White Paper

Centralization of Master Control Operations

Phil Wilton
Product Manager, Snell Ltd

Abstract

For many years now file-based transmission playout has been a key part of most broadcast platforms. Growing consolidation among the content aggregators that run such platforms, however, raises the question of how to achieve even greater efficiency in the operation of multi-channel playout platforms. This paper presents an overview of the implementation of a hub and spoke broadcast architecture to support centralized acquisition, distribution and delivery of assets that also offers the flexibility for local programming of commercials and news.

A hub and spoke infrastructure should not increase operational costs in terms of duplication of ingest processes and storage nor should it increase operational or manpower costs. In addition the approach offers a suitable level of resilience without the need for increased or duplicated investment at the spoke sites.

The architected solution allows for resilience to be centralized within the hub site in terms of material archive and alternate playout platforms, which can also be utilized at local spoke sites. Centralizing core high value assets allows for maximum flexibility in terms of repurposing them across multiple regions and platforms, also costly processes such as QC only have to be performed once.

Introduction

At the enterprise level installing, operating and running an efficient multi-channel delivery platform is a balancing act. On one hand a sizable portion of the running costs will be consumed in three main areas, capital investment, operational costs and bandwidth capability. The capital investment required to suitably equip such an installation, with the necessary broadcast infrastructure and equipment, is still significant despite the on-going convergence of I.T., Telco and Broadcast technology. Coupled with this will be the usual on-going operational costs including technical, operational as well as the necessary support staff. Also a sizable part of the day to day running costs will be consumed by the relevant connectivity, or bandwidth, between the various broadcast sites typically covering I.P. fiber, Telco and satellite feeds. These can take the form of dedicated or ad-hoc circuits for the contribution and distribution of programme content. Whilst it is fashionable to be positioned as a completely file based operation it is impossible to ignore tape based content as virtually all media organizations have access to very large tape based archives. It is therefore inevitable that provision has to be made for the acceptance, ingest and QC of tape based material, plus metadata mark up handling.

On the other hand the continued success of any transmission facility or content delivery platform means that it must be able to evolve as market trends change and in some cases accommodate a total shift of its core business models. Currently this growth comes from two key areas. Firstly we are still seeing an increase in what is termed traditional broadcasting or linear delivery of programming content. This includes the exploitation of niche content in what is sometimes called "The long tail" as well a simple time delay or "+1" channels. In addition media companies as well as traditional playout operators are moving ever closer to the non linear delivery of content in what has been called the "Anycast" scenario. This means that consumers will choose what they want to see, when the want to see it, where they want to see it and on what device they want to consume it on. In addition we have the push to deliver linear programming onto dedicated platforms such as mobile devices with initiatives such as the Open Mobile Video Coalition.

Maintaining this strategic agility is key to survival but it has to be maintained and implemented in the face of the inevitable downward trend in budgets and the resulting reduction in available spend. Whilst at the same time maintaining or even increasing the day to day operational capabilities and flexibility in delivering programme content as effectively and efficiently as possible. Not forgetting the fact that as a commercial organization there is the inevitable pressure to be able to deliver 24/7 without the loss of any revenue streams from spots or commercials, so utmost reliability and resilience are also up there at the top of the wish list.

In the past operators of playout platforms and manufacturers of playout automation offerings have tried to implement this central casting model through what many now call a "hub-and-spoke" operation.

However, through a variety of shortcomings, such as shared components across the infrastructure resulting in a single point of failure, these vendors and products have failed to deliver a real solution to the demands of multi-channel playout in a central casting scenario.

In order to achieve the best possible solution centralized playout architectures must be studied in great detail, including technical and operational requirements, as well as each model's likely impact on all areas of customer's costs with respect to capital hardware investment and on-going operational costs. Through experience, analysis and real-world deployments, our team of engineers have refined the well-established Morpheus automation system to provide an enterprise level solution that enables a high degree of efficiency and resilience for hub-and-spoke architectures whilst addressing the demands of a rapidly changing media marketplace.

Centralization: Challenges and Benefits

Centralization of broadcast operations is not a new concept; indeed most content aggregation and content distribution companies already operate central or hub sites either as discrete entities or as part of their global operations. Centralization of hardware and common operations makes a lot of sense and brings a scale of efficiency to the whole operation as well as maximizing utilization in all areas.

Typically the central or hub site will provide for the majority of the operations that require either manual operation and/or large installation of hardware such as servers and ingest stations as well as near-line and deep archives.

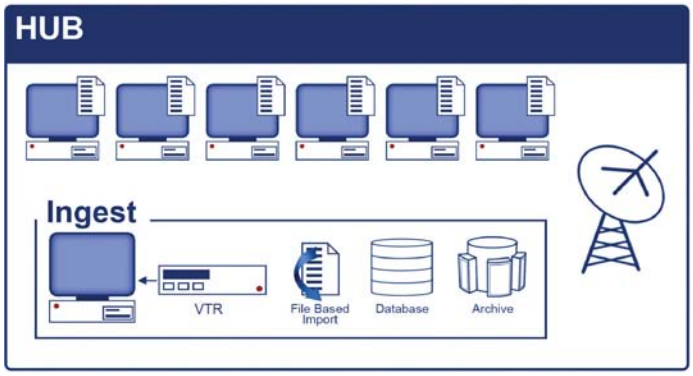

Figure 1: Hub components

Media organizations will either own or have access to a large amount of content, sometimes running into hundreds of thousands of hours or tapes, typically this will be long form content. This means that provision has to be made for tape based ingest. This will typically involve a number of operator driven ingest stations based around a VTR, ingest server and suitable ingest application including metadata mark up. Sometimes these are semi automated using tape robots such as the Sony Flexicart which allows a number of ingest streams to be imported in parallel. QC is an essential part of this process along with the capability to make any adjustments to the video signal. It is also possible that for international content different language audio tracks may be required, as well as subtitling. This can of course be handled off-line as a file based process with any suitably equipped dubbing or subtitling suite. As well as technical QC there will also be a requirement for content to be checked from a legal and licensing perspective as well as any other relevant compliance requirements. Ingested content can then either be used for transmission or moved to near-line archive or a hierarchical storage management system that might well include deep archive on data tape. This deep tape based archive and traditional video tape based archive means a lot of media management and manual movement of tapes to and from various processes; indeed one major broadcaster has indicated that it can take up to 15 hours of different process to deliver 1 hour of screen time. In addition there will also be the requirement to accept short form content such as promos and spots; these could arrive by tape but it is increasingly common for these to be delivered as files via various third party delivery suppliers.

Having dealt with legacy and archive issues there are the inevitable requirements for file based delivery and distribution. At the enterprise level this can easily cover anything from, say, 10 to 60 channels of content that has to be delivered to a number of different regions across different time zones. This inevitably means a large investment in playout servers, graphics engines, subtitling systems, routing matrices and master control devices all under automation control. However in this case study most of these resources are deployed within the spoke sites. In addition there will be interfacing to external traffic systems with schedule load and edit capabilities as well as back office functions such as, "as run logs", billing and rights payments etc.

A successful playout automation system must offer 3 key features:-

1. Utmost reliability is required as the automation system is responsible for generating and delivering the revenue stream. A reliability figure in the region of 99.9999% has been achieved. This level of resilience should cover playout database, media database, playout event store and playout hardware.

2. Maximum flexibility to cope with the ever changing architecture requirements of channel playout. Therefore it should be capable of handling multi channels, local opt outs, local regions, different time zones, flexible channels, preview channels and back up channels plus associated hardware.

3. Despite being an automation system with a very high level of reliability and resilience there will always be operational and technical requirements for manual intervention and control.

It is not uncommon for a sports channel to spend up to 50% of the peak schedule operating live via manual control. In addition there will be ad-hoc requirements to take manual control in order to pre-empt an on-air situation due to known hardware failure.

Incorporating all of the above elements and process within a single centralized site makes good business sense in that maximum use can be made of all resources, both component parts servers, VTR's tape robots, graphics engines, subtitling systems and master control facilities etc. Operational efficiency is maintained by ensuring best use of all operational, technical and support staff.

Hub-and-Spoke Architecture

The fundamental requirement for implementation of a successful hub-and-spoke architecture is to place at its heart an enterprise-level playout automation system. In one installation Morpheus automation has been integrated into just such an architecture, providing 43 channels of playout covering a total of 11 spoke sites, distributed across a number of different time zones. Within the hub Morpheus applications provide the tools necessary for all ingest operations related to central programming.

In providing these tools the automation system minimizes media movement and the burden on operational staff but also maximizes the use of hardware.

External traffic systems provide the central and local programming schedules which are loaded into Morpheus. Central playout control is handled across 8 workstations with each workstation operator handling a number of channels. Under normal network operation file based content is delivered centrally from the hub. Solutions architected with Morpheus provide resilience within the hub site by way of its content transfer capabilities, archive and management systems, and alternate playout platforms – all of which can be leveraged by local spoke sites. Centralizing core high-value assets and operations at the hub also yields maximum flexibility in repurposing content across multiple regions, with minimal cost implications, as the majority of these operations only happen once at the hub and ensures huge cost saving benefits as they do not have to be repeated at any of the 11 spoke sites.

Staff at the hub or at any of the spoke sites can take operational control over content delivery; broadcasting programs of particular interest to local and regional viewers.

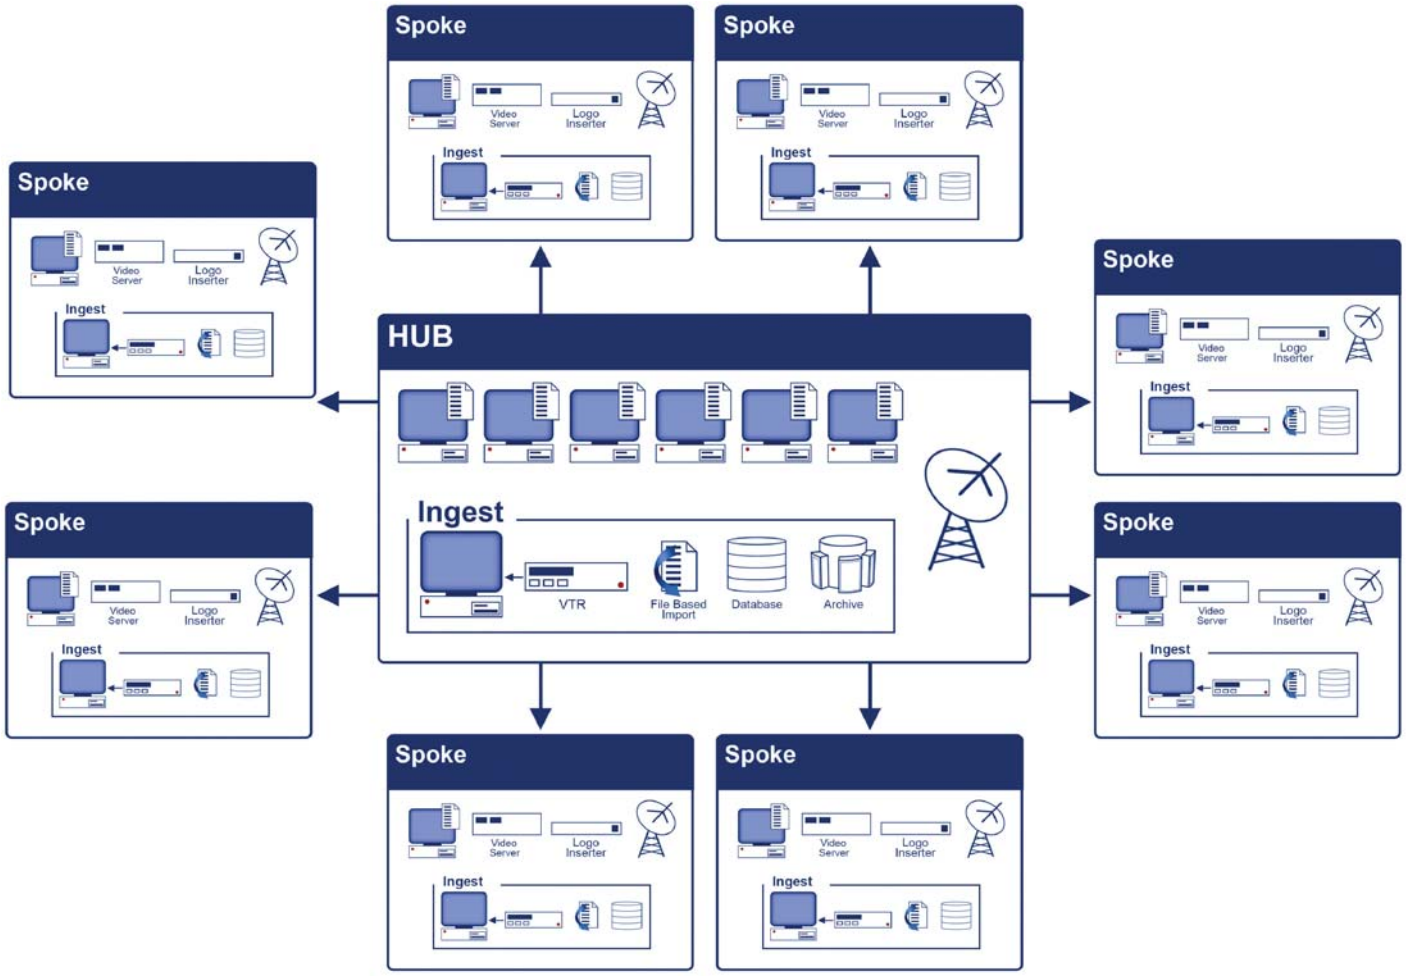

Figure 2: Overview of hub and spoke architecture

For example, during mainstream programming hours, control of area news and events can be handed off to various spoke sites where minimal operational support can oversee a local installation of Morpheus to handle creation and editing of community-type programming and news schedules. These schedules are also visible within the hub and any new content requirements trigger a transfer from the hub down to the spoke(s).

This workflow can also include ingest of content directly within the spoke for commercials, news, and programming specific to the viewing audience. Overnight, or at less busy times, each spoke can be operated totally "hands-free" with control being handed back to the hub site which maintains control and scheduling of all local programming as well as content ingest. Again this model offers huge savings in resources as all 43 channels across the 11 spoke sites can be controlled by around 8 operators within the hub facility.

Morpheus can push content out to any or all of the spoke sites at any time supplying up to three days worth of what would normally be centrally delivered programming. As each spoke site is configured with a complete standalone automation and media management system, including local ingest, Morpheus is able to offer a higher level of resilience that protects playout in the event of engineering or operational issues. These could be at the hub or if the network connection between the hub and one or more spoke sites is lost, it allows the isolated spoke site to continue autonomous operation for up to 2-3 days. In addition any spoke site and the local schedule can initiate a "pull" for material that still only exists at the hub site. There is also support for the ability to retrieve (pushed or pulled) material from the spoke sites so the transfers can be either way. Transfer agents/accelerators are required particularly if HD files are involved.

Spoke Requirements

While centralization of playout operations yields many benefits it also poses several nontrivial challenges. Among these is the challenge of providing localized content ingest and delivery. A key business and operational requirement of any affiliate or local / regional station is the ability to ingest and insert local or regional advertising, news, and other programming into the playout schedule. Any duplication of centralized investment on the local level is prohibitively expensive and therefore any deployment locally must be very lean. Furthermore the movement of media should be kept to a minimum in order to minimize the bandwidth and operational costs of distribution and contribution, as well as the impact of media transfers on "time-to-air" requirements. Therefore large content aggregators and distributors covering sizable geographical areas must find a way to address and enable local ingest and delivery of content without duplicating or increasing capital, operational, and bandwidth costs.

Spoke requirements will typically mirror a subset of the hub infrastructure. Provision must be made for local ingest of content. Local spots and commercials are usually typically centred on file based delivery whereas news and local programming could also come by tape; usually 1 tape ingest seat is required along with ingest server and the same ingest application as the hub with access to the main database to ensure that all material is registered.

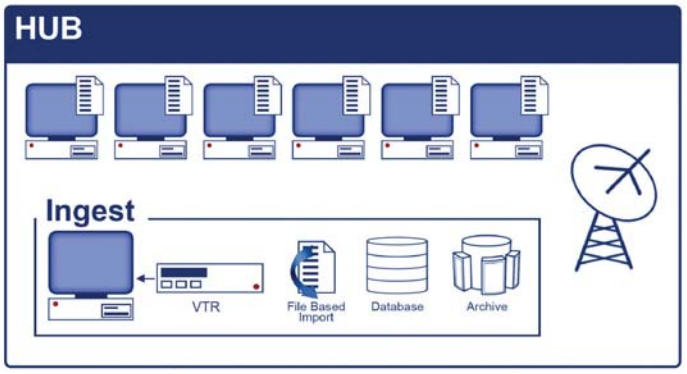

Figure 3: Spoke components

It is possible that a simple tape robot could be installed at the spoke thereby allowing remote ingest operation by the hub leaving the spoke to operate at a very lean manpower level with no specialized operators. Non tape based acquisition formats from Panasonic and Sony are also becoming popular with the advantage that they both offer simple file based import of completed "air-ready" content. This can be accomplished by FTP delivery from a remote connection or a simple docking connection in the spoke site. Playout facilities will encompass servers, routers, master control mixers as well graphics engines for local branding. In addition there could be local subtitle or language insertion depending on geographical requirements.

Operation

Under normal daily operations external traffic systems provide the central and local programming schedules which are loaded into the playout automation system, Morpheus. Central playout is handled across 8 workstations with each workstation operator handling a number of channels. The main long form material is distributed from the hub to the spoke sites typically 2-3 days in advance and stored on the local playout servers based on the scheduled playlist.

The Automation Edit Workstations under operator control within the hub site are networked to the specific spoke sites that they are controlling. Operators at the hub site watch ASI return feeds from the spoke sites for monitoring purposes. All engineering and control applications can be remotely managed.

An automated material checker reviews the content listed in the transmission schedule and will check against the content in the automation database to ensure that it exists. Any missing material is highlighted and a missing material transfer request is made back to the hub to deliver the content to the spoke. At the hub site and/ or any of the spoke sites the playout schedule can be edited to accommodate local programming such as spots or news. The control and operation of any spoke site can be taken either centrally within the hub or locally within the spoke at any time.

Therefore manning of the spoke site can be kept to a minimum with even totally unattended operation at non peak times such as during weekdays or overnight operation. Staff at the hub or at any of the spoke sites can take operational control over content playout broadcasting programs of particular interest to local and regional viewers. For example, during mainstream programming hours, control of area news and events can be handed off to the local spoke sites where minimal operational support can oversee a local installation of Morpheus to handle creation and editing of community-type programming and news schedules.

For live programming across the network a multi site "Gang Take" feature ensures frame-accurate regional content insertion across multiple sites with all selected sites being triggered by just one transmission controller. Both these operational features ensure a high level of resilience within the total system such that control and editing of the playlist can be carried out at any site for any other site covering the hub and all the spokes.

The nature of centralized playout operations makes system resilience a critical factor in system design. Correct implementation of the chosen architecture is another challenge that must be met properly if the full channel output is to be maintained continuously. Given these challenges any hub-and-spoke model designed for multi-channel playout over multiple regions must offer a high level of resilience and incorporate tools and functionality that empowers centralized staff to manage and deliver local content effectively. By implementing this concept the local spoke sites retain the feel of a "local TV station" but with less local spend and all the benefits of the central resources. Local programming can be acquired locally but played out nationally including different regions and time zones.

The level of resilience at the spoke can also be a sub set of what had been implemented in the hub.

Typically the servers running the database and media movement transactions will be configured as Main, Mirror and Witness. In terms of ensuring continuity of the playout channel there are typically two main architectures used:-

1. Complete duplication of the playout servers with what is termed a "guard channel". This is a back up server which can support the main output stream should it fail. The back up device will be pre-loaded with the same material as the "on-air" server and will be controlled by the playout automation system in parallel. This allows for selection of the guard channel at any time either in anticipation of an up and coming technical problem or at any time an operational situation should arise and ensures that there is no loss of continuity in programme output. This form of resilience is expensive due to the additional cost of the required hardware on a 1:1 basis and is mainly used to support high value channels.

2. However it is also possible to implement guard channels at a lower cost threshold, so for example a VTR can guard a server port or a caption generator could guard a server port. These alternatives do not provide for the same level of continuity as a guard server but they can be implemented at a much lower cost and can be configured to provide back up for high or low revenue channels.

The most economical solution for the spoke operation is what is generally termed N+1 operation. This is typically a hot spare server on which there is no pre-loaded content. This means that the N+1 channel can be used as a back up to cover all of the channels in the spoke operation. Therefore in anticipation or on the realisation of an on-air problem the selection of an N+1 channel will initiate a transfer of material to the hot spare server and then subsequent playout to air. This action does of course lead to a short loss of continuity but this has to be balanced against the cost of implementation and the value of the channel being protected.

The use of the N+1 architecture does offer other operational capabilities and benefits such as preview, where it can be used to either preview any part of the up and coming schedule, or it can be used to preview the transitions between the end of the current programme and the start of the next programme, or to ensure that the branding from say a graphics engine is correct.

Overall this means that the implementation of the installed hardware is the same as for any large broadcast operation but you have the advantage of geographical separation which means effective disaster recovery deployment.

Conclusion

By properly implementing a hub-and-spoke architecture for multi-channel playout, content aggregation and distribution companies whether broadcasters or other media companies can balance the demands of on-going capital investment in broadcast infrastructure and operations with the need to remain flexible within a continually evolving media delivery market.

Already installed at the heart of large enterprise playout architectures Snell's Morpheus automation serves as a proven solution for efficient operation along the hub-and-spoke model. In doing so the system enables users to centralize media management and operational processes within a hub site and, in turn, to achieve a significant degree of resilience through centralized delivery and archive support for local spoke sites. Centralization of media storage at the hub also eliminates redundant tasks – such as multiple QC checks at different sites – while making it more cost-effective to repurpose that content across multiple regions and platforms.

The distributed architecture of the hub-and-spoke model does not increase operational costs in terms of duplication of ingest processes and storage nor does it increase a facility's personnel costs. Rather the hub-and-spoke model offers flexibility that allows for tailoring of playout operations to meet specific functional or business needs. This model of centralized multi-channel playout also provides a high level of end-to-end resilience without requiring an increased investment at the spoke sites thereby offering maximum operational versatility and excellent value for money. This architecture also allows for inherent distributed disaster recovery deployment without duplication of largely redundant resources.

Acknowledgements

I would like to thank my colleagues at Snell for their comments and contributions and the directors of Snell Ltd for permission to publish this paper.