ethernet, for the compute, the higher the better.

RDMA Fabric (OPA/ InfiniBand) isn't necessary for render farms since the problem they solve is 'embarrassingly parallel' i.e. nodes work independently on one problem, there is no need for MPI communication. That means that in media and entertainment workloads, jobs are independent, though they read the same data, they only work on one image at the time, this image is neither dependent on previous nor on the next one.

To help you build the best Fabric Design Centre Dell has created an online Network builder Tool fdc.emc.com/#!/network-fabric



Dell EMC PowerSwitch S Series 1GbE Switches



Dell EMC PowerSwitch S Series 10GbE Switches



Dell EMC PowerSwitch S Series 25+GbE Switches
25/40/50/100GbE ToR switches built for building high-capacity network fabrics that are easy to deploy, cost-effective and provide a clear path to a software-defined data center.



Dell EMC PowerSwitch Z Series Switches
1000GbE and 400GbE core/aggregation switches designed for building optimized data center leaf/spine fabrics of virtually any size.

www.delltechnologies.com/nl-nl/networking/data-center-switches/index.htm#accordion0

Isilon Leaf-spine network architecture

OneFS 8.2 enables the deployment of a leaf spine backend network switch architecture that increases the size, scale and performance of Isilon clusters.

The OneFS 8.1 operating system introduced the use of Ethernet switches for the backend node to node communication. Further, OneFS 8.1.1 introduced a choice of Dell Ethernet switches for the backend to simplify configurations and provide a Dell-on-Dell solution.

Leaf spine backend network architecture using a Dell switch solution the Isilon cluster can scale up to 144 nodes.

Advantages

A leaf spine network minimises latency and the likelihood of bottlenecks in the backend network, its architecture is highly scalable and built with redundancy.

Description

In a leaf spine network switch architecture, the access layer of the network is referred to as the leaf layer. The Isilon nodes connect to leaf switches at this layer. At the next level, the aggregation and core network layers are condensed into a single spine layer.

Every leaf switch connects to every spine switch to ensure that all leaf switches are no more than one hop away from one another. In addition, leaf switch to spine switch connections need to be evenly distributed meaning there should be the same number of connections to each spine switch from each leaf switch.

Example

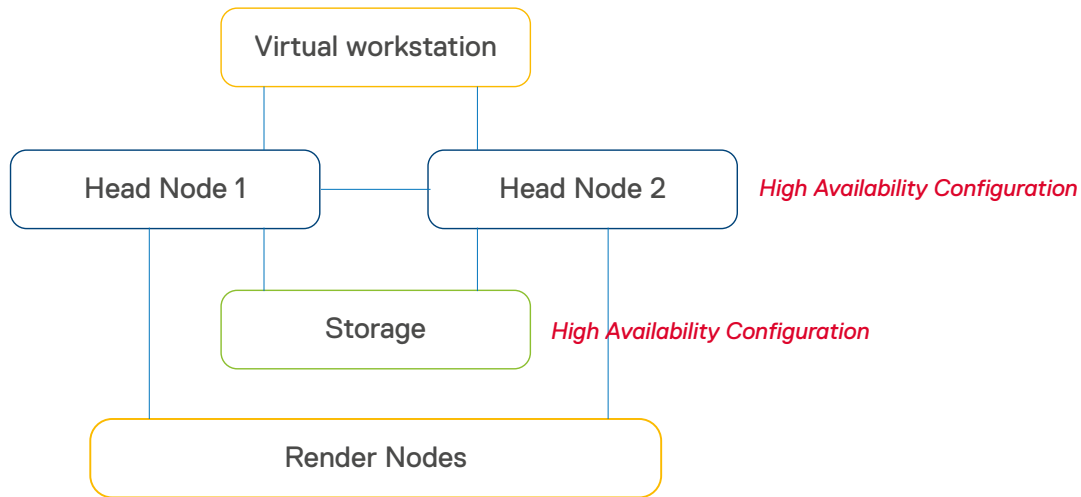
Component	Description	Connection considerations
Network Spine Switch	Dell Z9100-ON 32-port switch	Back-end network with 100GbE (uplink) connects to the leaf switch
Network Leaf Switch	Dell Z9100-ON 32-port switch	Downlink from the leaf switch to 40GbE performance nodes (32x40GbE maximum), or four 10 GbE archive nodes with a breakout cable
Isilon Performance Nodes (Flash and Hybrid)	F810, F800, H500, H600 and H5600	Performance nodes support a 40GbE connection to the leaf switch
Isilon Archive and Hybrid Nodes	A200, A2000, and H400	Archive nodes support a 10GbE connection to the leaf switch using breakout cable
Uplink	Leaf to Spine switch connection	100GbE connection
Downlink	Leaf to Isilon node connection	40 or 10GbE

[Layer 2 topology configurations](#)

[Leaf-Spine Cluster Installation Guide](#)

[Dell EMC PowerScale: Leaf-Spine Network Best Practices](#)

Render farm



Head node

The head node is the 'brain' of the render farm, this where the admin takes places; e.g. installing/ running software packages, managing, deploying and provisioning the render farm.

The head node can be where the user logs in, or another node can be provisioned to separate the two. Coupling a second node for high availability is recommended to make sure that the render farm is always running even in case of issues.

Considerations to select the head node:

CPU:

It is not necessary to have high-end CPU/ RAM but a mid-tier option is more suitable. Processors with good turbo modes are preferred, so that the node can run with fewer cores/ higher frequency. The Intel® Xeon® Scalable Processors Silver range is well suited for the task.

Storage:

A storage array is added for high availability. Otherwise, the head node can host several TB of storage, therefore smaller render farms can actually have the home storage (NFS solution) directly hosted in the head-node. In a typical environment, the storage array attached to the head node will store a NoSQL database such as the Dell MongoDB solution. Virtualisation software VMware, Citrix, applications and the cluster manager and render farm manager will be hosted in the array. NVME drives can be used for improved performance and reactive applications. Having two servers accessing the storage array at the same time enables high-availability, meaning if one node fails, the other node takes-over without any shortage or the user realising.

www.mongodb.com/partners/emc

GPUs:

For virtualisation purposes GPUs can be fitted in the head node

Compute nodes

The horsepower of the render farm, performance here is key, nevertheless before designing the most computational capable solution, there are a few considerations regarding the host data centre that must be taken into consideration;

- Rack Max Weight
- Rack Max Watt
- Cooling
- Potential development (planning for cluster extension)
- Cabling space and prices
- Switches

Server consideration

Performance/ density:

One should aim for best performance per server/ rack four nodes in 2U is the HPC standard, which is the sweet spot for best-performance/ density (translated into the Dell C6420).

Memory:

Perhaps the most limiting factor in terms of performance. Make sure that you have a balanced memory configuration (e.g. on 2nd Generation Intel® Xeon® processors six or 12 DIMMS).

Traditionally scenes need no more than a few 100MB to a few GB with HDR textures, therefore on 2nd Generation Intel® Xeon® with 12x32GB should be ample to get the right performance.

Choosing the best CPU

Choosing #cores and frequency

It all depends on the application;

- Licensing fees → If you pay per socket, you don't care how many cores you use and reversely. Often licences are based on VCore (virtual cores = threads) in which case a high frequency clock and turbo capacity makes most sense.
- If the application can run on one core → frequency is the most important factor
- If the application can run on multiple cores → core count is the most important factor

Flop Calculation

Finding the optimal CPU can be a challenging task, the first thing to do is using the below formula to come-up with the appropriate SKU for expected performance. The below example is based on a 2-socket server with Intel® Xeon® 6248, the result is in GigaFlop calculation per second.

$$GFLOPS = \text{sockets} * \text{cores per socket} * \text{TDP clock speed in GHz} * \text{Floating point units per core} * \frac{\text{vector width per floating point unit}}{\text{data width}} * \text{FMA capability}$$

Formula

GFLOPS =

sockets
 cores per socket *
 TDP clock speed in GHz *
 Floating point units per core * (FPUs/Core)
 (vector width per floating point unit)/(data width) *
 FMA capability

Example: Intel® Xeon® Gold 6248 Processor

GFLOPS =

- 2
 - 20
 - 2.5 GHz (Base Frequency)
 - 2 (Threads / Cores)
 - (512 (+ AVX 512)) / 64 (if Single Precision = 32) = 8
 - 2
- = 3,200**

Playing with Overclocking; Dell's BIOS allows you to select the number of cores and block the frequency level. That enables you to choose CPUs based on boost performance whilst only leveraging a fraction of the total core count.

				# of active cores / maximum core frequency in turbo mode (GHz)																															
SKU	Cores	LLC (MB)	TDP (W)	Base																															
				non-AVX Core Freq. (GHz)	1	2	3	4	5	6	S	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28		
6248	20	27.5	150	2.5	3.9	3.9	3.7	3.7	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.6	3.4	3.4	3.4	3.4	3.2	3.2	3.2	3.2										
6246	12	24.75	165	3.3	4.2	4.2	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1	4.1																		

Example: selecting the Intel® Xeon® Gold 6248 Processor and running it in turbo mode @3.6GHz with 12 cores enabled. Advantages include lower TDP, improved/ price performance, and most interestingly using one CPU type for different applications whilst only tweaking active cores.

Vectorising: Intel® Advanced Vector Extensions 512 (Intel® AVX-512) can provide additional performance if the code has been optimised.



C6420 Two-Socket Server Node (Intel)

PowerEdge C-Series

A modular server node provides a dense and balanced computing platform optimized for high-performance environments.

www.dell.com/en-us/work/shop/dell-powerededge-servers/sc/servers/modular-infrastructure?



2-Socket (Supports up to 2 Processors)

C4140 1RU High-Density Server (Intel)

R740 2RU Server (Intel)

R640 1RU Server (Intel)

R740xd 2RU Server (Intel)

R740xd2 2RU Server (Intel)

www.dell.com/en-us/work/shop/dell-powerededge-servers/sc/servers/poweredge-rack-servers

Four-Socket Rack Server	Memory
PowerEdge R840	48 (15.36TB)
PowerEdge R940	48 (15.36TB)
PowerEdge R940xa	48 (15.36TB)

www.delltechnologies.com/nl-nl/servers/poweredge-rack-servers.htm#accordion0&accordion1&accordion2

Extreme Density Server

PowerEdge FX Chassis

The Dell EMC PowerEdge FX converged architecture provides a flexible, modular platform that can be easily customized to match the requirements of specific data centre when density is an issue.

FX2/FC640 Provides Exceptional Workload Flexibility In 4 Servers In 2U Rack Space.

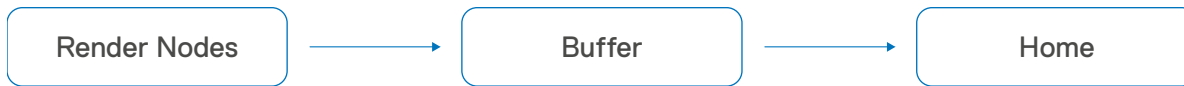
4 x FC640 equals Up to 224 Cores/6TB Memory/2xM.2 Raid1/8x2.5"/FN IO's

The FX2 is a 2U platform that combines the density and efficiencies of blades with the simplicity and cost advantages of rack-based systems. The FX2 hosts flexible blocks of server and storage resources while providing outstanding efficiencies through shared power, networking, I/O and management within the chassis itself.

[The PowerEdge FX architecture - scaling business performance without compromise](#)



Storage



The need for a Buffer? Media and Entertainment workloads are very IOPS demanding. Therefore, creating a Buffer based on NVMEs between render nodes and Home storage to offload files that are often being accessed greatly improves performance.

When selecting your storage options, consider several points;

- Price/ performance
- Tiering capability (e.g. archiving) and Auto-tiering
- Ease of scaling-up
- Cloud capability

Option_1: Dell EMC PixStor and Excelero Solution



Home: Dell EMC PixStor Ready-Solution

PixStor is A Parallel File Systems (PFS) that provides concurrent access to a single file or a set of files from multiple nodes, very efficiently and securely by distributing data to multiple LUNs across several servers.

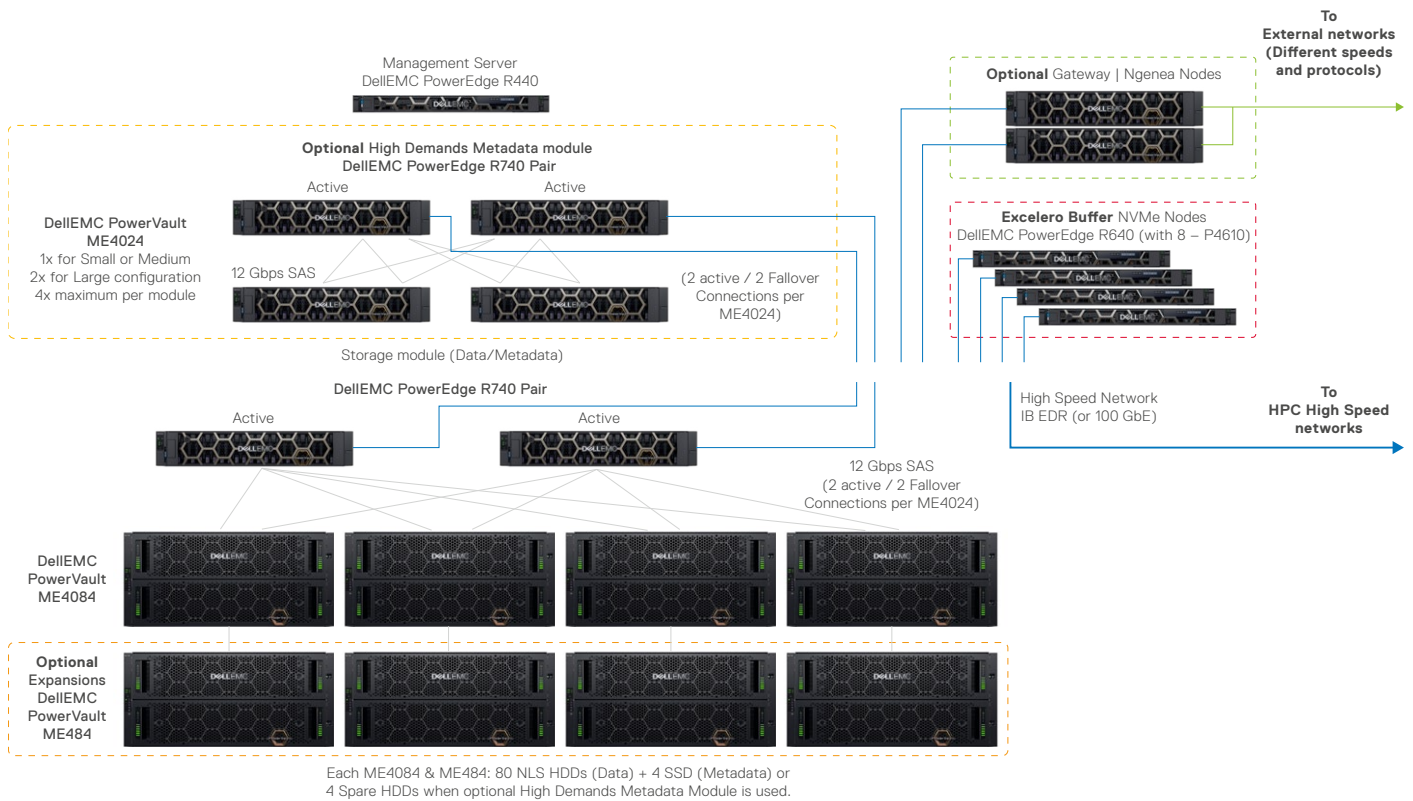
Protocols: very frequently data cannot be accessed using the native PFS clients normally used to access data, but instead other protocols like NFS or SMB must be used. Point in case is when customers require access data from workstations or laptops with MS-Windows or Apple macOS, or research/production systems that only offer connectivity via standard protocols The DellEMC Ready Solution for HPC PixStor Storage uses Gateway nodes as the component to allow such connectivity in a scalable, efficient, and reliable way.

Auto-tiering: the PixStor Solution can provide tiered access to other devices using enterprise protocols, including cloud protocols, using the Ngenea node with ArcaStream proprietary software that allows that level of integration while staying very cost effective.

Support: The solution comes supported by Arcastream and which supports and fully integrates the Excelero solution as a Buffer storage.

[Dell EMC Ready Solution for HPC PixStor Storage](#)

[NVMe in the Studio with Pixit Media PixStor and Excelero NVMesh®](#)



Buffer: Excelero

M&E applications very often hit the IOPS performance wall. To overcome this latter creating a high-speed storage Buffer is necessary. The Excelero Solution enables NVME performance for hot files that need performance.

The solution is aimed for organisations that need high IOPS whilst keeping costs down. The below solution can easily be scaled-up and is fully supported with enterprise class features.

[Read more about NVMesh architecture here](#)

Performance Optimized Platform – 1U with up to 8 NVMe Drives

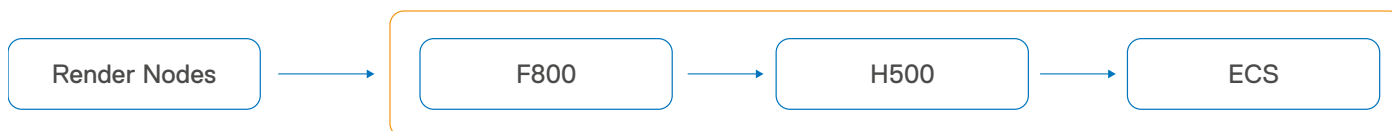
Dell EMC	PowerEdge R640
NVMe Drives	8 x 6.4 TB Dell NVMe (Samsung 1725a)
Chassis Form Factor	1U, 1 Node
Drive Density	1 Drives/RU
CPUs	2 x Intel® Xeon® Silver 4114 per node
Memory	Minimum 96 GB (12 x 8 GB) per node
Usable Capacity Mirrored	25 TB
Usable Capacity 8+2	40 TB
IOPS (4K, Mirrored)	Up to 4.8 Million Read, 720,000 Write
Throughput	Up to 23 GB/s Read, 10.1 GB/s Write
Latency (block)	<100µs Reads, 30µs writes
Network Connectivity	4 x 100GbE Ports (2 NICs)
Budgetary Estimate	
\$/TB Mirrored	
\$/TB 8+2	
\$/GB/second	



Video Resolution	FPS	Read Streams	Write Streams
4K	30	23	10
4K	60	11	5
8K	30	5	2
8K	60	2	1
16K	30	1	0

Option_2: Isilon – Home and Buffer in one

Isilon: one namespace for 3 tiers



Isilon offers a single approach to the whole data management pipeline with different tiering capabilities whilst supporting all your different protocols. Additionally, Isilon offers many other advantages;

- It is Self-managed– that means you don't need staff to run it
- Data life cycle management, Isilon applies data moving policies that you predefine
- It automatically creates one single data lake for all your unstructured data
- Isilon supports all your protocols (NFS, Hadoop, HTTP , Rest API...)
- Auto-tiering & ECS (cloud archiving): automatically switches between hot tier for high IOPS performance (F800 Flash) to higher capacity storage for long-term archiving (ECS)
- [SynclQ \(automated data replication\)](#) software that automatically replicates your data between tiers

System Configuration

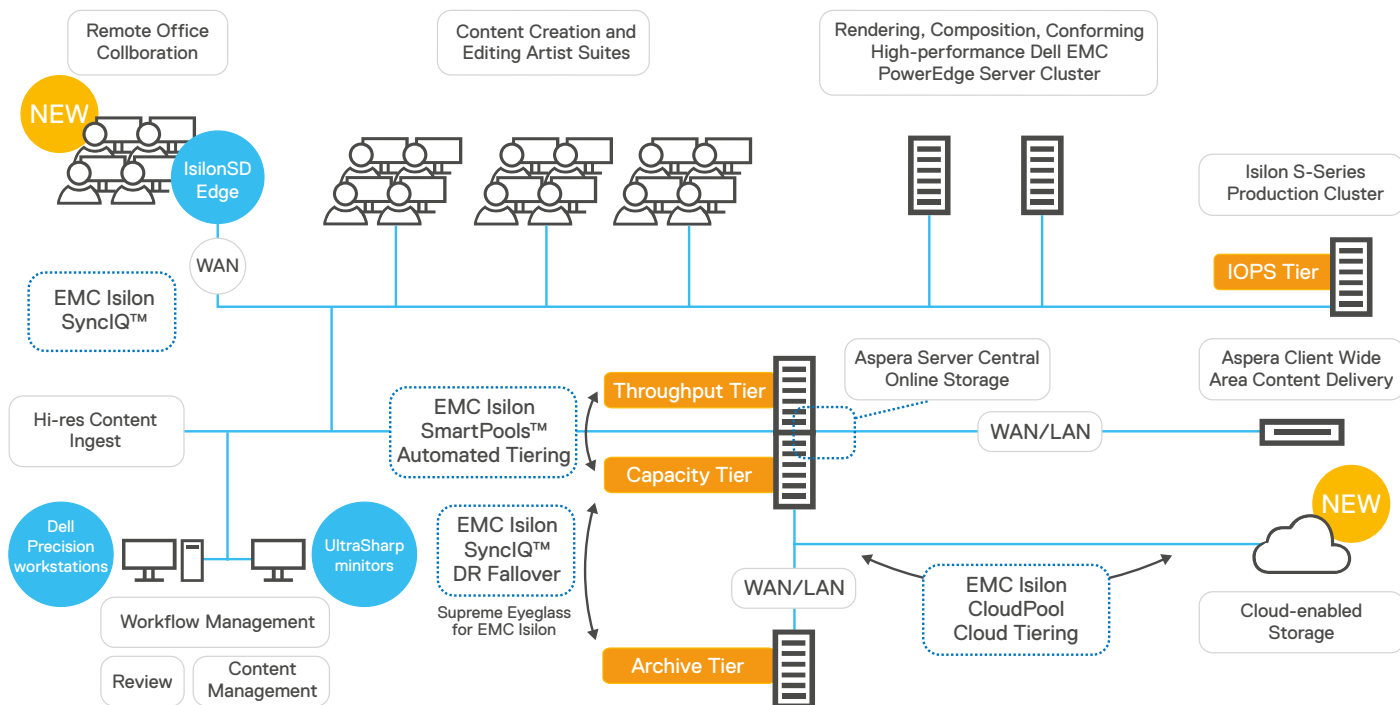
- 2x F800 (All Flash node)
- 2 x H500 (Hybrid node)
- ECS cloud archiving

www.delltechnologies.com/en-hk/solutions/media-entertainment.htm

Alternatively, by using Power Scale NVMe's on the head node, one can start with only two nodes which makes it a more approachable entry point by removing the need of using NVMe's storage arrays.

Isilon Gen 6 Family

		Performance (Per Chassis)	Capacity (Per Chassis)		
All-Flash Nodes	F810/ F800	250K ops/ss, <1ms	230 TB	924 TB	Extreme Performance & Compression
		15 GB/s read	96 TB	924 TB	
Hybrid Nodes	H600	120K ops/s	72 TB	144 TB	Performance, Capacity and Value
	H5600	12 GB/s read		800 TB	
	H500	5GB/s read	120 TB	480 TB	
	H400	3 GB/s read	120 TB	480 TB	
Archive Nodes	A200		120 TB	480 TB	Capacity & Economics
	A2000			800 TB	



Isilon Compression

OneFS in-line data reduction combines both real-time compression and deduplication. Compression uses a lossless algorithm to reduce the physical size of data when it is written to disk and decompresses the data when it is read back. More specifically, lossless compression reduces the number of bits in each file by identifying and reducing or eliminating statistical redundancy. No information is lost in lossless compression, and a file can easily be decompressed to its original form.

Deduplication differs from data compression in that it eliminates duplicate copies of repeating data. Whereas compression algorithms identify redundant data inside individual files and encode the redundant data more efficiently, deduplication inspects data and identifies sections, or even entire files, that are identical, and replaces them with a shared copy.

Both compression and deduplication are transparent to all applications that sit on top of the filesystem including protocol-based services like NFS, SMB, HDFS, S3, etc. The primary purpose of OneFS in-line data reduction is to reduce the storage requirements for data, resulting in a smaller storage footprint, reduced power and cooling requirements, and a reduction in the overall per-TB storage cost.

Additionally, in-line data reduction also helps to shrink the total amount of physical data written to storage devices. This is particularly beneficial for solid state drives (SSDs) and other media with finite overwrite limits, by significantly reducing flash drive wear rates.

The OneFS In-line Compression - Dell Technologies

PowerScale Dedupe

Accelerate demanding file workloads with extreme performance and capacity. Inline data reduction and all-flash NVMe provide greater efficiency all in 1 U.

PowerScale F200
PowerScale F600

[Dell EMC PowerScale | Dell Technologies](#)

PowerScale SmartDedupe is a native data reduction capability which enables enterprises to reduce storage costs and footprint and increase data efficiency, without sacrificing data protection or management simplicity.

Dell EMC PowerScale SmartDedupe maximizes the storage efficiency of a cluster by decreasing the amount of physical storage required to house an organisation's data. Efficiency is achieved by scanning the on-disk data for identical blocks and then eliminating the duplicates. This approach is commonly referred to as post-process, or asynchronous, deduplication.

After duplicate blocks are discovered, SmartDedupe moves a single copy of those blocks to a special set of files known as shadow stores. During this process, duplicate blocks are removed from the actual files and replaced with pointers to the shadow stores.

With post-process deduplication, new data is first stored on the storage device and then a subsequent process analyses the data looking for commonality. This means that the initial file-write or modify performance is not impacted, since no additional computation is required in the write path.

OneFS SmartDedupe - Dell Technologies

Option_3: Dell EMC Ready Solutions for HPC NFS Storage

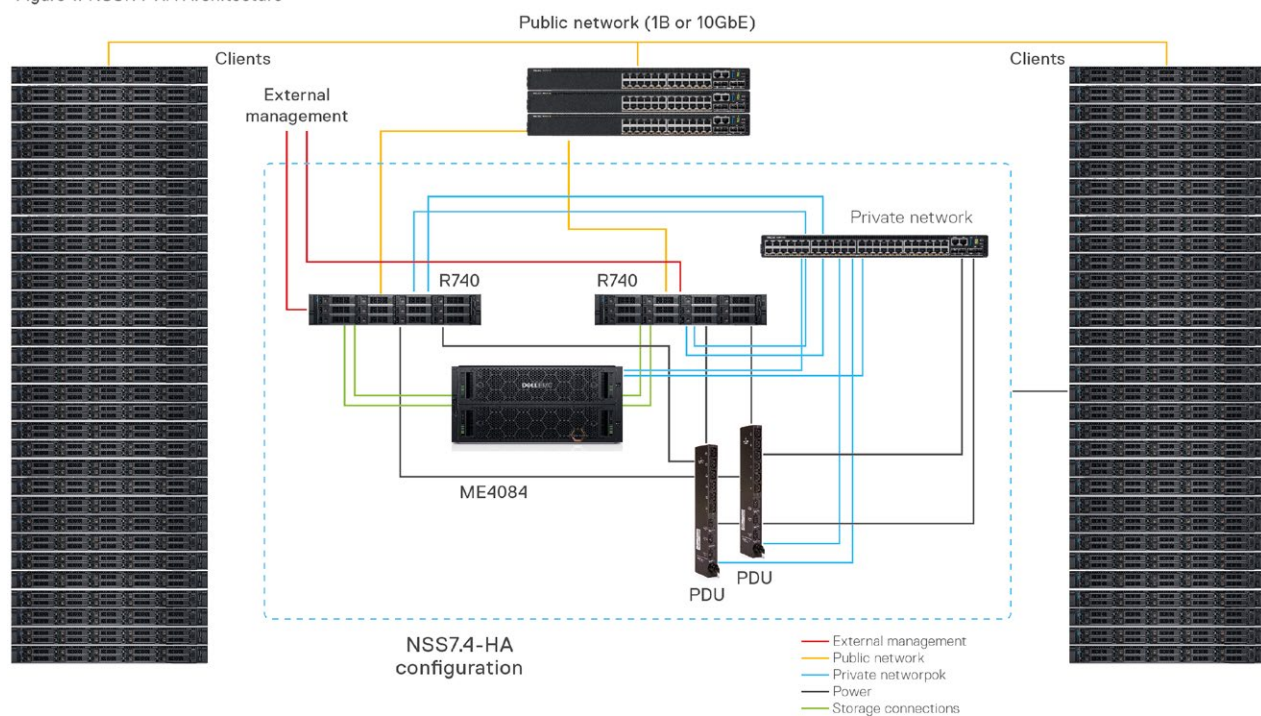


Budget friendly option for NFS storage without the features, this has the added benefits to support higher-bandwidth ethernet network of 100 Gbe and higher.

www.dellemc.com/resources/en-us/asset/briefs-handouts/products/ready-solutions/dell_nfs_storage_solution_brief.pdf

File servers	Storage subsystem	Operational storage	System networking	Operating system
PowerEdge R740 Server with dual Intel® Xeon® 6240 processors	PowerVault ME4084 with support for up to 1,008TB raw space	PowerEdge R740xd Server with dual Intel® Xeon® 4110 processors	<ul style="list-style-type: none"> PowerSwitch S3048-ON (internal) Mellanox ConnectX®-5 InfiniBand EDR/100 GbE and 10 GbE (external) 	Red Hat Enterprise Linux with Red Hat Scalable File System and Red Hat High Availability Add-On

Figure 1: NSS7.4-HA Architecture



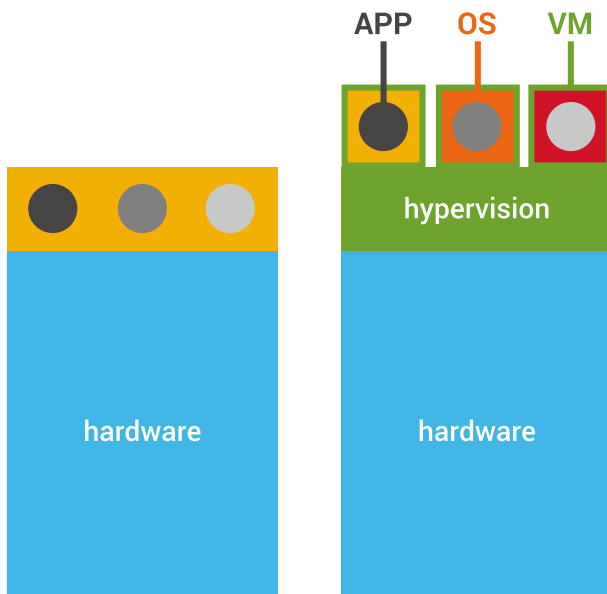
www.dell.com/support/article/en-us/sln317325/dell-emc-hpc-nfs-storage-solution-high-availability-nss7-4-ha-configuration?lang=en

Cloud and virtualised computing

Dell Virtualisation solution

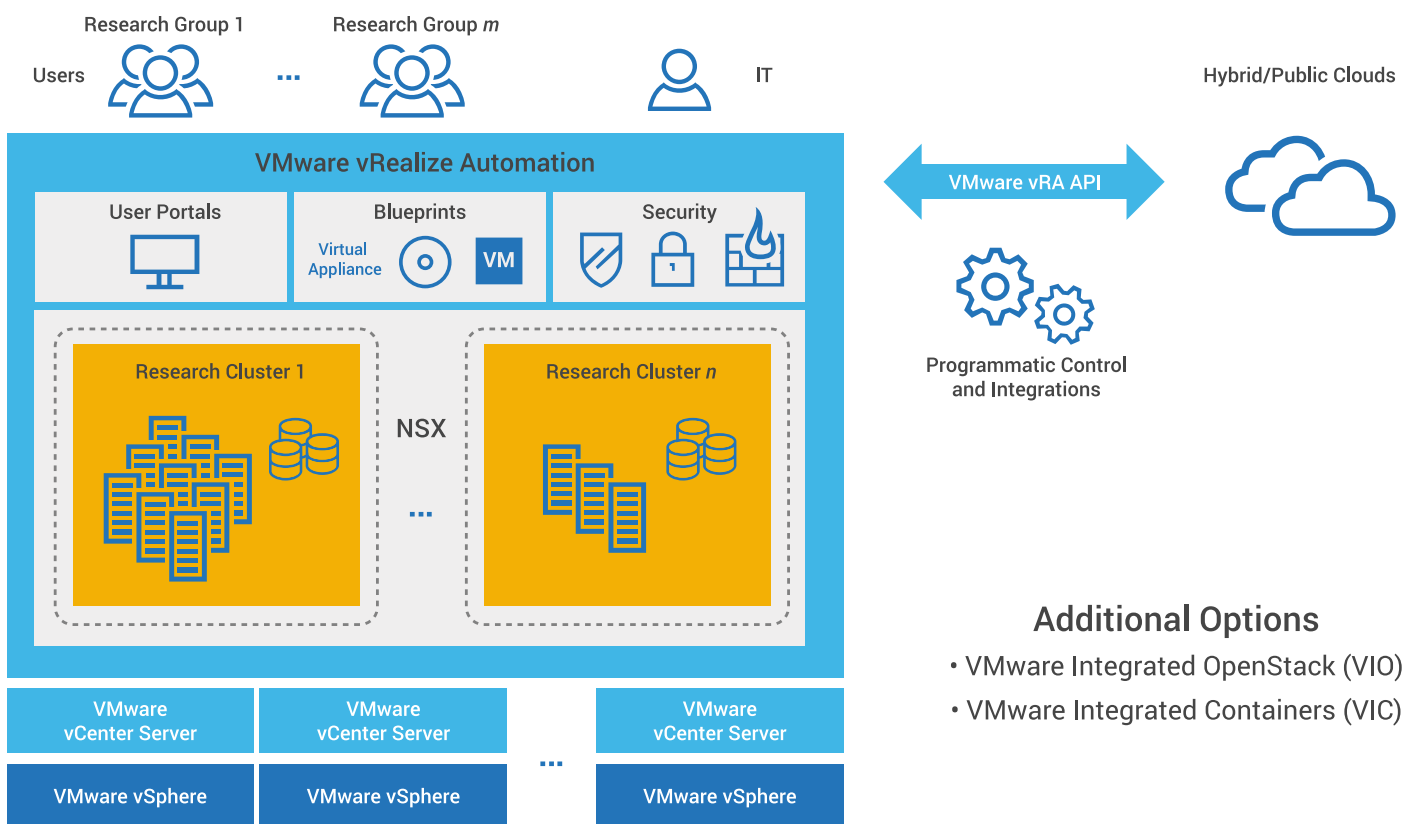
Virtualising render workloads with [VMware vSphere Scale-Out](#) gives you the ability to create a centralised compute service that simultaneously supports teams with disparate software and compute requirements, including different operating systems. Virtualising means a cost reduction and increased efficiency via more effective resource pooling and increased security multi-tenancy to ensure that while hardware may be shared to increase efficiency, sensitive data will not. Finally, VMware enable a multi-cloud strategy with [VMware Cloud Management Platform](#) for a best-practices hybrid multi-cloud environment.

VMware



Virtual machines offer:

- Heterogeneity
- Multi-tenant data security
- Fault isolation
- Reproducibility
- Fault resiliency
- Dynamic load balancing
- Performance



Additional Options

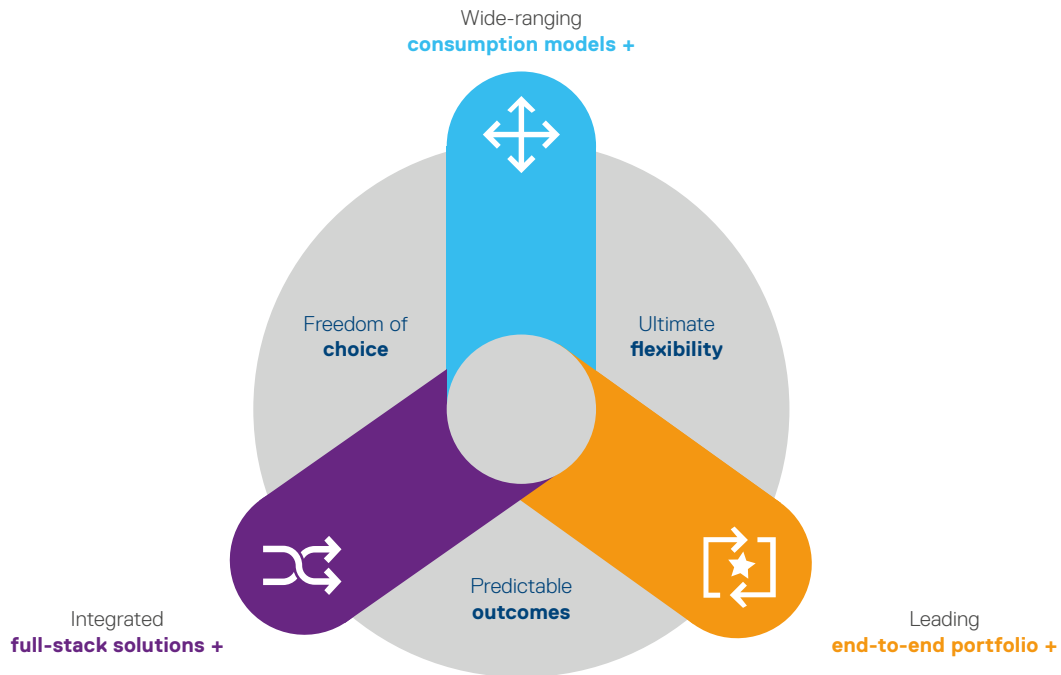
- VMware Integrated OpenStack (VIO)
- VMware Integrated Containers (VIC)

www.vmware.com/nl/solutions/high-performance-computing.html

www.vmware.com/nl/solutions/virtualization/cloud-management-platform.html

Dell Technologies on-demand

Dell Technologies On-Demand offers flexible consumption payment solutions that enable customers to pay for technology as it is used (or consumed). The featured pay-per-use model is known as **Flex on Demand**, which is well-suited for rack-level solutions. This usage-based model utilises a metering tool to measure processing, memory and storage utilisation so that payments are aligned with actual usage.



www.delltechnologies.com/en-us/solutions/dell-technologies-on-demand.htm#accordion0

Dell Technologies offers multi-cloud choices

Brochure—AI, HPDA and HPC Hybrid Cloud

Dell Technologies provides flexibility for your hybrid cloud HPC, AI and HPDA deployments. Choose where to deploy your clouds and how you'd like to pay for them, then add deployment and management options to craft a solution that's just right for you without vendor lockdown.

Mix and match the options below to create a hybrid cloud HPC environment that best suits your needs			
	Purchase options	Deployment options	Management options
Private on- or off-premises cloud	Buy or lease hybrid cloud solutions using Dell Technologies Financial Services	Use internal resources or Dell Deployment services for on-premises. Use Dell or hostingprovider deployment services for off-premises	Use internal resources or outsource with Dell or partner managed services
Public cloud	Pay-per-use for HPC resources (IaaS or PaaS)	No deployment required	Managed by the cloud providers
Hosted and managed public or private cloud	Consume hosted and managed HPC services in a PaaS/HPCaaS model	No deployment required for PaaS/HPCaaS	Managed by the cloud providers
Public/private cloud partnerships	Pay-per-use for supercomputing resources offered by academic institutions	Typically, no deployment required	Managed by the organization or institution

The Dell Technologies advantage

Increase agility

With multiple choices for purchase, consumption, deployment and managed services allowing you to pick and choose the right set of capabilities for your needs.

Reduce complexity

Focus on using HPC, not deploying and managing it, by leveraging Dell Technologies and partner deployment assistance and managed services or public cloud alternatives. Dell Technologies can help you protect intellectual property and maintain compliance with an on- or off-premises IT resources.

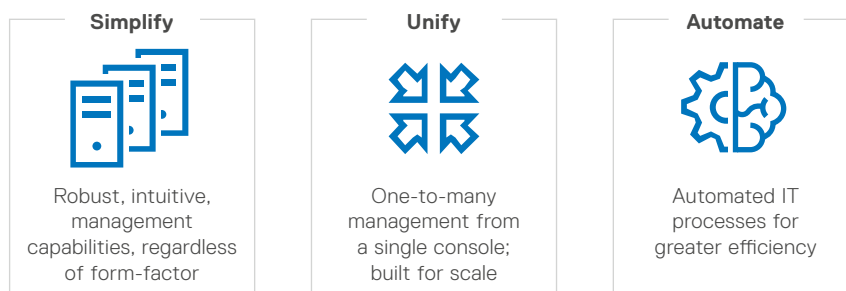
Grow affordably

Enforce policy-driven governance of HPC resources to avoid “sticker shock” from lines of business tapping public cloud HPC on an ad hoc basis. Increase cost efficiency by partnering with hosting providers that have lower energy and facilities costs. Enjoy flexible payment models that may allow you to shift some capital expenses (CapEx) to operational expenses (OpEx).

Services & software

OpenManage

Dell EMC OpenManage Enterprise is an intuitive infrastructure management console. It is designed to take the complexity out of IT infrastructure management. It delivers better results with less time and fewer steps.



iDRAC

The iDRAC also provides both a web interface and command line interface that allows administrators to perform remote management tasks. It has many advantages, such as;

- Set RAID, Access SupportAssist and Check firmware remotely
- Security:
 - Factory generated password
 - System Lockdown
 - SSL Encryption
 - Dual Root of Trust
 - User Permissions
- Virtual Console
 - Click directly into OS remotely rather than having to be at server or KMM
 - Use iSM to allow iDRAC and OS to communicate information back and forth
 - Extends SupportAssist functionality to include OS & application data

Get a free demo to understand the many advantages of iDrac Democenter.Dell.com

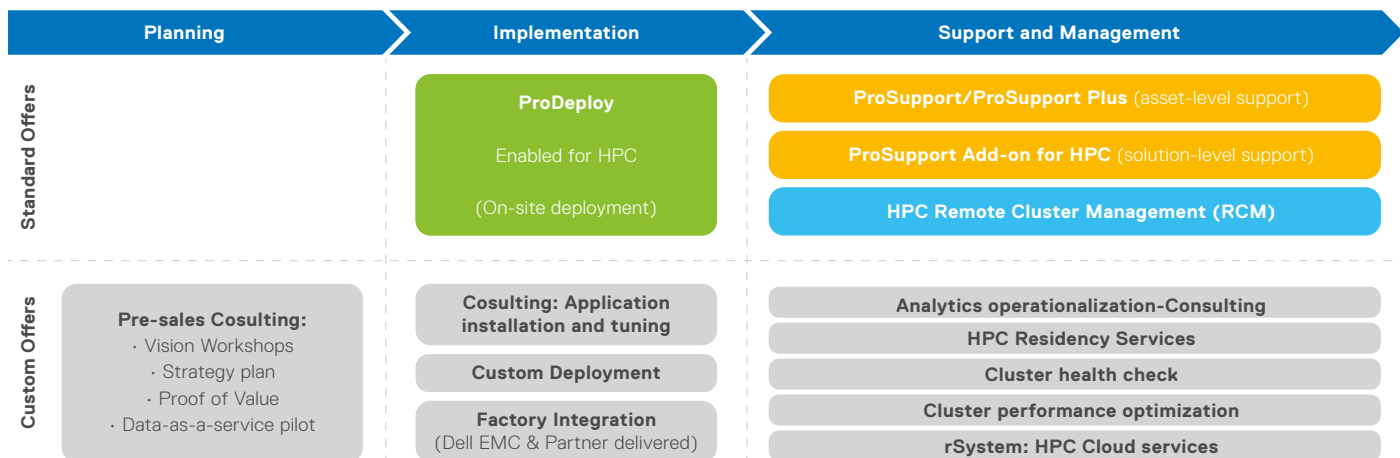
Dell Technologies Services and Security

Make Dell Technologies Services your transformation partner for the digital future. From the edge, to the core, to the cloud—our industry experts offer strategic guidance and proven practical capabilities to help you accelerate time to value of your transformation objectives. www.delltechnologies.com/nl-nl/services/index.htm

Dell Endpoint Security. Trusted Devices. A comprehensive approach to endpoint security, for the industry's most secure solutions www.dell.com/ms-my/work/learn/software-security-data-security

From design and implementation to support and systems management, Dell EMC offers a comprehensive services portfolio for render farms/ High Performance Computing (HPC) clusters, including on-premise and managed systems, as well as those in the cloud. www.delltechnologies.com/en-us/solutions/high-performance-computing/services/index.htm

Dell EMC HPC Services Portfolio



Cluster Manager

Dell EMC supports the OpenHPC Collaborative Project, and is working hard to provide a new, open source framework to support the world’s most sophisticated HPC environments. Visit openhpc.community to learn more.

Bright Cluster Manager provides infrastructure management technology for Dell EMC’s HPC Systems to enable small and medium-sized enterprises in accelerating their science, engineering and analytics, and is included in the Dell EMC HPC System for Life Sciences, the Dell EMC HPC System for Manufacturing, and the Dell EMC System for Research. www.brightcomputing.com/technology-partners/dell

	Bright Cluster Manager*		OpenHPC	
Cluster Management				
Provisioning	Bright Cluster Manager		Warewulf / xCAT	
Monitoring	Bright View	Bright Cluster Manager (cmsh)	Nagios	Ganglia
Management			Clustershell, mrsh, conman	
Operating System	RHEL/CentOS 7 RHEL 8	SLES 12 SLES 15	RHEL 7 / CentOS 7 / SLES 12	
Platform	OpenManage, IDRAC, IPMI			

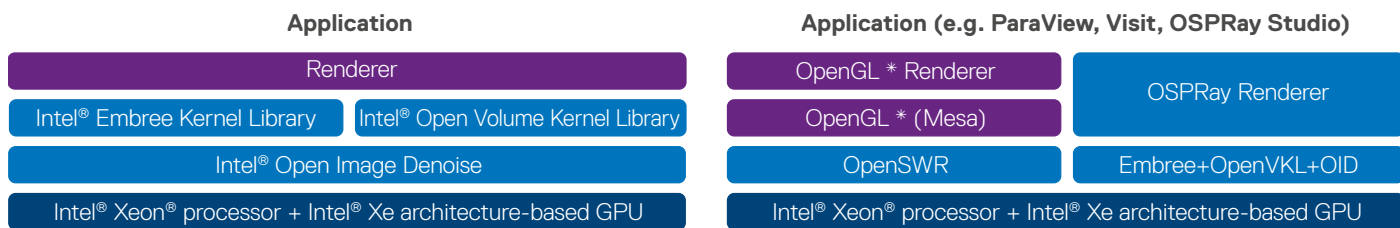
Intel® oneAPI Rendering Toolkit

Open Source Libraries Developed by Intel® that are used in leading applications for professional rendering & visualisation, optimised for parallel processing architectures.

It provides an access to all system memory space for largest data sets and improved visual fidelity via ray tracing with global illumination. It is Cost efficient and provides interactive performance for any size data.

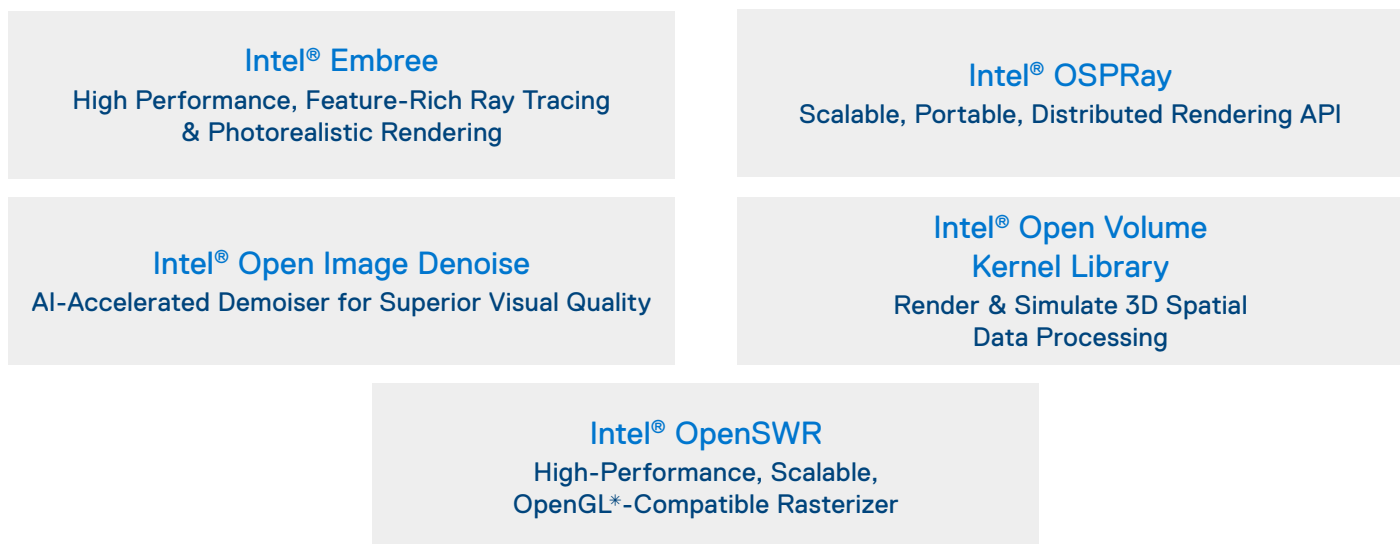
software.intel.com/content/www/us/en/develop/tools/oneapi/rendering-toolkit.html

OneAPI Rendering Toolkit




Using optimized libraries to enable high-performance, high-fidelity application on Intel® platform

Intel® oneAPI Rendering Toolkit



Intel® oneAPI Rendering Toolkit (coming soon)



Decision factors when choosing a technology partner.



Choosing the right technology partner that compliments not only a studio's technology needs but also its strategic goals is very important. Studios are looking for technology vendors that listen to the industry and take the time to understand the challenges. They need to work alongside studios to solve current challenges as well as continue to innovate for future challenges. Trust in the people, the hardware, the software, security and the support expertise are just as vital.

Why Dell Technologies?

Why Dell Technologies has an "unfair advantage" over the competition

- 1 Coherent vertical focus
- 2 End-to-end portfolio
- 3 Our People focus on what matters – Relationship, Trive, Trust
- 4 Success in M&E is a consequence of the adoption of a "Collaboratio First" strategy

IV. The Future, Trends & Challenges

Machine learning in areas such as;

- Motion capturing
- Deep Learned image denoising
- Set / Character building via Script analysis
- Image Upscaling instead of rendering high resolution
- Motion blur processing

Cloud render farms: Given the barrier of entry to huge amounts of compute power is getting easier and cheaper via cloud rendering companies and vendors. Cloud rendering in the future will allow independent artists and studios access to compute power at a large scale enabling for increases in rendering throughput and quality of work.

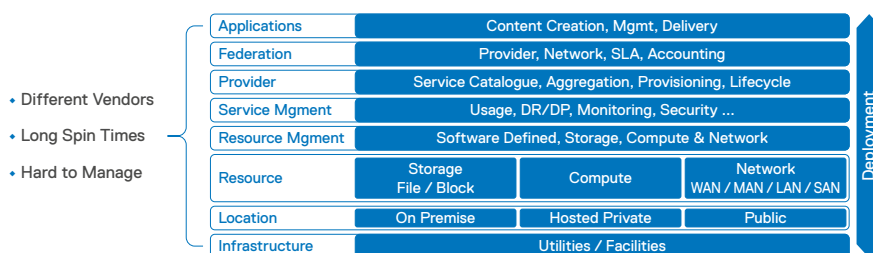
According to a publication by movie labs *"The Evolution of Media Creation – A 10-Year Vision for the Future of Media Production, Post and Creative Technologies"* The following foundational principles the top 10 future trends;

1. All assets are created or ingested straight into the cloud and do not need to be moved
2. Applications come to the media via a high-bandwidth, low-latency, cloud-enabled production world where these files do not move, and software tools come to the content instead of the other way around
3. Propagation and distribution of assets is a "publish" function
4. Archives are deep libraries with access policies matching speed, availability and security to the economics of the cloud
5. Preservation of digital assets includes the future means to access and edit them
6. Every individual on a project is identified and verified, and their access permissions are efficiently and consistently managed
7. All media creation happens in a highly secure environment that adapts rapidly to changing threats
8. Individual media elements are referenced, accessed, tracked and interrelated using a universal linking system
9. Media workflows are non-destructive and dynamically created using common interfaces, underlying data formats and metadata
10. Workflows are designed around real-time iteration and feedback

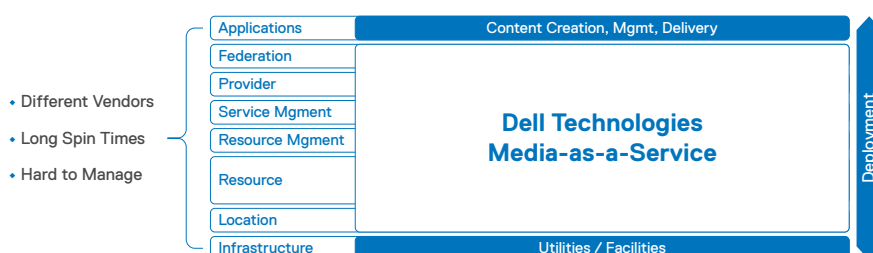
Collaboration; with emphasis on media creation workflows, tools, infrastructure and techniques that make collaboration easier and more efficient, such as cloud storage.

Automation of mundane or repetitive tasks such as provisioning services and extending infrastructure in the cloud with emphasis on eliminating technical debt.

What Studios Manage Today



What Studios will Manage in the Future



Challenges in the industry

- Content creation standards are being driven to even higher resolution (8K)
- Trends such as OTT, VR and HDR are creating all new pipelines
- Studio has to rapidly expand and contract to accommodate spike projects
- Studio's seek low or no commitment consumption models
- An inability to call on high-power and fit-for-purpose infrastructure leads to costly outsourcing
- Trying to stay ahead of the competition means that a studio may need to open new facilities, close, or merge existing ones
- Being forced to move large parts of the infrastructure and render farms to accommodate electrical power constraints
- Discerning and increasing sophistication of viewing audience requires more realistic VFX
- Supporting legacy and/or linear pipelines

Security and Best Practices

Media and Entertainment – Security considerations and difficulties.

Most M&E facilities (especially VFX facilities) will generally be catering to film studios, broadcasters or online streaming services.

Each of these has their ideas on what a secure workflow looks like, but in recent years, they have all come together to try and formulate one standard.

The [CDSA \(Content Delivery & Security Association\)](#) and the [MPA \(Motion Picture Association\)](#) have banded together to form the [Trusted Partner Network \(TPN\)](#).

These associations contain the majority of the largest “content owners” in one group.

These include:

- Walt Disney Studios Motion Pictures
- Netflix Studios, LLC
- Paramount Pictures Corporation
- Sony Pictures Entertainment Inc.
- Universal City Studios LLC
- Warner Bros. Entertainment Inc.
- BBC Studios
- Amazon Studios
- Lionsgate
- CBS

The main idea of the TPN is to make it easier for facilities to ensure they are adhering to industry security standards, whilst also limiting the need for continued security audits from each of their customers.

Previously, if you worked on a show for Warner Bros. they could come and audit you yearly. If you then did a Disney show, then they would come and audit you separately.

Audits (and preparing for audits) can be very disruptive to some facilities. Especially those with smaller IT teams. The TPN is attempting to simplify this by creating one standard that all studios etc agree to, which would mean only one audit per year.

Wrangling all these content owners to agree to these standards appears to be tricky as the TPN has only slowly gained momentum over the last two years, however, most studios are no longer conducting their facility security audits and relying on vendors to do them via the TPN.

These standards all revolve around one single purpose, and that is to protect the content of the films or shows that are being worked on.

Data flows need to be controlled and audited:

- Where is content data stored and how does it get there?
 - Create a separate IO (In, Out) network for use for ingesting media only
 - Lockdown internet access to only whitelisted transfer mechanisms (e.g. Aspera)
 - Lockdown logins on the systems connected to this network to only approved users
 - Audit trails for login and transfers
 - Have a content storage system that is separate from business or corporate data
- How is content data sent back to clients?
 - By utilising the IO network and whitelisted mechanisms as dictated by clients
- What is in place to ensure that content data can't be exfiltrated?
 - Separated networks for content/production use - Firewalls
 - Heavily restricted access to this network
 - No general internet access in this network
 - Logins restricted to this network
 - All data that comes in or out of this network must be via the aforementioned IO network

Other security guidelines are all standard security best practices that most companies would have to honour.

At the time of writing, there is still no published security standard for Cloud workflows from the MPA.

They have promised to publish guidelines specific to cloud use, but until that is formalised all facilities can do is apply some basic best practice principles and/or simply document their proposed workflows and the business cases for them.

Generally, if some attempt has been made to secure a workflow (a "compensating control") and this can be demonstrated during a security audit, this should be reflected in the resulting TPN security audit report.

The [MPA Common Guidelines](#) assume no cloud services are in use (and these guidelines form the majority of the TPN audit), and the CDSA only touches briefly on them in a small section of their [guidelines](#).

This makes it difficult for facilities to adopt more modern workflows, although some are forging ahead.

By way of example of highlighting how these "Standards" are behind the times, the most up to date [MPA guidelines](#) expect that any network that a facility is using to create "content" must not have any access to the internet! (*Ref DS-2.0 - Internet*)

In the real world, however, DS2.0 refers only to "direct internet access" and if there is a business case for a facility to connect to a cloud/internet service to do their work (a common example is licensing Adobe software via Adobe Creative Cloud), as long as this is via a proxy and they can demonstrate that the risk to losing content is minimal, this can be deemed ok.

General internet will need to be restricted in these content networks, however, and some basic productivity apps (O365, GSuite) that facilities rely on can get caught up in this "General internet" guideline as DS2.0 also directly mentions prohibiting "email".

This raises several questions:

- How does one access email?
- What about instant messaging?
- How can you work as a team without some sort of connectivity to productivity apps?
- Are they expected to maintain a slew of on-premise services?
 - On-prem email servers
 - On-prem chat service
 - What does that mean for their IT workforce?

The only way to adhere to this guideline (and they are only "guidelines") is to implement some sort of web isolation (like [Broadcom's Web Isolation](#) product) or spend thousands on VDI (VXRail/VMware), remote desktop services or simply second machines (laptops) that employees use to access the internet. This just simply isn't viable, especially for smaller facilities who are unlikely to have an IT team that can implement and maintain these systems.

Another example (particularly relevant in our current COVID WFH environment) is remote access to these facilities. (DS3.2 – Restrict Remote Access)

How can facilities work from home if there are heavy restrictions on remote access?

It certainly isn't an option to have the content be saved offsite locally on an artist's workstation in their home, not to mention this is very inefficient and would breach any TPN security obligations.

Going via bastion host (the recommendation) simply isn't an option for certain Craft Groups. Connecting via a bastion host won't support Wacom pressure-sensitive tablet which is a staple for artists like these.

Productivity will suffer!

Hopefully, this will be a temporary issue, but if there's a "new wave" of working from home requests, it will continue indefinitely.

Netflix is rumoured to be building a VFX studio entirely in the cloud, so WfH and WfA workflows are here to stay.

www.awn.com/news/arch-platform-technologies-launches-powerful-cloud-based-vfx-platform

Facilities can either ignore these guidelines and risk non-compliance or forge ahead with expensive solutions (VPN, Teradici, VDI) to try and remain secure and productive and rely on documented business cases to help their TPN audit results.

These problems have solutions, but these solutions aren't always viable for all facilities, especially smaller facilities.

VFX work is cyclic. Most people working in these facilities are contractors who are only there during the period of a show being worked on.

This is tricky to manage as facilities can't always commit to yearly (or even half-yearly) commitments to infrastructure that will only be utilised at peak during the life of a show.

The above also assumes that facilities have the appropriate controls in place to create isolated networks. (Another suggested requirement – *DS3.0 – Isolate content/production networks from non-content/production networks*).

Powerful firewalls and/or switches to be able to segment networks and control and audit data coming in and out of them are not cheap and require specialists to maintain.

Surely, they can use cloud

Is a workforce full of contractors, being able to spin services up and down as you need them all scream "cloud" right?

Well, how do you do that if you're being told you can't have access to the internet from your content networks?

It's not strictly true that you can't use cloud services. The TPN certainly has a cloud on their mind and they do offer a certification to cloud service providers (which is effectively just a TPN audit like any other facility, just of the data centres the cloud providers use). The "Best Practice" documentation just hasn't caught up yet. The industry has been promised cloud guidelines for over two years now. We're still waiting.

For those providers that DO have a TPN certification (AWS, Azure and others) it means that VFX facilities can safely partner with these providers and use their services. That's not to say there still needs to be added security placed on top, however.

The guidelines (that we're still waiting on) will likely have some of the following recommendations:

- Avoid public internet
 - Make use of direct connect to known data centres
 - If public internet is required, adopt in-line encryption
- Adopt a CASB
 - For SaaS services, look at utilising CASB's like Bitglass or Netskope
- MFA and IdP will be a must
 - Services like Okta will likely be a hard requirement

Media labs have released a white paper called the “Evolution of production security” that tries to bring to light and open up the discussion around the challenges production studios are facing when it comes to securing studios and allowing collaboration they list six primary security principles, which are the foundation of the security architecture as

1. Security is intrinsic to every component of every workflow and does not inhibit the creative process
2. The security architecture addresses requirements specific to cloud-based workflows.
3. Production workflows, processes, and assets are kept secure, even on untrusted infrastructure.
4. Content owners control security and workflow integrity.
5. Security can be scaled to appropriate levels and can integrate with existing security policy and management systems.
6. The security architecture limits the spread of any breach and adapts to the evolving threat and response landscape.

Until we get the official TPN/MPA cloud guidelines, facilities can only attempt to adhere to the best practice guidelines and document any workflows they implement that differ. A studio can use cloud services as long as the content owner is happy with the controls in place.

V. Glossary and References

Glossary

Chunk - Chunking jobs is highly beneficial when application start-up and scene load times are significant compared to processing times. Chunking loads the scene once and renders many frames saving time on application start-up and scene load.

Client/worker/slave – A node designated to be managed by the supervisor for executing a package of work.

Dependencies – Jobs that require other jobs to start or finish before they can start processing are dependency constraint.

Render job/sub-job/task/instance – Is an encapsulated package of commands to be executed by the worker. A Job can contain one sub-job with one task or a job can contain many sub-jobs with many tasks. A job with many sub-jobs is sometimes referred to as a chunk.

Render wranglers – An individual or department that is responsible for the render farm in its entirety and is the first level of render support for Artists. Ensuring render jobs are scheduled correctly, failed to render jobs are debugged, faulty nodes are dealt with and several other various tasks needed in monitoring and maintenance of the render farm.

Server/supervisor/master – A central server that is responsible for the submission, dispatch, monitoring and coordination of jobs.

Reference

www.researchgate.net/publication/262404570_A_Distributed_Render_Farm_System_for_Animation_Production
Yao, Jiali & Pan, Zhigeng & Zhang, Hongxin. (2009). A Distributed Render Farm System for Animation Production. 5709. 264-269. 10.1007/978-3-642-04052-8_31.

Quote source www.indiewire.com/2010/11/harry-potter-producer-david-heyman-talks-deathly-hallows-tentpoles-gravity-238059/

filmora.wondershare.com/video-editing-tips/best-special-effects-software.html

Renee Dunlop (2014) "Production Pipeline Fundamentals for Film and Games" CRC Press

studio3dm.com/2018/12/10/3d-animation-production-the-step-by-step-guide/

vfxplatform.com/

"The Evolution of Media Creation – A 10-Year Vision for the Future of Media Production, Post and Creative Technologies"

3D Modeling

Sidefx Houdini - www.sidefx.com/products/houdini/

Autodesk Maya - www.autodesk.com/products/maya/overview

Autodesk 3Ds Max - www.autodesk.com.au/products/3ds-max/overview

Cinema 4D - www.maxon.net/en-us/products/cinema-4d/overview/

Blender - www.blender.org/features/

Lighting

Katana - www.foundry.com/products/katana

Maya - www.autodesk.com/products/maya/overview?support=ADVANCED

Compositing

Nuke - www.foundry.com/products/nuke

Fusion - www.autodesk.com.au/products/fusion-360/overview

Adobe After Effects - www.adobe.com/au/products/aftereffects.html

HitFilm Pro - fxhome.com/hitfilm-pro

Rendering Engines

3Delight - www.3delight.com/

Arnold - www.solidangle.com

Maxwell Render - www.nextlimit.com/maxwell/

RenderMan - renderman.pixar.com/product

Redshift - redshift3d.com

V-Ray - www.chaosgroup.com/

Customer stories

Mavel Studio

Dell Brings Tech to New Heights with Marvel Studios' "Ant-Man and The Wasp" this Summer

Soho VFX

Soho VFX relies on Dell technology to create visual effects

Dell EMC Servers Boost Visual Effects Studio's Productivity by 250%

Technicolor

Dell Technologies assisted Technicolor in a digital transformation to develop high-end production using private cloud virtualisation and workforce transformation.

Cinesite

Cinesite stays agile with cloud-hosted render farm and storage, plus powerful workstations, meaning artists can create award-winning scenes four times faster

Toonz Animation

Successful studio's rapid growth gains momentum with high-performance storage

Important Looking Pirates

Important Looking Pirates creates stunning, breakthrough visual effects and locks in workflow and productivity gains with fire-powered, scalable IT solution

Animal Logic

Award-winning digital animation and VFX studio builds on the success of The LEGO Movie with Dell EMC Isilon

Toei Animation

Japan's pioneering animation company draws in a new era of film with high-performance storage

Powerhouse Animation

Powerhouse Animation improves animation quality and drives corporate growth by speeding rendering processes by 40% and giving artists better digital tools with Dell and Intel®

Huevocartoon

To bring next-generation animation and aspiring, young animators together, Huevocartoon relies on scalable, high-performance technology solutions from Dell EMC

TV Asahi

Dell EMC supports computer-generated effects for 4K broadcasting

Silver Trak Digital

Silver Trak Digital delivers 4K and other media faster while cutting costs and risk, and speeds staff efficiency by 40 percent with a seamless Dell EMC Isilon All-Flash storage platform

Studies

Arista Media and Entertainment Study



Dell Technologies Global Headquarters is located at One Dell Way, Round Rock, TX, 78682

©2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries.

Dell Corporation Limited. Registered in England. Reg. No. 02081369.

Dell House, The Boulevard, Cain Road, Bracknell, Berkshire RG12 1LF, UK. Other trademarks may be trademarks of their respective owners.

Intel® and the Intel® logo are trademarks of Intel® Inc. in the USA and other countries.