



# CLOUD-EXTENSIBLE TRANSCODING

AVOID THE TRAP OF SINGLE CLOUD OVER-DEPENDENCY

---

# TABLE OF CONTENTS

---

Transcode Workload Challenges .....	3
Cloud Dependency .....	4
Cloud-Extensibility .....	5
Elastic Transcoding .....	6
Scale-Out Transcoding .....	7
Compute Agnosticism & Awareness .....	8
Critical Or Nice To Have? .....	9
The Concurrent Solution .....	10

---

# TRANSCODE WORKLOAD CHALLENGES

---

Handling video transcoding workloads with unpredictable, rapidly evolving scale is an operational challenge in today's world of larger files and growing SVOD demand. Broadcasters need elastic resources to efficiently and quickly make original broadcast content available for on-demand consumption. Studios working with master files need large amounts of compute resource for distributing programs on-time to their content distributors. As a result, the flexibility, speed and scalability of the content transcoding function has become a critical success factor in media workflows.

But what does this mean for Media businesses that use transcoding solutions and services? What type of transcode solution is best-positioned to help Media businesses meet their objectives? This paper explores how the latest cloud-extensible transcoding solutions can contribute to faster and more cost-effective content workflows, while also avoiding the emerging "cloud-trap".

---

# CLOUD DEPENDENCY

---

Cloud-based compute seemed like a great solution to solve your transcoding problem just a few short years ago. You had an aging, hardware based transcoder well integrated in your workflow, and were experimenting with a few software applications. The hardware vendor was increasing their support costs years over year, and your procurement team was questioning the value in maintaining a platform almost a decade old. Meanwhile the software was difficult to deploy, often required dedicated computing hardware anyway, and wasn't well integrated. Cloud-based transcoding came onto the scene and solved your problems with a relatively light outlay of operational expense. As your reliance on this solution to your problem increased, new problems emerged. You realized you have a dependency and cost problem.

As businesses use cloud-hosted transcoding they are learning the limitations and the issues. What is emerging is that businesses need something slightly, but fundamentally, different.

The customer requirement is changing from "I need flexible transcoding capacity to handle my peak transcode workloads" to "I need flexible transcoding capacity to handle my peak transcode workloads that does not tie me to a single service provider's environment **and** helps me better manage cost."

This requirement has emerged due to the disconcerting over-reliance on a single service provider. Cloud-hosted services are appealing, but when all your VOD asset workflows are tied up with a single, very large, cloud service provider, a set of new business risks emerge.

The primary risk tends to be commercial, and is generally described as "cost creep." The reality is that cloud-based services often start as a cost effective solution, but change over time due to increasing scale or pricing structure modifications.

For example, today the policy might be that uploads are free and downloads / deliveries are chargeable. When the uploaded file is a large XDCAM 50mbps file, and the output is a smaller MPEG-4 10mbps files, this is attractive. What happens if the pricing policy changes to charge for uploads rather than downloads? Or, what happens when the scale of your operations increase? What impact does 4K content have?

In the subsequent risk mitigation planning, you can try to negotiate lower prices with your cloud service provider, or accept and absorb the higher costs, or move your transcoding workloads to another location. The best option would be to have no dependency on remaining with the cloud service provider so you can control your own destiny and manage your costs and margins. One technical problem about moving your transcoding platform is that media asset creation workflows can be certified by the receiving video platform. So your transcoding solution has to remain compliant. What if you've used a transcoding solution that is an intrinsic part of your cloud service - can you find this same solution elsewhere?

# CLOUD-EXTENSIBILITY

So what do you do? First you need to be able to move to a new environment that allows you to manage your costs, while also being compliant with your customer's requirements.

Over dependency on the cloud and high costs are avoidable with the right transcoding solution. Forward-looking media businesses evaluating their long-term infrastructure strategy are proactively figuring out how to achieve the best of both worlds - maximum control at optimized cost.

The concept of **cloud-extensibility** is fast emerging. We can define cloud-extensible as orchestration of compute agnostic resources between single or multiple public and/or private clouds. It's already common practice to diversify video transport workloads across different service providers such as satellite providers and CDNs. So why not video processing workloads? With software-defined transcoding solutions it's possible to blend private and public cloud environments for your transcoding workloads.

Cloud-extensibility avoids the business risk of being trapped with a single service provider. Software-defined enables you to deploy common transcoding software across any hardware infrastructure environment. A cloud-extensible solution means you can utilize any environment you wish in order to protect your business from uncontrollable changes by a single cloud service

provider. It also means that you can balance workloads, and costs, between private cloud / on-premise infrastructure and public cloud infrastructure.

To complete the solution you'll use the workflow manager of your choice to control, via API, a transcode manager that will allocate work to different transcode environments.

This all sounds fairly simple, but the business requirement is only now emerging as organizations come to terms with cloud-based transcoding costs and realize there can be a better way. The good news is that cloud-extensible solutions are available today.

## Solution characteristics

What are the key characteristics a cloud-extensible solution must have in order to make a real difference to the business's bottom line?

- **ELASTICITY** - capacity can be scaled up or down as required and, if desired, can also be paid for on a per-use basis.
- **SCALE-OUT** - extra capacity can be "bolted on" to support as many instances of the software-defined transcoder as required – no forklift upgrades, no proprietary hardware.
- **COMPUTE AGNOSTICISM AND AWARENESS** - the transcoder must be able to be run in a "general compute" environment and alongside other applications in a containerized or virtualized ecosystem. Workload can be distributed to consistent transcoding functions in any compute location.

# ELASTIC TRANSCODING

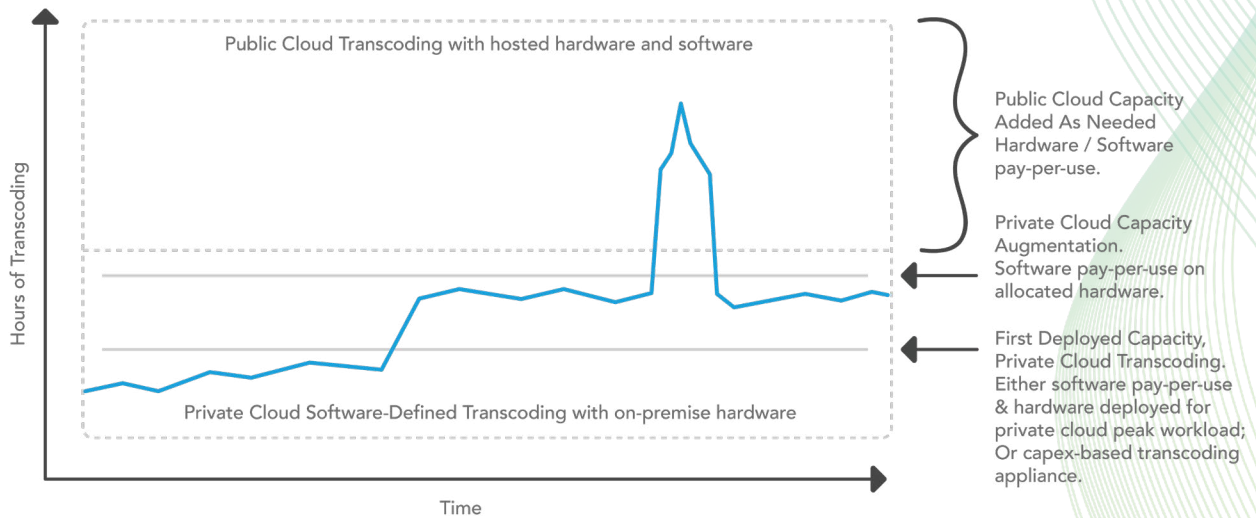
Elasticity has become synonymous with expanding capacity. But it is more than that. Elasticity really means using for and paying for capacity only when required. There are both technical and commercial considerations.

The ideal transcoding solution is a pay-per-use model. Like many other "on-demand" video processing functions, media businesses are looking for transcoding capacity when they need it rather than building out systems for peak load and then leaving capacity idle most of the time.

The best solutions therefore include software and hardware on a pay-per-use basis. Hardware is "just" compute power and is generally the lowest cost part of the solution, making it relatively simple to run a hardware platform. Software is where the differences are found in processing quality and speed when hardware platforms are equal. But software is relatively easy to instantiate on-demand. Whether transcoding is performed in a public cloud environment or on-premise, the software needs to be deployable and charged for as used. Billing engines within the transcode software communicate with the workflow manager to track usage and create invoices. This gives the most accurate method of billing, to the minute or the pixel. This also provides a level of predictability for managing what can be very lumpy operational costs.

Elasticity is a necessity in today's world of fast-changing demands and technologies. Transcoding is an ideal use case for elastic capacity.

## Elastic Transcoding - Capacity & Costs



# SCALE-OUT TRANSCODING

Scale-out refers to the ability of the video infrastructure to expand without limits. Software-defined applications, like a transcoder or storage system, have a scale-out capability that recognizes extra hardware capacity simply as somewhere to do their work. Compare this to the big iron, proprietary hardware solutions of the past.

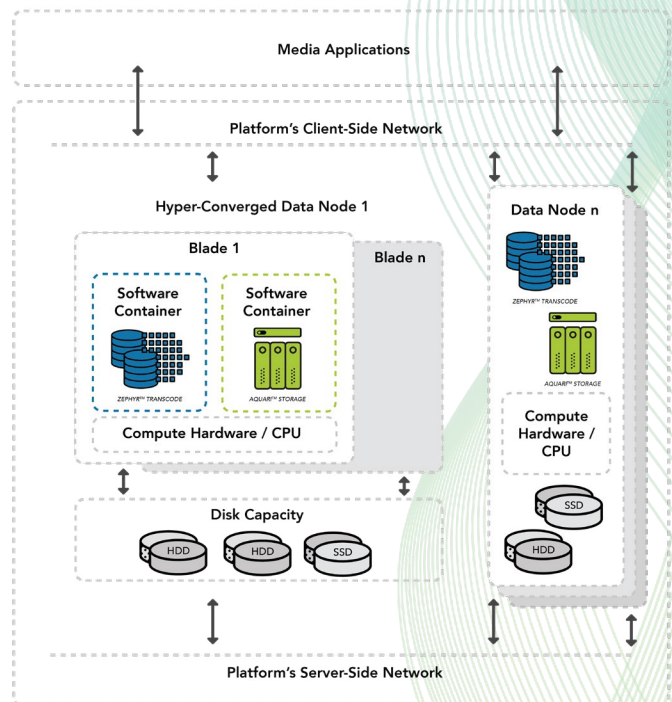
In scale-out transcoding, a Scheduler function analyses incoming jobs and allocates them to Worker nodes that perform the actual source content fetching, processing and writing out the output to its target destination. Adding extra capacity to the system involves creating additional “worker” instances on hardware with similar processing capabilities - in a public cloud environment this means allocating additional hardware. In a private cloud environment, the hardware capacity can be allocated in the same way if enough hardware is available. Once created, the new instances automatically join the pool of capacity available for the Scheduler to use.

One area of scale-out innovation from leading vendors is to use under-utilized compute resources in software-defined, distributed scale-out storage platforms to run transcoding jobs. In other words, it is possible to run storage and transcoding in a **hyper-converged platform** for maximum efficiency.

Vendors that are able to work in depth in both storage and transcode environments have a natural advantage to deliver this hyper-converged solution. In scale-out storage environments, where each additional data node adds storage, network and compute capacity, the compute resources are typically higher than the storage application requires, and so it is a natural target for optimizing on premise compute efficiencies.

This hyper-converged scale-out process is mostly intended for transcoding files rather than live streams. However, the best transcoding solutions use a blade architecture that can be expanded for file or live content simply by adding additional blades to the system.

## Hyper-Converged Storage & Transcode Platform



# COMPUTE AGNOSTICISM & AWARENESS

As compute resources become more easily accessible in cloud environments, it is possible for transcoding workloads to be flexibly moved between the best locations. The compute location-awareness and intelligence to automatically select the best location based on a trade-off between processing costs and job output deadlines can be embedded in business rules within an existing workflow manager. With a common software-defined solution available in multiple locations - including on-premise - the task of the workflow manager is to choose the best location to complete each job.

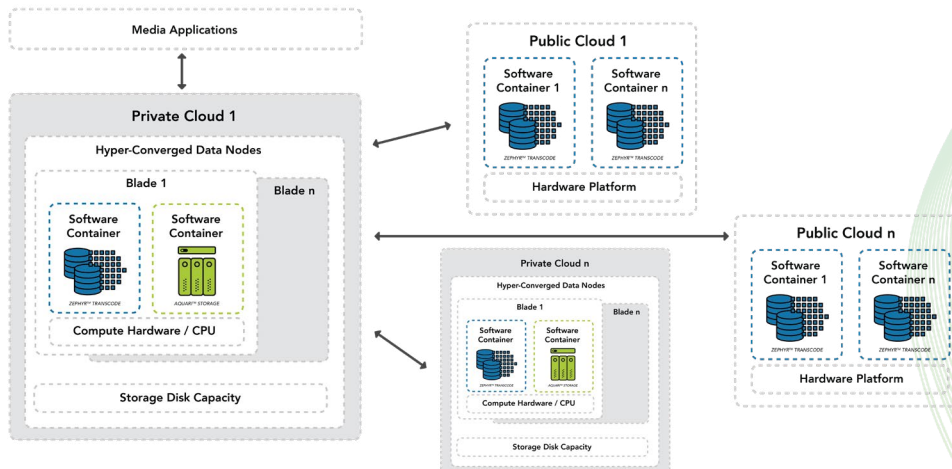
While public cloud services are well-suited to handling workload peaks because of the ease of scaling the compute resources (i.e. allocating more hardware and virtual machines), there are several downsides in terms of speed and efficiency:

1. Content uploads can take hours and days.
2. Massively virtualized environments can reduce compute efficiency.

From a workflow speed and efficiency perspective, on-premise transcoding in a general compute-based private cloud can yield best results. For the most part, commercially off the shelf (COTS) hardware is used in these environments. As internal IT departments become more adept at managing private cloud environments, it is possible to have it all: speed, horsepower, and operational efficiency.

An ideal solution therefore combines an intelligent workflow manager/scheduler, scale-out capabilities on-premise, and access to transcoding capacity with either public cloud services or private cloud services running the same software to handle any peak workload. Transcoding processes can be cloud-extensible.

## Cloud-Extensibility





---

# CRITICAL OR NICE TO HAVE?

---

Transcoding is a critical part of any content distribution workflow, to prepare content for each delivery platform and optimize distribution efficiency. Making the transcoding function as agile, efficient and cost-effective as possible is a business requirement.

Companies that have already used cloud-hosted transcoding services or cloud-hosted compute to run software-transcoders have learned some lessons about managing their dependency on, and costs with, cloud service providers.

Many of these companies have evolved their strategy to invest in solutions that are cloud-extensible. Core to that strategic approach is the employment of a pay-per-use elastic capacity model to complement their on-premise scale-out capacity. Coupled with an appropriate workflow management tool, this strategy optimizes the automated allocation of transcode jobs to the best compute locations which is critical for an efficient and predictable operational model.

Transcoding workloads are set to grow. Larger file sizes and more VOD outlets are the trend every media business should expect. Choosing the best transcoding solution to manage your costs, meet your deadlines and improve your profitability is a business-critical decision.

---

# THE CONCURRENT SOLUTION

---

Concurrent's software-defined transcoding solution, Zephyr™, is designed from the ground up to be cloud-extensible. With support for live and file workflows, deployment as Linux containers in hyper-converged environments, and easily integrated via API into any workflow manager, Zephyr can help your business manage any type or scale of transcoding workload in the most flexible and cost-effective way.

- [Zephyr Transcode Webpage & Product Brief](#)
- [MAM Storage & Transcode](#)
- [Solutions for Content Providers](#)
- [Solutions for Service Providers](#)

TO LEARN MORE ABOUT HOW CONCURRENT CAN HELP YOU MANAGE YOUR TRANSCODING WORKLOADS MOST COST-EFFECTIVELY, PLEASE VISIT [WWW.CONCURRENT.COM](http://WWW.CONCURRENT.COM) FOR MORE INFORMATION OR CALL +1.877.978.7363