



Cloud Video Platform: Realizing broadcaster-grade Cloud TV

Introduction

The digital transformation of TV is in full swing. However, a key driver of successful digital transformations is the ability to realize the benefits of moving technology to the cloud without compromising the mission-critical nature of TV.

Most broadcasters target at least four 9s (99.99%) of availability, which means at most 4.32 minutes of air a month — about the length of a single ad break in a single program broadcast in a single month. With many cloud services only offering three 9s, equivalent to 43.2 minutes on air a month, achieving the broadcaster target requires more than just a back-to-back service-level agreement (SLA).

However, even where a four 9s cloud platform is available, the availability of the underlying cloud technology doesn't automatically translate to achieving four 9s at the service level. When an issue occurs, the key metric is mitigation time — the duration of the issue occurrence, the service's operations team realization of it, the issue triage, and the successful restoration of service. That's a lot to do in 4.32 minutes.

The innovation described in this paper is therefore the combination of technologies and business processes deployed by the Cloud TV Suite from Comcast Technology Solutions that enables us to offer a broadcast-grade video platform with four 9s availability. The paper demonstrates this by describing four headline components within the broader architecture.

No solution of this type will ever be perfect, and there are many new innovations in the development pipeline to further strengthen our architecture. We therefore conclude with a look at three such innovations in the works.

Cloud TV Suite today

Innovation 1: Multi-application fairness

Given that the Cloud TV Suite is a multi-tenanted platform, in our context, this innovation primarily means prioritizing fairness between different tenants (i.e., broadcaster customers). But it still applies even in single tenanted platforms where the platform serves multiple device applications and a small subset of those devices encounters an issue.

The challenge here is when a small percentage of calling devices suddenly accounts for a very large share of requests. This could be due to a bug, but it could also be due to a sudden spike in traffic, for example, all users connected via a particular ISP suddenly reconnecting following a glitch in that network.

The best way to meet a four 9s SLA is to prevent outages in the first place, and this is very true in this area, where we have built three key defenses against traffic spikes:



Preventing resource monopolization: Runtime analysis groups callers together and identifies when one group is taking a disproportionate share of resources. If this happens, requests from that group can be dequeued to help maintain overall platform stability.



Responding with low latency successes while overloaded: Our goal here is to successfully serve as many requests as possible. This means dequeuing requests when overall load is high, preventing requests from being accepted in the first place, since if high load has already caused response time to increase beyond specified limits, then there is little point trying to process additional calls.



Targeted request rejection when overloaded: Most services process requests via a pipeline, with queuing between stages of the pipeline. When one stage becomes overloaded, then other requests are analyzed to identify requests with closely matching criteria. This means requests that need highly loaded resources can be rejected without impacting traffic for more lightly loaded resources.

All the techniques highlighted in this innovation demonstrate how a four 9s SLA can be achieved using bespoke engineering to go beyond simply relying on the inherent elasticity and scalability of cloud platforms.

Innovation 2: Failover

Continuing the theme in which the best way to achieve four 9s is by preventing issues before they occur, managing failover is critical.

The use of availability zones within a region is standard practice. This means ensuring microservices have resources in each of three zones. For databases, it means building scaling based on at least two readers and one writer service, and building logic to promote a reader to a writer in the event of failure with the original writer.

However, these are table stakes measures and of themselves are insufficient to ensure broadcaster-grade availability. Therefore, our key innovation here is building a geo-balanced active/active architecture on top of our cloud services provider. This helps balance traffic between regions, and in the event of an issue, enables one of the active paths to absorb the traffic from the failed path.

Running active/active helps minimize the latency when failing over regions. However, manual intervention based on anomaly detection is still too slow for four 9s availability, and so our platform uses extensive custom logic to automatically detect anomalies, which automates the failover trigger.

This innovation demonstrates that simply moving to the cloud often isn't sufficient to maintain broadcaster-grade availability norms. Digital transformations therefore extend beyond the technology layer and require integrated developments within an organization's business processes and tools.

Innovation 3: Precomputation

Although digitally transforming broadcasting means moving to the cloud in some form, this migration isn't of a single monolithic system. At the very least, it is a combination of compute, databases, and other enabler technologies (such as API gateways, load balances, and DNS). More likely, each of these high-level concepts are themselves comprised of multiple sub-systems. For example, many cloud compute architectures are based on microservices, to facilitate easier maintenance and scaling.

In this context, an issue with one sub-system can have a domino effect on other systems. Thus, a service performing outside of norms but within expectations can have an amplifying effect on dependent services, until one sub-system eventually fails. Such problems are difficult to triage within 4.32 minutes. In addition, recovery to a healthy state can create further issues as backed up requests from other services suddenly cascade onto the affected service.

The Cloud TV Suite mitigates this using precomputation to reduce or eliminate runtime dependencies between systems. This helps make issues easier to triage. An example of this is the Grant/Deny Service, which precomputes end-user entitlements by determining every combination of end user and content in advance, meaning entitlement checks no longer need to look up user information, content information, and rights information from other systems.

This in itself is still insufficient, since large-scale changes that impact many or all users (for example, adding a new Christmas Movies collection) can create a cascade of changes that need recomputing. Users requesting content in this period must not be served with decisions based on stale data. Hence, further innovations examine inbound traffic so that recalculating specific combinations can be prioritized via a high-priority queue.

Innovation 4: Business processes

As highlighted previously, achieving four 9s with cloud architectures requires more than just migrating services to the cloud and relying on the underlying elasticity and scalability. It also requires bespoke technology to overlay the cloud service to ensure resource allocation fairness and gradual, targeted degradation.

However, success depends on going further, since even the best technology platform will occasionally have problems. As with many IT and digital transformation projects, success requires the transformation of business processes as well as the underlying technology.

In this context, the business process transformation means ensuring a culture of reliability is driven into the organization. This translates into key documentation such as maintaining an up-to-date risk register, which is proactively worked on as a form of tech debt, and runbooks, which can be easily parsed during issues that describe how to recover from typical fault scenarios, ensuring services are quickly restored.

Cloud TV Suite tomorrow

Building a resilient platform capable of four 9s availability isn't a "one and done" project. It is an ongoing, continuous investment. We believe the above developments are innovative not just individually but also as part of a holistic approach to resiliency. They are just some examples of how we have progressed the state of the art with regard to resilient architectures.

However, we are not resting on our laurels and have more developments in the pipeline. We close with a brief look at some of them:



Chaos engineering: This is viewing our system from the outside in, continuously calling it to detect service degradation and recording typical call sequences to mimic real-world traffic.



Automated tooling: As our business processes are refined, the next step is to automate as much of them as possible. Automated tooling can be used to trigger workflows based on particular types of issues, initiating key mitigation steps (such as proactively triggering service restarts and automatically paging support teams), freeing up operations teams to focus on triage and mitigation.



AI/ML: Machine learning will help push reliability engineering on to the next step of performance. As highlighted, there is a lot to do in 4.32 minutes, but using machine learning to synthesize large datasets to spot trends indicating an issue can help give operations teams a head start on triaging and mitigating before service degradation occurs.

Conclusion

In summary, our submission is that the resiliency architecture of our Cloud TV Suite has progressed the state of the art in the field of broadcaster cloud transformations in three ways:

1. We have successfully developed multiple theories of reliability engineering into practical reality.
2. Individual approaches are not enough to achieve four 9s, and so, in addition, we have built technology that combines those individual approaches into a single holistic solution.
3. A single holistic technology layer in itself is still insufficient to achieve four 9s, and so we have wrapped that solution with a set of business and operational processes.

The field of reliability engineering is often overlooked in a market that is rapidly pushing the envelope with innovations in business models and using AI/ML to enhance workflows and processes. But it is a critical field, and mastery of it will be essential in ensuring that the digital transformation of the TV and media industry successfully translates into the online era of consumer expectations that were established by previous TV generations.

Power up your platform with CTS's Cloud TV Suite

Comcast Technology Solutions, a division of one of the world's leading media and technology companies, brings Comcast Corporation's proven technologies to an evolving list of industries worldwide. We believe in continuous innovation, always looking for new and better ways to connect with our customers, as well as aggregate, distribute, and secure our own content, advertising, and data. We invest in and test these solutions, so you don't have to — freeing you up to focus on accelerating your business, not your tech stack. We bring these innovations to the global marketplace, enabling our partners to think big, go beyond, and lead the way in media, technology, and cybersecurity.

Find out more

+1 800-824-1776 | ComcastTechnologySolutions.com
comcasttechnologysolutions@comcast.com

