Universal production workflow

With MVR

Casting

- The Multicam Video Recorder ..... MVR
- The Media Production Manager..... MONITOR
- A Non Linear Editor (alias NLE)..... XEDIT
- A Metadata Source (text, scores, sensors, content mining) ..... XDATA
- A render node (rendering, encoding, transcoding) ..... XRENDER
- An output server (streaming, VOD, playout)..... XPLAY

XEDIT, XDATA, XRENDER and XPLAY are generic workstations for respectively editing, data ingest, processing and content output. Multiple actors may play in a given workflow. Each module may run remotely with a temporary connection.

XEDIT can be a “rough” edit for journalists, a remote edit via a low-res proxy, or a high-end edit for finishing of sophisticated contents.

XDATA is a client workstation with data inputs from operator, or automation. It may come from different vendors for virtual imaging, scores, audio mining or scene detection. Data may come from various phases of the production (before live, during live, or after capture).

XRENDER and XPLAY are generic processing nodes. For live analog broadcasts XPLAY could be a standard broadcast playout server.

Depending on required connectivity and processing power, a single processing node may host more than one client.

The Big Picture

Universal production model

The production model is based on a pure IP-based network. Gigabit Ethernet is used as the current standard and 10 Gigabit is foreseen. This model allows increased collaboration: ingest systems, logging & auxiliary data inputs, editing applications and content production can all work together simultaneously for faster over-all production.

Multiple simultaneous outputs are targeted:
- Traditional broadcast "to air" (analog, DTV, DVB, DVB-H…)
- Streaming (DSL, DVB-H…)
- DVD authoring, VOD download (DSL, DVB-H…)
- Virtual imaging feeds (electronic billboards, sponsoring, editorial)
- Coaching, Electronic referee, Analysis
- Security, Commissar's playback (races)
- Archive, Judicial recording

These outputs may require:
- Editing, DVE, Slow-motion
- Format change & transcoding
- DRM & watermark
- Virtual imaging embedded metadata (personalization, hyper video)
- Other metadata for editorial enhancements, statistics, links
- Etc.

Including, but not limited to

The list above includes **non-broadcast** applications, and the workflow welcomes non-broadcast video sources. Additional cameras are required for specific applications, that may need the slow-motion replay (commissars reviewing a race) and the archiving (judicial recording).

In the same way, the director may choose a non-broadcast source for a special editorial need (for instance, a surveillance camera providing an **original image** not covered by the production van.

In theory, any combination of the "layers" listed above is usable. For instance, we may wish to edit the live program again, or add "late" metadata (statistics) to a downloadable content.

In practice, Multicast of MVR multiplex may allow unlimited parallel processing (unreachable with centralized storage).

IP Multicast & duplexing

Multicast is a key feature of IP networks (called "broadcast" and available since the early days): one data packet is sent, and **many receivers** can read it at the same time.

Multicast allows duplexing of signals, exactly like **video distribution**, with no additional cost. Hence, it is much easier to have **processing nodes grabbing their own content**, and creating their **own local storage**, rather than trying to write once and read many times from a critical central storage.

Another advantage of duplexing is the **natural redundancy** of storage available throughout the network. This redundancy is known by the MONITOR which can use the **local resource** when distributed tasks are spread over the network (see "farming" below).

In addition, each node may pick-up only the required elementary streams (video, audio or data) from the MVR Multiplex.

Live switching … from a downstream location

Live switching based on compressed signals has been described by LVS since 2004; it may run on a single dedicated processing node for basic edit (CUT). The live switcher is **splicing** directly the MVR Multiplex, and does not need local storage.

Remote live switching is accessible (like remote-edit in post-production), with no visible delay for the director at a downstream location ("time machine production"). Nevertheless this workflow requires some **cooperation from the infrastructures** (two-way satellite link or cable with data path).

Farming: reconfigurable & fault tolerant

Farming is also used as a benefit of the distributed architecture for CPU intensive tasks. **Distributed processing** is in perfect harmony with the **distributed storage** born from duplexing (see above).

The resources are handled automatically by the MONITOR in nominal operation as well as in downgraded modes (node failure or unavailability).

Workflow "wiring"

Again, the analogy with "video cable" is usable to describe the signal routing.

The workflow is designed on MONITOR which interface offers:
- A **graph editor** for cabling workflow and handling profiles
- A **scheduler** to stick tasks to the timeline (record, processing, outputs)
- A **process controller** for handling priorities, load balancing, workstation failure
- An **SQL-based MAM server** responsible for managing feeds and metadata

The graph editor is directly inspired from **Microsoft Direct Show**, with some enhancements:
- **Networking**
- **Time management**
- **Live control** (runtime graph modification)

Conclusion

Our vision is quite different from the "high performance centralized storage" generally described in most workflows. This expensive centralized storage is generally coupled to distributed "low-cost" editing because there is no other solution than having multiple seats. The centralized storage needs to be secured, because everybody relies on it: redundancy is added to high performance and leads to the most expensive server solution.

We believe in **pure distributed processing**, based on proven **farming infrastructures** familiar to telecom operators. Pure distributed processing means **no server bottleneck** and offers at no cost **natural multiple redundancy**. It is based on the **live duplexing** of the source material as many times as needed by the workflow. Each node has its own "cache" on its own **lowest-cost local storage**, which can be shared with other nodes upon request from the MONITOR.

The capability for **live switching** in parallel with **any other post-production workflows** is a proof of the concept for MVR: a production "head-end" for the lowest-cost universal platform.