Architecture of a Broadcast and Media Production Learning Center Integrating Cloud-Based Ingest and Automated News Classification

Adrian VINTILĂ, Constanța-Nicoleta BODEA
Bucharest University of Economic Studies, Romania
adi.vint@gmail.com, bodea@ase.ro

This paper presents a reference architecture for a learning center designed to support engineering education in television broadcast and media production. The proposed environment replicates a small-scale but operationally complete television broadcast facility and is intended to give students hands-on access to the equipment, software, and workflows that characterize modern newsrooms. The architecture is organized around four interconnected components that follow the content lifecycle: media acquisition, processing and planning, production, and distribution. Two original subsystems, previously designed and evaluated by the authors, are integrated into this architecture in order to expose students to current technological trends. The first is a cloud-based automated ingest subsystem, embedded within the media acquisition component, which removes manual steps from the file-transfer pipeline and accelerates the availability of mobile journalism media assets. The second is a supervised machine-learning subsystem for the automatic classification of Romanian news stories, embedded within the processing and planning component, which supports editorial organization and newsroom decision-making. Five learning scenarios are defined to illustrate how students engage with the architecture across all four components, ranging from mobile journalism capture and cloud-based ingest to control room operation and on-air delivery through a streaming platform. The paper's contribution is the consolidated reference architecture and the integration of these two previously evaluated subsystems into a coherent educational framework. A systematic assessment of student learning outcomes is planned for future work.
Keywords: *Practice-based learning, Media asset management, Hybrid cloud workflow automation, Supervised text classification, Television production*
DOI: 10.24818/issn14531305/30.2.2026.02

1 Introduction

Theoretical knowledge is essential in any engineering program, but in broadcast and media production it is not sufficient on its own. Students need to work directly with the equipment, the software, and the workflows used in real newsrooms in order to develop the skills expected by employers in the field. This need is reinforced by the pace at which broadcast technology changes: hardware, software, and production practices evolve quickly, and curricula has to keep up [1] [2].

For most universities, setting up and maintaining a fully equipped broadcast laboratory is expensive and difficult to sustain over time [3]. A shared learning center, designed as a miniaturized operational television facility, offers a practical answer to this problem [4]. It brings equipment, software, and staff together in one place that can serve several programs and institutions. Such a center can also be used to test new functionalities, hardware integrations, and software upgrades, and to support the continuing professional development of engineers already working in the industry.

This paper is a reference architecture paper. Its contribution is twofold. First, it provides a consolidated description of the architecture of a learning center for engineering programs in broadcast and media production, organized around four main components that follow the lifecycle of media content from capture to distribution, and offers practical learning scenarios for each main component. Second, it integrates two original subsystems into this architecture: a cloud-based automated ingest pipeline for mobile journalism content, and a supervised machine-learning system for the classification of Romanian news stories. Both subsystems were designed, implemented, and evaluated in previous work by the authors, and are not re-evaluated here. The ingest pipeline was tested against several existing workflows in an operational news television station and reduced end-to-end ingest time by more than an order of magnitude while also eliminating the need for a dedicated ingest operator [5]. The classification subsystem was trained on a corpus of approximately 24,000 labeled Romanian news stories collected from the same station over a fifteen-month period and reached near-perfect accuracy for the best classifier and vectorizer combinations after class balancing [6]. The present paper describes how these components fit into the broader architecture of a learning center and how they are exercised in structured learning activities.

The architectural novelty of the paper is concentrated in the acquisition and processing and planning components, which host the two integrated subsystems. The production and distribution components complete the reference architecture and are described through state-of-the-art broadcast building blocks, with focus on automation systems, thereby providing a modern operational context for the learning scenarios. The architecture is offered as a reference design intended to guide implementation at higher education institutions with engineering programs in broadcast and media production. A systematic evaluation at the level of student learning outcomes is left as future work and is discussed in Section 5.

The remainder of the paper is organized as follows. Section 2 introduces the concept of the learning center and its main components. Section 3 describes the technical architecture, including the two integrated subsystems. Section 4 presents the hands-on learning scenarios designed for the architecture. Section 5 concludes the paper and outlines directions for future development.

2 The Broadcast and Media Production Learning Center Overview

2.1 Concept and Educational Rationale

The learning center described in this section is a proposed reference configuration rather than a completed institutional deployment. It is thought as a small-scale, fully working, television station that lets students move through the full content lifecycle of a broadcast facility, from capture to distribution. The idea behind it is simple, students cannot learn this field only by watching professionals work. They need to operate the equipment, configure the systems, handle failures, and build or operate end-to-end workflows in conditions that are close to those of a real newsroom [7]. The proposed center combines professional broadcast equipment, general-purpose IT infrastructure, and software tools. Some of these tools are commercial broadcast applications, while others are developed specifically for the learning center.

2.2 Architectural Components

The architecture is organized into four main components, each of which corresponds to a distinct stage of the content lifecycle [8], [9]:

- The media acquisition component is responsible for the ingestion of media assets from various sources into the central media management system. It constitutes the input boundary of the proposed architecture;
- The processing and planning component is the operational engine where media assets are edited, scripted, organized, planned, and prepared for for broadcast. It forms the core of the Media Asset Management (MAM) system;
- The production component covers live event coverage based on the processing and planning phase. It transforms the output of the previous component into a ready to air broadcast signal;
- The distribution component delivers the finished product to downstream channels, including cable operators and streaming platforms. It represents the output boundary of the proposed architecture;

Figure 1 provides an overview of the proposed architecture and indicates which practical learning scenarios, defined in Section 4, are associated with each of the four main components.

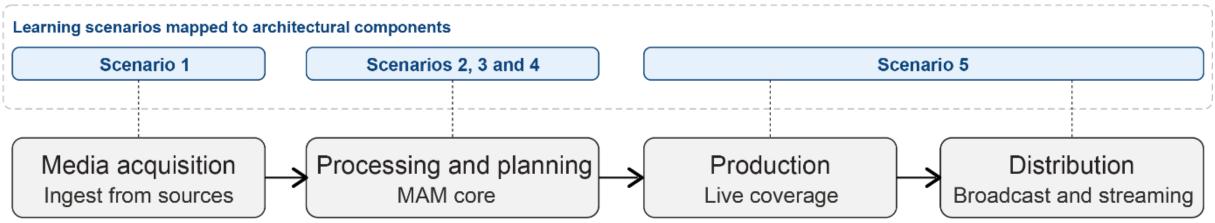

Fig. 1. Architecture overview

Underlying all four architectural components is an IT infrastructure layer composed of routers, switches, domain controllers, FTP servers, and virtual machines. Together with dedicated video signal links, this layer provides the connectivity required for the components to communicate and operate as a coherent system.
The remainder of this paper focuses on the architecture of each of these components and on the educational scenarios they support.

3 Architecture of the Learning Center

This section describes the technical configuration of the four architectural components. Two original subsystems are integrated where appropriate: a cloud-based automated ingest subsystem in the media acquisition component, and a supervised machine-learning subsystem for news classification in the processing and planning component.

3.1 Media Acquisition

The media acquisition component, commonly referred to as the ingest stage in broadcast operations, is responsible for capturing media content from a variety of sources, transcoding it to the format required internally, and storing it into the Media Asset Management (MAM) system [10]. The component must support the simultaneous handling of multiple streams and the real-time import of incoming video and audio material. Transcoding would be performed either manually, through operator-driven workflows, or automatically by custom scripts based on the FFmpeg open-source toolkit [11]. A user interface allows operators to control the acquisition process and to append each asset with descriptive metadata, which becomes essential for later retrieval and editorial use.

3.1.1 Acquisition Sources and Infrastructure

The acquisition component aggregates content from several distinct source categories, each with its own technical interface and operational pattern.
Field cameras and mobile phones constitute the primary capture devices. The learning center proposed architecture places particular emphasis on smartphone-based capture, both because of its growing role in professional newsrooms and because of its accessibility for educational use. Modern smartphones incorporate optical and computational improvements such as multi-lens systems, high dynamic range capture, image stabilization, and AI-assisted image processing for noise reduction, color correction, and object-aware focusing [12]. These devices offer a favorable cost-to-quality ratio and enable rapid deployment to remote events, which makes them well suited to the agile workflows associated with mobile journalism [13], [14].
Another source for asset acquisition are IP video streams from news agencies that provide a continuous flow of pre-recorded and live international footage. In the learning center, these streams are converted into SDI (Serial Digital Interface) video signals through dedicated decoders and routed to the central video router as standard broadcast sources.
A cloud-based platform supports the storage and distribution of audiovisual material exchanged with external partners. The operators can retreive shared video assets from the cloud, transcode them and manually ingest into the main storage.
Viewer-generated content, transmitted over

instant messaging services or other online channels, complements editorial material [15], [16]. A configurable workflow downloads such submissions, transcodes them to the internal format, and inserts them into the central storage with minimal operator intervention.
Finally, a video archive, implemented as a network-attached storage system, hosts prerecorded material that can be reused as visual support during story composition. Bidirectional workflows between the archive and the MAM system allow students to export or import media assets.
The hardware infrastructure underlying these sources comprises a set of ingest servers, which record incoming streams, and transcoding servers, which convert assets to the format required by downstream stages. Figure 2 shows how these sources and the supporting infrastructure are interconnected within the media acquisition component.

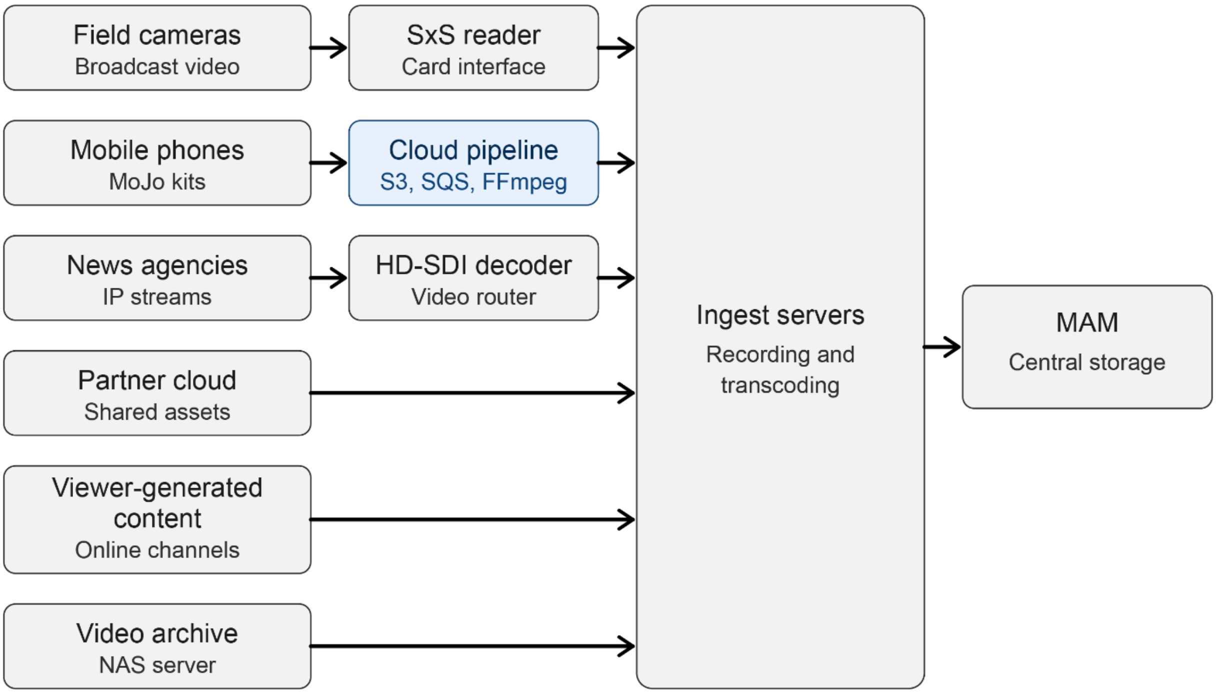

Fig. 2. Media acquisition

3.1.2 Cloud-Based Automated Ingest Subsystem
Conventional ingest pipelines remain heavily dependent on human operators. Recording devices must be carried to the facility, files copied through wired or removable connections, transcoding launched manually, and MAM imports queued for processing. The resulting latency is problematic whenever the editorial value of a video clip depends on its rapid availability, as in breaking news coverage [17], [18]. To familiarize students with a modern alternative and to support practical experimentation with cloud-native automation, the learning center embeds an automated ingest subsystem optimized for mobile journalism content. The subsystem was designed and evaluated in prior work by the authors [5].
The design combines a public cloud entry point with on-premise processing resources in a hybrid configuration [19], [20]. Reporters in the field record material on a smartphone and use a dedicated mobile application to push the selected clip to a cloud object storage bucket over cellular or Wi-Fi connectivity. Every successful upload generates an event notification routed to a managed message queue in the cloud. An on-premise listener subscribes to this queue, retrieves new notifications as they arrive, and pulls the corresponding file into a local staging directory. A second on-premise component watches that local directory, detects when the download has finalized, applies a transcoding profile aligned with the broadcast format expected by the MAM, and places the result on a network share that the MAM

monitors. A media migration policy then imports the asset automatically, making it ready for editorial use.
Once the reporter selects a clip in the application, the chain runs end-to-end without further intervention. Field measurements conducted in an operational news television station showed that the subsystem makes material available in the MAM roughly three times faster than alternative cloud-based workflows that still rely on manual sharing or operator-driven import, and up to fifteen times faster than conventional workflows that require returning the recording device to the facility [5]. The subsystem also removes the need for a dedicated ingest operator at this stage. For students, it provides a concrete setting in which to examine how cloud storage, event notifications, managed message queues, and local watch folder automation can be combined to replace a process that has traditionally been treated as manual. Beyond the technical aspects, an intended outcome is that students recognize the value of improving established workflows by redesigning them for efficiency. Figure 3 illustrates the progression of a media asset through the subsystem.

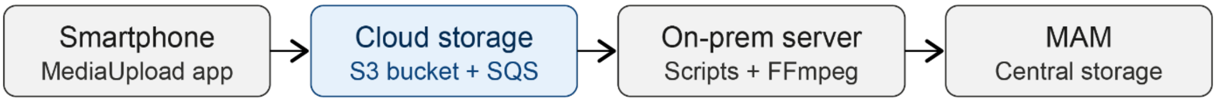

Fig. 3. Cloud-based automated ingest subsystem

3.2 Processing and Planning

The processing and planning component is the operational core of the Media Asset Management system (MAM), centralizing and coordinating the flow of every story from acquisition to distribution. At this stage, students access the raw footage produced during ingest and operate a full range of non-linear video editing tools for cutting, trimming, and color correction, as well as for adding text, graphics, and sound effects. They also write the narrative script of the story and the corresponding script for the news anchor prompter. All content is stored in a database cluster configured for redundancy and is orchestrated by a dedicated software application known as the Newsroom Computer System (NRCS) [21], which manages content creation, content management, real-time collaboration, planning, and automation. In the proposed architecture, the NRCS provides the main operational interface for story creation, metadata management, planning, collaboration, and automation.

3.2.1 Core Infrastructure

From a technical standpoint, the processing and planning stage is the most complex component of the learning center and relies on several interconnected subsystems that can be grouped into a data layer and a services layer. The data layer should comprise of a central storage facility for media assets, built from storage controllers and one or more enclosures of hard disks reaching a total capacity of hundreds of terabytes or even several petabytes. Alongside this, a main database cluster serves as the central repository for all media assets and their associated metadata, allowing users to search, locate, and manage assets, while also tracking their status throughout the production pipeline, from acquisition through post-production to distribution. Media files are not stored directly in the database but referenced through pointers to the actual content on the central storage [9], and role-based access control ensures that only authorized users can view or modify the media assets. Graphical content is handled separately by a dedicated graphics database, which stores templates and visual elements, together with the graphics design system used to create them.
The services layer is organized around several specialized categories of servers. Core servers run dozens of agents responsible for authentication and security, search, printing, import and export workflows, and the job broker that coordinates every feature of the system. Control servers act as a gateway between the MAM and external devices such as the studio automation system, the teleprompter, and the

graphic engines, while migration servers handle media migration policies used to import/export media or move assets through the system. Rendering servers combine multiple visual and audio elements into a final video product that can be used in the production phase. This process can be time-consuming, particularly for longer sequences or clips involving complex visual effects.
All of the above components are exposed to the user through the NRCS, which functions as the graphical user interface of the media asset management platform and integrates the underlying infrastructure into a coherent production environment.
Figure 4 illustrates the main components of the processing and planning technical infrastructure, with the services layer on the left and the data layer on the right.

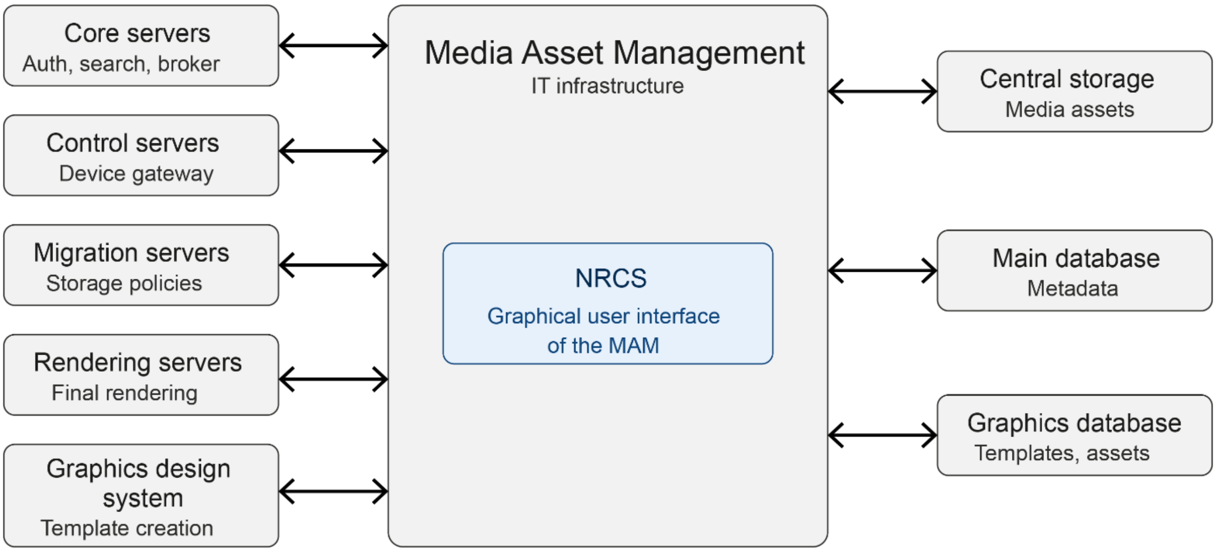

Fig. 4. Processing and planning diagram

3.2.2 Automated News Classification Subsystem

Within the processing and planning component, assigning editorial categories such as politics, economy, sports, or international affairs to each story is a routine but important task. The resulting category labels inform broader newsroom decisions: assessing newsroom performance, allocating editorial resources across desks and topics, shaping editorial content strategy, and contributing to audience insight. In conventional newsrooms this work is performed manually by journalists, which is costly in time, prone to inconsistency between editors, and difficult to scale as the archive grows. To address this and to give students a concrete entry point into applied machine learning, the learning center incorporates an automated news classification subsystem that automatically assigns editorial categories to completed stories. The subsystem was developed and evaluated in earlier work by the authors [6].
The classifier was trained on a purpose-built corpus of around 24,000 labeled Romanian news stories, collected from the MAM of an operational news television station over a fifteen-month interval and distributed across fifteen editorial categories [6]. Several text vectorization techniques and a representative selection of supervised classifiers are evaluated in combination [22], [23], [24]. The dataset is split 80:20 into training and testing partitions, and a balancing step based on random oversampling is applied to compensate for the marked skew of the editorial label distribution [25].
Empirical evaluation on this dataset shows that the balancing step has a strong and uniform effect on performance. On the original imbalanced data, accuracy varies considerably across configurations: the strongest combination, support vector machines paired with tf-idf, reaches around 0.80, while several oth-

ers fall well below that level. After balancing, accuracy rises across the board, and Random Forest and the Multi-Layer Perceptron approach 0.99 on accuracy, precision, recall, and F1-score under every vectorizer tested [6]. The differences between vectorizers, which matter visibly on imbalanced data, become marginal on the balanced corpus.

In the proposed integration, the trained classifier would be deployed as an editorial reporting application. As stories are finalized in the NRCS during the week, each is automatically assigned an editorial category by the classifier and appended to a weekly report listing how many stories were created in each category. Given the high accuracy reported in [6], this report can be treated as a reliable basis for decision-making. The weekly view of editorial output supports newsroom performance assessment, informs the allocation of editorial resources across desks and topics, feeds into content strategy decisions, and contributes to a more structured understanding of audience-facing output. For students, the subsystem provides a self-contained applied machine-learning environment in which they can examine how the dataset was assembled, retrain models under different feature extraction and balancing choices, and observe how those choices propagate into operational behavior. Beyond the technical mechanics, the subsystem is intended to make visible to students how decisions in the data pipeline, and class balancing in particular, can have a larger effect on real-world performance than the choice of model itself.

3.3 Production

The production platform contains two closely integrated environments, the studio set and the production control room (PCR). The studio set constitutes the physical backdrop for television production. Within the proposed architecture of the learning center, the studio set is specified to include a robotic camera, a microphone, décor elements and displays, lighting fixtures, and teleprompter monitors, all configured to replicate a professional broadcast environment.

The production control room is responsible both for live production and for recording programs or shows for later use. It receives the plan prepared during the processing and planning stage and outputs a finished program ready for broadcast. At the heart of this environment is the studio automation server, which functions as the core of the entire production workflow. The director operates through the automation system, which in turn coordinates the specialized subsystems that together generate the final broadcast output.

Through the automation, the director triggers the video switcher that selects between camera feeds and other video sources for the broadcast output, and cues the graphic engine to render on-screen graphics in real time, including lower thirds, location tags, and other visual elements required by the production. The automation also manages the audio mixing console, governing the levels and distribution of audio signals throughout the program. It operates the robotic camera to control motion parameters, such as zoom, focus, iris, pan, and tilt, and to recall stored image presets, enabling consistent camera behavior throughout a production. Lighting is driven through the lighting console, which uses the DMX512 (Digital Multiplex with 512 pieces of information) standard to configure the studio lighting for optimal visual impact and atmosphere. The automation further orchestrates the playout servers responsible for storing and delivering pre-recorded video and audio content, and commands the video processors that feed visual content to the monitors embedded in the studio set décor.

The studio automation server receives its instructions from the processing and planning component of the MAM system. It ingests the playlist prepared in the NRCS in the planning phase through the Media Object Server (MOS) protocol [26], and executes each cue accordingly, while the director supervises the execution and ensures that the broadcast follows the original plan. This centralized coordination streamlines the production process, reduces human error, and improves overall operational efficiency.

Alongside the automation system, several additional subsystems support signal distribu-

tion, monitoring, and quality control. A video router matrix allows video signals to be distributed and switched between the various sources and destinations within the platform. Multiviewer monitors display multiple video sources simultaneously on a single screen, providing the director and technical staff with a comprehensive view of all the feeds being used in the production. Signal quality is continuously verified through a waveform monitor, an instrument for measuring and displaying the characteristics of video signals.
The platform also integrates components for external live contributions and on-set communication. Live transmissions video streams can be sent over the internet directly to the production control room and allowing remote reporters to contribute live material to the program. Interruptible Foldback (IFB) lines provide a private communication channel between the control room and the production crew without the audience hearing the exchange. The teleprompter system displays scripted text to presenters, enabling them to read from prepared material while maintaining eye contact with the camera. All components in the production control room are interconnected via Ethernet interfaces, which ensure reliable communication between the automation server and each of the controlled subsystems. Figure 5 illustrates the production studio automation system workflow.

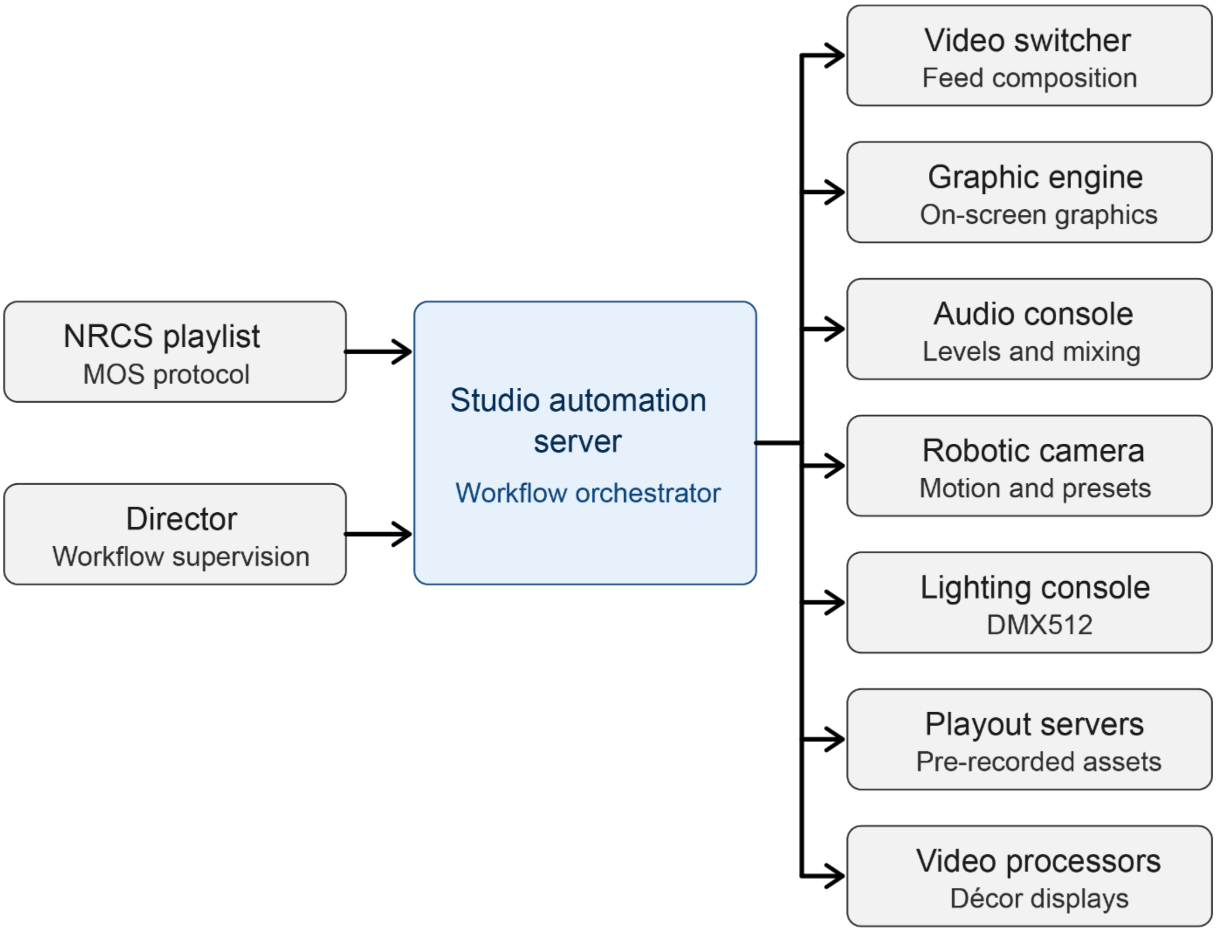

Fig. 5. Production studio automation workflow

3.4 Distribution

The distribution platform represents the last stage of the broadcasting process, responsible for delivering the final product to viewers through a variety of channels, including cable TV companies, video streaming platforms and social media platforms. The proposed distribution component is organized around the Master Control Room (MCR), the central hub where all content is received, managed, and distributed for broadcast, and where the final output is verified against the required quality

and technical standards. The MCR includes local asset handling functions for video media assets playout and graphical elements such as the station logo or other branding elements, but it remains distinct from the central MAM.
The distribution platform integrates several technical components that operate in a coordinated manner. The starting point is the scheduling system, a software application that enables the creation of a programming schedule for a defined period of time, typically a week or a month. It supports the selection of programs to air, the ordering of their broadcast, and the assignment of time slots, while also maintaining a database of program metadata such as event titles, descriptions, and ratings. For live or simulated-live productions, the production control room (PCR) output enters the MCR as a live source and is passed through the playout/distribution chain according to the active schedule. The scheduling database runs on a dedicated server that is separate from the core MAM infrastructure.
Once the schedule has been defined, the playout servers take over the transmission process by automatically playing out the scheduled content according to the playlist generated by the scheduling system. In addition to storing a large library of video media assets, these servers can also pass through the signal received from the production control room during live productions, ensuring a seamless transition between pre-recorded and live programming. Content is not limited to locally produced or archived material, as a FTP server is used to receive video assets over the internet from external partners and contributors. A dedicated compliance recorder is used to capture and store the final broadcast output in real time for regulatory and legal purposes, retaining it for a legally defined period so that it remains available in case of legal disputes.
The final stage of the distribution chain consists of real-time IP-based transcoders, which receive the video signal from the playout servers and encode it into an RTMP (Real-Time Messaging Protocol) stream that is sent out to the viewers. This represents the last technical component of the broadcaster, marking the point at which the media product leaves the facility and reaches the audience. Figure 6 illustrates the diagram of the distribution platform.

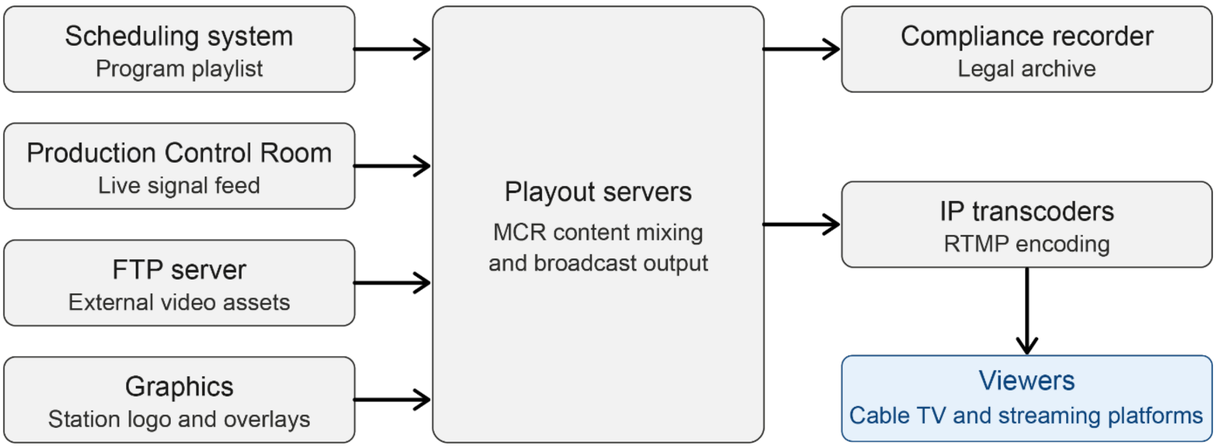

Fig. 6. Distribution platform

4 Learning scenarios
The architecture described above supports a structured progression of learning activities that mirrors the lifecycle of a news story in an operational broadcast facility. Five scenarios are presented below that are designed to cover all four architectural components. The first three follow the natural sequence from acquisition to packaging, the fourth focuses on the educational use of the news classification subsystem, and the fifth closes the loop by taking the packaged segment through production and distribution to on-air delivery.

4.1 Scenario 1: Mobile Journalism with Cloud-Based Automated Ingest

The first scenario introduces students to the practical realities of newsgathering and to the cloud-based ingest subsystem described in Section 3.1.2. Students are dispatched to a remote location and asked to capture raw footage using a mobile phone equipped with the mobile journalism kit. They learn the capture techniques associated with this filming style, including framing, audio capture, and stabilization, all of which have become widely adopted in modern newsrooms because of their accessibility, low cost, and rapid deployment [17].

Once the footage is recorded, students use the dedicated mobile application installed on the device to upload the selected clips to the cloud. The upload is performed over a cellular data connection, which simulates the operational conditions of a reporter working in the field. From the moment the upload completes, the rest of the pipeline operates without manual intervention: the cloud storage service emits an event notification, an on-premise listener downloads the file to a staging area, the watch folder script transcodes it into the broadcast format, and the MAM ingests the resulting asset through its migration policy. Students follow the propagation of the file along this chain by inspecting the cloud console, the staging directory, and the MAM interface, which gives them a concrete view of how cloud event notifications, message queues, and local automation interact.

In parallel with this main exercise, students explore complementary acquisition channels. They retrieve material from news agencies and from the web, configure transcoding workflows, and simulate live transmissions by combining a remote team operating in the field with an indoor team operating the ingest stations. For every asset that enters the MAM, students provide structured metadata that will be required by the subsequent stages.

This scenario is the foundation of the entire learning sequence and is treated as a prerequisite for the activities that follow. Assessment focuses on whether students have successfully imported usable assets into the MAM, whether the metadata they associated with each asset is coherent and complete, and whether they can describe the role of each stage of the cloud-based ingest chain.

4.2 Scenario 2: Processing the raw media assets

The second scenario builds on the assets ingested during the first scenario and is focused on the video and audio construction of a coherent news segment. Students start a new editing project whose scope is to refine the raw media captured in the first scenario that explores the media acquisition phase.

Within this scope, students first review the entire body of available footage in order to internalize its content and to identify the segments that contribute most directly to the story being constructed. They then trim and assemble the selected clips, removing redundant or irrelevant material and ensuring that each segment is of an appropriate duration, usually one to two minutes. A voice-over track may be recorded and added directly within the editing project at this stage. Students arrange the timeline clips in narrative order and apply transitions, such as fades or dissolves, to smooth the visual flow between scenes. They balance audio levels across the timeline, add ambient sound effects or background music where appropriate, and apply color correction so that the visual style remains consistent across the assembled segment.

The output of this scenario is a self-contained audiovisual sequence: a timeline with stabilized pacing, consistent audio levels, and a coherent visual treatment. The assessment focuses on whether the resulting sequence is cohesive, whether its narrative progression is clear, and whether its technical quality, in both audio and video, is suitable for downstream packaging.

4.3 Scenario 3: Newsroom Integration, Scripting, and Rundown Planning

The third scenario picks up the audiovisual sequence produced in Scenario 2 and integrates it into the newsroom environment. The focus of this scenario is no longer timeline editing but editorial packaging inside the

NRCS: scripting, graphic insertion, prompter preparation, and alignment with the on-air rundown playlist (planning). It is in this scenario that students experience the handoff between editing tools and the broader production environment, which is the boundary that Scenario 2 stops at.
Students begin by drafting a script that frames the visual material with narration or on-camera reporting. The script is authored within the NRCS and is simultaneously prepared for the teleprompter, so that timing, tone, and vocabulary are aligned with what the presenter will read on air.
Once the script is in place, students add the graphical elements that transform the sequence into a broadcast-ready segment: lower thirds that show the story title, full-screen graphics that provide context, and any additional visual cues. Because these elements are driven from the graphics database and rendered by the graphic engine during production, students also configure the cues for the automation system through which the segment will later request each graphic at the correct moment.
Finally, students integrate the segment into the rundown, assigning its position within the program, declaring its durations and cue points, and verifying that metadata, script, graphics, and video are coherent across the NRCS. The output of this scenario is a rundown-ready item, approved for handoff to the production control room. Assessment verifies whether the segment is editorially complete, whether its NRCS representation is internally consistent, and whether it is in a state that would allow the production stage to take it to air without additional editorial intervention.

4.4 Scenario 4: Automated News Classification in the Editorial Workflow

The fourth scenario introduces students to the news classification subsystem described in Section 3.2.2 and to its role in everyday editorial work. The scenario has both an applied and an analytical dimension. On the applied side, students operate the classifier as an editorial aid. On the analytical side, students examine the data and the models that underpin it, developing an understanding of how datasets are constructed and how supervised machine learning systems are trained and evaluated.
The applied component begins with the routine activity of authoring stories in the NRCS. As students complete the body of each story, the classifier returns a suggested editorial category. Students review the label assigned and they record the cases in which the classifier's recommendation diverges from their own judgment.
The analytical component takes place in a controlled experimental environment in which students work directly with the underlying dataset. They learn how the dataset was constructed and they examine the class distribution and identify the imbalance that is characteristic of editorial corpus, in which a few categories dominate the volume of stories while others remain comparatively rare. They then experiment with different feature extraction techniques, retrain a subset of supervised classifiers, and compare their performance using accuracy, precision, recall, and F1-score. They repeat the experiments after applying random oversampling to the training data and observe the potential improvement in performance, especially for the previously underrepresented categories.
Once the experimental phase is complete, students discuss which model–vectorizer combinations would be most appropriate for production deployment and how the classifier should interact with the planning interface of the NRCS.
The assessment of this scenario combines two perspectives. The applied perspective verifies whether students used the classifier coherently within the editorial workflow and whether their decisions to analyze suggestions were well reasoned. The analytical perspective verifies whether students can describe the experimental setup, interpret the resulting metrics, and motivate the choice of a particular configuration for production use.

4.5 Scenario 5: Production Control Room Operation and On-Air Delivery

The fifth scenario closes the loop of the content lifecycle by taking everything produced during the first two phases and turning it into a valid broadcast product. Where the first three scenarios exercise acquisition and processing and planning in depth, this scenario covers the remaining two components and gives students direct experience of how an editorial plan is transformed into a broadcast signal and then delivered to viewers.

Students operate the production control room automation system described in Section 3.3. The rundown assembled during Scenario 3 is transferred to the studio automation server through the MOS protocol [26] and becomes the operational playlist of the PCR. One student acts as director and drives the automation, while others monitor the controlled subsystems, including the audio mixing console, the graphic engine cues, the lighting console, and the robotic camera presets. The exercise is framed either as a recorded program or as a simulated live broadcast, and students are expected to recognize and recover from realistic faults such as late cues, missing graphics, or misaligned audio levels.

Once the production is set up, students follow the resulting program into the distribution chain described in Section 3.4. They configure the scheduling system to place the program at a defined time slot and verify that the playout server plays out the scheduled content correctly. They then route the feed to the encoder, which outputs the program to the YouTube streaming platform as a live event [27]. Students monitor this live event on YouTube as the final output of the entire chain, observing the program as it would be received by an external viewer.

The assessment of this scenario verifies whether students can run a complete production, end-to-end, without major editorial or technical failures and whether their recovery from faults is efficient. Then, we verify whether students can describe the full signal path from the rundown in the NRCS, through the automation cues in the PCR, into the distribution playout server, and out through the encoder to the YouTube live event.

5 Conclusions and Future Work

This paper described a reference architecture for a learning center supporting engineering programs in broadcast and media production. The architecture is organized around four components that follow the content lifecycle from acquisition to distribution and is intended to give students hands-on access to the equipment and workflows used in real broadcast facilities. Two subsystems, previously evaluated in separate work, were integrated into this architecture in order to expose students to current technological trends: a cloud-based automated ingest subsystem, placed inside the media acquisition component, which reduces manual steps in the file-transfer pipeline and accelerates the availability of mobile journalism material, and a supervised machine-learning subsystem for the classification of Romanian news stories, placed inside the processing and planning component, which supports newsroom organization and also serves as a practical applied machine-learning environment. The architectural novelty of the paper is concentrated in the acquisition and processing and planning components, whereas production and distribution are described with established broadcast building blocks so that the learning scenarios can run in realistic operational context and give students an overview of how a broadcast and media production center functions.

Five practical learning scenarios were defined to operationalize the architecture across its four components. The first scenario aligns the acquisition activity with the cloud-based ingest subsystem, and the fourth scenario engages students with both the practical and the analytical dimensions of the classification subsystem. The second and third scenarios cover editorial construction and newsroom integration respectively, now with a clearer boundary between timeline-level editing and NRCS-level packaging. A fifth scenario takes the resulting segment through the production control room and the distribution chain, so that the pedagogical coverage reaches all four architectural components and offers a very important insight of how raw media passes through all the stages of a broadcast media center.

The present paper proposes a reference design that integrates prior technical results into an educational framework, with the development of student competencies for the broadcast industry positioned as an intended design outcome of the proposed platform. A systematic study of learning outcomes, including the measurement of skill acquisition and course-level assessment results, remains part of planned future work. Also, the educational impact of the framework is presented as an intended direction for evaluation in future work. In parallel, two technical directions are also planned. The first is further automation of internal workflows in each architectural component, with particular attention to operations that still rely on human intervention. The second is proposing an architecture that migrates selected systems from the local datacenter to cloud platforms, in order to support scaling on demand and remote learning.

References

[1] Nicolaou, Constantinos, Maria Matsiola, and George Kalliris. "Technology-enhanced learning and teaching methodologies through audiovisual media." *Education Sciences* 9.3 (2019): 196.

[2] Reyes-de-Cozar, Salvador, Marta Perez-Escolar, and Pablo Navazo-Ostua. "Digital Competencies for New Journalistic Work in Media Outlets." *Media and communication* 10.1 (2022): 27-42.

[3] St Clair, Jeanti. "Doing it for real: Designing experiential journalism curricula that prepare students for the new and uncertain world of journalism work." *Coolabah* 16 (2015): 122-142.

[4] Tembrevilla, Gerald, André Phillion, and Melec Zeadin. "Experiential learning in engineering education: A systematic literature review." *Journal of Engineering Education* 113.1 (2024): 195-218.

[5] Vintilă, Adrian, and Constanța-Nicoleta BODEA. "A Framework for Automated Digital Media Asset Acquisition using Cloud." *Informatica Economica* 29.1 (2025).

[6] Vintila, Adrian, and Constanta-Nicoleta Bodea. "Romanian News Article Classification: A Multi-Model Comparison with Class Balancing." *International Conference on Informatics in Economy*. Singapore: Springer Nature Singapore, 2025.

[7] Abele, Eberhard, et al. "Learning factories for research, education, and training." *Procedia CiRp* 32 (2015): 1-6.

[8] European Broadcasting Union, The Digital Media Value Chain: Modelling Core Business Objects and Processes in Digital Media Enterprises. EBU Tech Report 041, version 2.0, 2024.

[9] Heitmann, J. "Content Management Systems for television production." *EBU Technical Review* (1999): 24-34.

[10] Wager, Skiff. "Ingest, manage and distribute." *Journal of Digital Asset Management* 1.3 (2005): 157-163.

[11] Subhash, V. *Quick Start Guide to FFmpeg: Learn to Use the Open Source Multimedia-processing Tool Like a Pro*. Apress, 2023.

[12] Morikawa, Chamin, et al. "Image and video processing on mobile devices: a survey." *The Visual Computer* 37.12 (2021): 2931-2949.

[13] Westlund, Oscar. "Mobile news: A review and model of journalism in an age of mobile media." *Digital journalism* 1.1 (2013): 6-26.

[14] Burum, Ivo, and Stephen Quinn. *MOJO: The mobile journalism handbook: How to make broadcast videos with an iPhone or iPad*. Routledge, 2015.

[15] Colistra, Rita, Eva Buchman, and Kevin Duvall. "User-generated content and television news stations: Perceptions, policies, and uses." *Electronic News* 11.3 (2017): 123-144.

[16] Tolmie, Peter, et al. "Supporting the use of user generated content in journalistic practice." *Proceedings of the 2017 chi conference on human factors in computing systems*. 2017.

[17] Rodrigues, Luis Pedro Ribeiro, Vania Baldi, and Adelino de Castro Oliveira Simões Gala. "Mobile Journalism: the emergence of a new field of journalism." *Brazilian journalism research* 17.2 (2021): 280-305.

[18] Usher, Nikki. "Breaking news production processes in US metropolitan newspapers: Immediacy and journalistic authority." *Journalism* 19.1 (2018): 21-36.

[19] Deb, Moumita, and Abantika Choudhury. "Hybrid cloud: A new paradigm in cloud computing." *Machine learning techniques and analytics for cloud security* (2021): 1-23.

[20] Liu, Fang, et al. "NIST cloud computing reference architecture." *NIST special publication* 500.2011 (2011): 1-28.

[21] Schultz, Fred. "Automation and live television news: Enhanced support for a complex workflow." *SMPTE motion imaging journal* 112.1 (2003): 29-35.

[22] Gasparetto, Andrea, et al. "A survey on text classification algorithms: From text to predictions." *Information* 13.2 (2022): 83.

[23] Kamath, Cannannore Nidhi, Syed Saqib Bukhari, and Andreas Dengel. "Comparative study between traditional machine learning and deep learning approaches for text classification." *Proceedings of the ACM Symposium on Document Engineering 2018*. 2018.

[24] Kowsari, Kamran, et al. "Text classification algorithms: A survey." *Information* 10.4 (2019): 150.

[25] Wongvorachan, Tarid, Surina He, and Okan Bulut. "A comparison of undersampling, oversampling, and SMOTE methods for dealing with imbalanced classification in educational data mining." *Information* 14.1 (2023): 54.

[26] MOS Protocol Working Group, MOS Protocol 4.0. Official protocol documentation. https://mosprotocol.com/

[27] Google/YouTube Help, "Choose live encoder settings, bitrates, and resolutions,", https://support.google.com/youtube/answer/2853702.

Adrian VINTILĂ has graduated the National University of Science and Technology Politehnica Bucharest. He is currently a PhD student at the Bucharest University of Economic Studies. With a strong background in the field of information technology and broadcast engineering, his research focuses on intelligent systems for innovation in media and television production within educational centers.

Constanta-Nicoleta BODEA is Professor Emerita at Bucharest University of Economic Studies, Romania, holding a Ph.D. in Economic Cybernetics and Statistics. She is Senior Researcher at the Centre for Industrial and Services Economics, Romanian Academy of Sciences, and contributes to the Doctoral Study in Project Management at Alma Mater Europaea University, Slovenia. Her research focuses on project management, service innovation, and technology-enhanced teaching and learning. She published more than 50 books and book chapters, and 370 papers in journals and conferences.