THE BROADCAST BRIDGE
Connecting IT to Broadcast

Cloud Best Practices

Essential Guide

EG

ESSENTIAL GUIDES

Introduction

By Tony Orme, Editor at The Broadcast Bridge

Delivering highly resilient workflows is nothing new to broadcasters as 24/7 operations have been with us for over twenty years. However, the packet switched technology that underpins the cloud is new, and there are many additional working practices that must be considered when building cloud and hybrid infrastructures.

Packet switched networks have the potential to deliver greater efficiencies when compared to the SDI and AES circuit switched infrastructures. Essentially, the nature of packet switched networks allows designers and innovators to take advantage of the statistical nature of data distribution. That is, we rarely need to distribute full bandwidth video and audio over all the SDI and AES cables all the time. Instead, a subset of all the available distribution cables is used at any one time. Although the SDI and AES circuit switched approach guarantees bandwidth availability, it does so at the expense of flexibility and scalability.

In broadcast facilities where the video and audio are distributed in baseband, the data requirements appear to be consistent and static, but the reality is quite different. For example, removing the syncs creates data blocks that exhibit peaks and troughs. Admittedly, the data is highly cyclical but there are gaps where ethernet switches providing multiplexing operations can insert other frames of data.

Modern broadcast infrastructures no longer need to specify line, field, and frame syncs as they've technically been relegated to the history books. The matrix-based layout of monitors and image gathering devices means we only need one reference pixel per frame as all other pixels can be easily derived from this, and we can even derive the frame sample timing from the reference pixel. This methodology forms the basis of formats such as ST2110 where the essence has been abstracted from the underlying sync pulses. Alas, we still insist on using line, field, and frame syncs, but this is to maintain an element of backwards compatibility.

Compressing video and audio takes this idea one stage further as statistical peaks and troughs naturally appear in the data stream as the compression systems look to exploit similarities in the video and audio samples. Then there is the fact that not all the point-to-point circuit switched infrastructure components are used 24/7 within the broadcast facility, so there are other efficiency savings to be had.

Moving to IP and the associated packet switched networks allows broadcasters to gain greater efficiencies as the statistical peaks and troughs can be amortized over short periods of time. However, even with the best designed and optimized IP network in the land, there will always be times when viewers generate more demand for programs than the broadcaster was expecting. And this is one area where cloud systems excel as they provide additional scalability and flexibility that is hard to replicate in a datacenter. When employing the cloud, broadcasters must be aware of the working practices that need to be considered.

At the core of all technological developments is the need to keep businesses operating efficiently and reliably. Consistency of operation is fundamental in the cloud, and this leads to methods of working that may not be immediately obvious. For example, chaos engineering may seem like a contradiction in terms, but it has many benefits for broadcasters who see the need to stress their infrastructures in controlled environments so that their reliability can be further improved.

Broadcasters adopting cloud technologies, whether in isolation or as part of a hybrid solution must think carefully about the working practices they adopt. Luckily, we can learn a great deal from the working practices already established by IT.

Tony Orme.

Tony Orme
Editor, The Broadcast Bridge

Cloud Best Practices

By Tony Orme, Editor at The Broadcast Bridge

Moving to cloud computing is more than just a technical challenge, it has the potential to embrace the whole needs of the broadcaster's business. And whether a broadcaster decides to move completely to the cloud or adopt a hybrid approach, the consideration of best practices should be at the forefront of their minds.

Although cybersecurity is important with on-prem datacenters, the introduction of cloud computing has further escalated how we approach security. It's not that security has never been an issue for broadcasters as it certainly has, it's just that with traditional broadcast infrastructures media assets were physical devices that had to be obtained by breaking into a building or physically intercepting a delivery. Now, we must be more vigilant as criminals no longer need access to a physical asset stored on video tape. Consequently, we must spend much more time thinking about security and making sure the relevant systems are in place.

At the heart of cloud computing is scalability and flexibility which in turn leads to dynamic computer systems that increase and decrease the resource to meet the needs of the business. But this flexibility also has potential implications for resilience that must be addressed.

With cybersecurity and resilience at the heart of any broadcast cloud infrastructure, it soon becomes clear that the best practices adopted to meet these requirements must be built into the system at the beginning of the design, and not as an afterthought. Consequently, many of these principles are driven by the business requirements.

Business Continuity

The days where we could bury our heads in the sand and pretend we could build systems that don't fail are well and truly gone. Instead, we must take a more pragmatic approach and assume something will go wrong, and this is demonstrated in today's agile working practices which assumes things will go wrong. This should be taken one step further so that every component in each workflow should be analyzed and assumed to fail, and if it does, then have some form of remedy in place to fix it.

The same methodology applies to security. No matter how well designed a security system is, there is always going to be a small chance that a cybersecurity breach will occur, and this is especially true for broadcasters who are high profile targets for international cybercriminals. Adopting a zero-tolerance approach will significantly reduce the risk of a security breach, but the very act of allowing users access to a system will always weaken it slightly.

Core to making any broadcaster's system secure and resilient is the need to provide business continuity. That is, have a plan in place for every workflow that can fail. This also includes how a broadcaster backs up media assets as they must be careful of how the files are synchronized to the backing store. For example, if the on-prem storage is mirrored to the cloud storage and an on-prem file is infected with a virus, then, at some point it will be mirrored to the cloud storage and could possibly infect the cloud backup. A method of alleviating this is to provide incremental backups to the cloud storage so that historic copies of the files can be retrieved. This is potentially an expensive method of working as more storage will be used, but this must be weighed up against the cost of losing the asset.

Employing cloud technologies is an exercise in risk management as there are many solutions to workflow and storage resilience and backup strategies. Therefore, moving to the cloud completely, or as part of a hybrid approach allows broadcasters to take a deep look into their business continuity needs and build the required system. In other words, the technical workflow requirements of cloud deployment work together with defining the business needs and parameters.

Availability Zones

Public cloud service providers use the concept of availability zones, and these can be thought of as hardware backup systems that not only provide resilience, but also have the potential to improve latency. For example, an AWS availability zone will consist of at least two physically separate buildings with mirrored infrastructure (from the point of view of the user) so that if one building fails, then the other will take up the load.

With each availability zone, each facility has independent power, cooling, and networking infrastructure. And each facility within the availability zone connects using high-speed private networks with very low latency. Although facilities connect and share within an availability zone, and a region may have multiple availability zones, no availability zones are shared with different regions.

Load balancing solves the problem of scaling-out resource and making it appear as a single destination IP address to the user. When trying to deliver greater speed to a web-type application, the natural tendency may be to scale-up the webserver, that is, increase its processing and storage capacity. This is both expensive and cumbersome as the server will always reach a natural limit. Also, when the traffic demand is low, the broadcaster will find themselves with a

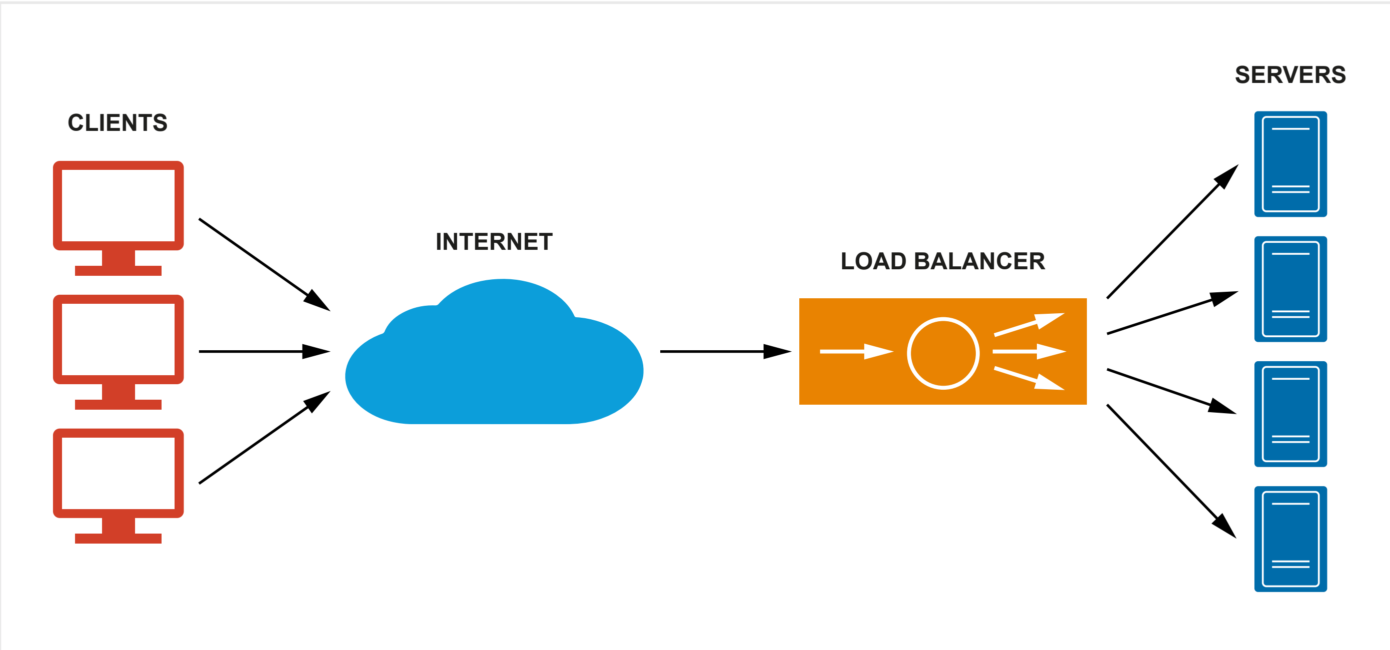

Figure 1 – Load balancing provides a single IP address for services running on multiple servers using RESTful methodologies. Also, Load balancers can be scaled-out so they do not provide a single point of failure.

very expensive server sat doing nothing. Scaling-out is the ability to increase the number of moderate sized servers, usually through virtualization, so that they can be reduced when the traffic demand is low.

Availability zones facilitate scaling-out of server and storage resource so more machines can be added and removed as required. Furthermore, the load balancing service can also scale-out so that it doesn't become a single point of failure.

Diversifying Infrastructure

A natural adaptation of availability zones is diversifying infrastructure. This works with both on- and off-prem systems as well as considering the use of multiple cloud service providers. A broadcaster could use one cloud provider for their main system with another cloud provider being available, but not used, for their secondary backup.

In the true agile method of working, scripts would be available that could spin up cloud provider-two at the drop of a hat. This is not a trivial task and maintaining two cloud infrastructures, even if only one is being used, is time consuming and hungry on DevOps time, but it does provide incredible resilience.

One of the challenges broadcasters face with this strategy is keeping the number of cloud vendor specific resources to a minimum. For example, a database design may be specific to cloud supplier-one that is not compatible with the equivalent database for cloud supplier-two, so this causes an incompatibility in the code base. Common APIs go a long way to resolving this, but broadcasters cannot assume an SQL database working on one cloud provider will port to another cloud provider.

Consequently, maintaining multiple cloud vendors requires a lot of effort and this must be balanced with the costs involved.

Operating diversity within a broadcaster's own infrastructure is much more straight forward, even if they use multiple vendor hardware as they have more control over the equipment they procure. Even this has its challenges as some equipment, such as networking routers and switchers, often have vendor lock-in through engineer training and support. Providing diversity in the internal network in terms of multi-vendor equipment is therefore not as straight forward as it may first seem.

Care must also be taken with external network providers, especially when considering diverse routing. Due to the logistical challenges of laying and installing cables across oceans, under roads, and through buildings, only a relatively small number of companies own and administer the physical cable, leading to a business model where many different service providers may be sharing the same physical cable without realizing it. Therefore, the broadcaster cannot assume that the two network vendors they have chosen for their diverse routing are using physically separate networks, understanding the true diversity of network interconnectivity requires the broadcaster to conduct intensive due diligence investigations as part of their procurement so they can truly understand the risk they are taking.

Microservice Implementation

The physical restraints of monolithic designs make scaling them very difficult, especially in the highly dynamic systems that are now prevalent in broadcasting. Microservices are continuing to advance, and not only do they deliver exceptional scalability and resilience, but they can also be cloud service provider agnostic, thus allowing broadcasters to diversify their infrastructures much easier.

Cloud vendors do provide microservice and container specific components, but they do not have to be used. If the broadcaster is willing to invest in the microservices learning curve, then they can deploy their own virtual machines and build the required containerized infrastructure on them. This further allows broadcasters to either distribute the servers across different multiple cloud regions or distribute over entirely different cloud service providers. Other than learning how to deploy dynamic microservice architectures, one of the other significant challenges is one of latency.

Using internationally distributed cloud regions allows broadcasters to position the physical datacenter hardware closer to their clients, and this has the potential to significantly reduce latency. Also, interconnected regional datacenters from the same vendor tend to be connected with high speed and dedicated networks, and this allows data to be moved with relatively low latency, much lower than could be achieved with internet connectivity.

However, to achieve the best results for container and microservice architectures, system architects must design with distributed processing, low latency, and security, from the very beginning of the build. This also facilitates the ability to build high levels of redundancy into the broadcast system at all levels of the workflows.

As viewer requirements change, broadcasters can see how the systems stress and where, so they can increase capacity as required. And as systems progress and broadcasters learn more about microservice architectures, they can even automate much of the scaling to further improve resilience and flexibility.

Chaos Engineering

It might seem a bit of a contradiction that having spent so much time and effort making broadcast workflows operate efficiently and reliably, that we should then purposely try and break them. But this is exactly how chaos engineering works.

The principle is to introduce controlled chaos into a system so that system weaknesses can be quickly identified allowing DevOps teams to find strategies to strengthen them and improve overall resilience. Although this is a powerful tool in improving reliability, it should only be conducted as part of a planned process. Deliberately deleting the microservice system that plays the opening heads of the six-o-clock news with two minutes to transmission, with no planning, is clearly a bad idea.

However, this does demonstrate the power of distributed systems such as microservice architectures that have been built from the ground-up with resilience and backup in mind. With adequate planning, it should be possible to remove network cables, switch off routers and servers, and delete applications. However, no matter how resilient a designer thinks the architecture should be, it's only when it's tested using a chaos engineering type approach does the true validity of their design become apparent.

Chaos engineering isn't just a one-off test, instead, it's an integral part of the system with "chaos" being injected into the architecture on a regular basis (with adequate planning). And this isn't just restricted to the microservice components but expands to the whole broadcast system. For example, what would happen if the electrical power supply was switched off to the datacenter? Would the UPS take the load for enough time until the stand-by generator was switched-on and stable? It's better to know where the failures are within a planned test than find them during the middle of the night when a real fault occurs.

Just like agile processes that govern DevOps, chaos engineering is a way of life that should be embraced as it will continuously improve performance and resilience of an IP broadcast infrastructure.

Continuous System Performance Monitoring and SLAs

Monitoring IT systems goes way beyond the video and audio signal monitoring broadcasters have become accustomed to. Although this is still important and plays a big role for broadcasters, monitoring IP networks, the resource they're using, and the costs being incurred, are equally important.

With dynamic infrastructures, especially those using public cloud services, new resource is added as required which incurs extra cost. Not only is it important to know when more resource is being allocated, but the effectiveness and efficiency of the algorithms deciding on the new allocation must also be continuously scrutinized so that they're not over or under scheduling cloud services. Also, DevOps teams can learn a lot from how the system is behaving overall and pre-empt any problems that may be about to occur.

One example is a low API response time, this could be due to an unusually long database query or network congestion, but the monitoring will give a more focused indication to where the latency is occurring. Research has demonstrated users can adapt to reasonable amounts of constant latency, but variable latency is difficult to work with and problematic. Due to the dynamic nature of the infrastructure, the latencies can develop and concatenate without warning, hence the need to provide continuous monitoring.

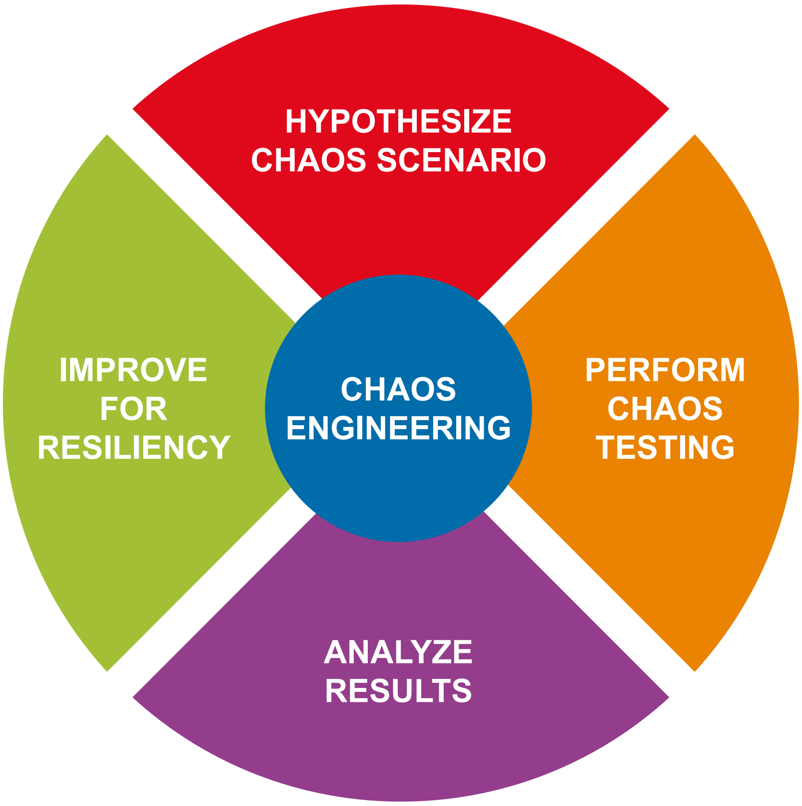

Figure 2 – Chaos engineering is a continuous practice that aims to improve resilience by testing and stressing infrastructures for unexpected events.

As well as providing deep insight into the operation of the system, continuous monitoring highlights areas of the infrastructure that would best benefit from service level agreements (SLAs) and determine the extent of their cover. Providing blanket SLAs for every piece of equipment or service is costly and inefficient, especially as many vendors now provide varying levels of SLA. Being able to allocate the best SLA for each part of the infrastructure components will deliver optimal resilience and efficiency.

Conclusion

Determining best practices for building hybrid, on- and off-prem cloud and virtualized infrastructures is a methodology that must be considered at the very beginning of the design. These practices further develop as the functional aspects of the workflows expand, and therefore need to be under constant review. Due to their dynamic nature, best practices, like so many other aspects of broadcast IP infrastructure design and maintenance, should be considered an ongoing process that is always open to development and improvement.

The Sponsors Perspective

What Are You Waiting For?

By Chris Merrill, Director Strategic Marketing, Grass Valley

Cloud based workflows are here to stay. So, there are no longer any technical or production-based reasons not to take full advantage of the opportunities this can bring to broadcasters.

Taking Advantage of a Force Multiplier

Everywhere you look there are articles and testimonials discussing the benefits of cloud-enhancement to Media & Entertainment workflows. Most of this information, however, is from vendors to the industry rather than consumers of the technology. It's therefore fair to ask why we aren't hearing more from technology consumers.

At Grass Valley, our experience is that adopters do see great benefit from using our AMPP platform. In fact, they see their use of the technology as a competitive advantage and therefore aren't eager to share what they are learning with other media organizations. Consequently, they are quiet in the public press.

A recent survey by Deloitte of 500 senior decision-makers in the United States found that 90% of respondents agree, or completely agree, that cloud combined with other technologies such as AI, IoT, and analytics serves as a "force multiplier" for their digital strategy (Deloitte, US Future of Cloud, June 2022). Overall, respondents think that incorporating cloud technology drives positive outcomes for businesses across a range of key measures including:

- Mitigating business and regulatory risk
- Building net new product/service revenue
- Expanding into new markets
- Providing better sustainability in support of environmental issues

Our own AMPP usage data validates this survey. Media & Entertainment companies are enjoying the value of cloud-enhanced technology. As a whole, the industry has moved past the Covid-19 induced panic, but the POCs and trials were valuable experiences. Many have retained what they learned as a result. The use of cloud-enhanced workflows is not a flash-in-the-pan experiment.

Over the past 12 months, AMPP adopters are consistently expanding their use both in terms of time and functionality. The average usage of AMPP has tripled to multiple millions of workload hours per month, creating a massive force multiplier across the industry.

Overcoming Status Quo Bias

It's natural to approach any change with caution. Studies on business decision making show that the most common response when faced with a choice to implement a new option is to stick with the status quo. In physics we learn that an object in motion stays in motion unless acted on by an external force. It's a similar tendency in the psychology of decisions making except this status quo bias is based on emotion rather than data. Change involves risk and humans tend to avoid risk whenever possible. The downside of that bias is that it can create a state of no change until an even greater risk becomes so compelling that it forces us from our previous position.

The time to reconsider the market risk of not adopting cloud-enhanced technology is now. The pace of technology adoption is speeding up.[1] What used to take many decades to become widely adopted is now happening in a few years. The COVID Pandemic further accelerated that rate with 75% of all firms adopting digital productivity-enhancing technologies.[2] By waiting to see how cloud-enhanced technologies help others increase their revenue per piece of content, Media & Entertainment companies run the significant risk of losing their own competitive advantage.

[1] *https://hbr.org/2013/11/the-pace-of-technology-adoption-is-speeding-up*
[2] *https://blogs.lse.ac.uk/businessreview/2021/11/22/covid-19-spurred-a-wave-of-new-technology-adoption-by-uk-businesses/*

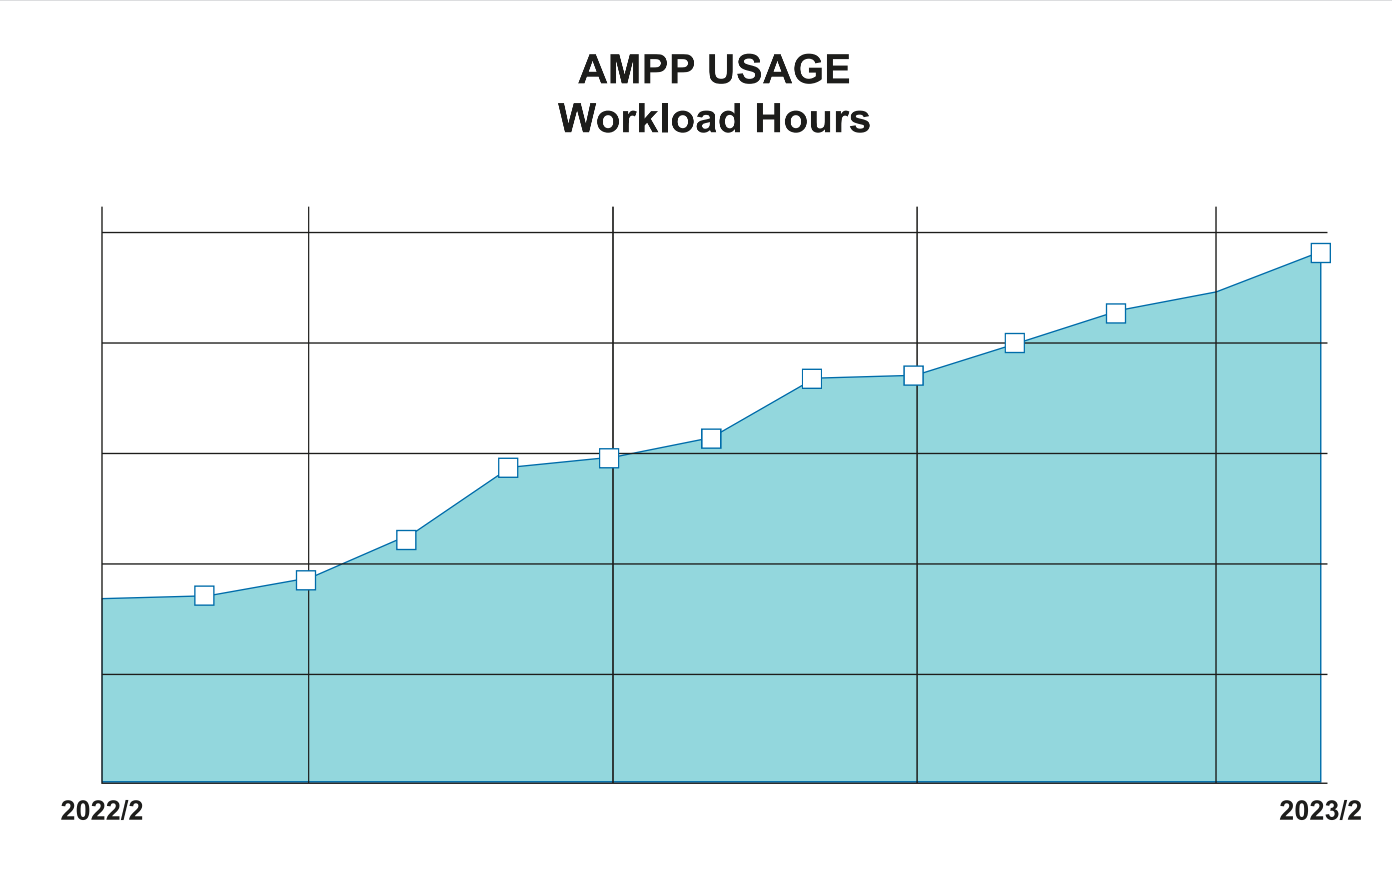

Supported by

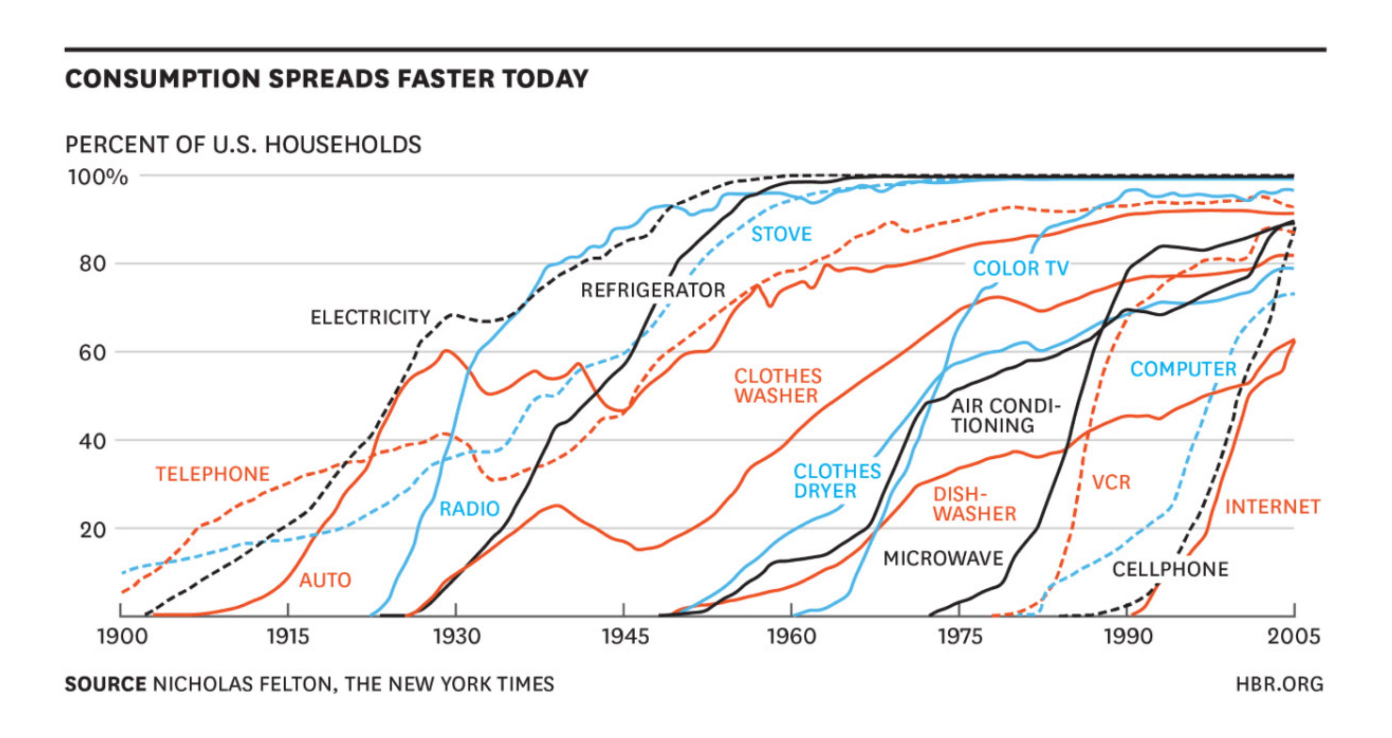

Incremental Adoption

The good news about deploying the AMPP platform is that it isn't an all or nothing proposition. Most new entrants to the platform start small. For example, start with remote monitoring, or signal conversion and routing.

The platform can easily be integrated into any existing SDI or IP infrastructure. Rather than completely renovating an existing plant, expand an existing workflow to add new capability. Because the barriers to entry are low and the user interfaces are familiar, adoption can occur gradually.

They key to getting the most out of this gradual adoption is a willingness to experiment. If AMPP is used like other technologies which are wired in the rack and never touched again, then you are missing one of the greatest advantages. AMPP is a platform with many possible applications. It doesn't have to always be the same. Functions from as simple as router control panel to as complex as a production switcher can be quickly spun up or down as required.

The trend we see is that as AMPP adopters become familiar with the ecosystem, they begin to expand its usage at their own pace moving, for example, from monitoring to local versioning to full live production.

Recommendations

Regardless of the cloud-enhancements you actually select, we recommend the following;

First, look for extensible cloud platforms rather than single-purpose apps. Trying to string together multiple apps from different vendors is not as simple task. Nor is it ideal to have to work with multiple, dissimilar user interfaces and tool sets.

The lack of common standards and interfaces is the prime driver behind the GV Alliance. Solutions selected from GV Alliance members, regardless of who originally developed them, have been tested by the vendors and are guaranteed to work together.

Second, look for solutions that will work with your existing hardware and software. The transition is much simpler if the system users don't have to change the way they are accustomed to operating.

Finally, wherever possible talk to actual users of cloud technology. You'll discover that its use is more widespread and easier to adopt than you might anticipate.

THE BROADCAST BRIDGE
Connecting IT to Broadcast

Themed Content Collection

EG
ESSENTIAL GUIDES

MEDIA

ARTICLES

For hundreds more high quality original articles and Essential Guides like this please visit:

thebroadcastbridge.com

04/2023

Supported by

Grass Valley
WE LOVE LIVE