



By Andy Minuth and Daniele Siragusano

Introduction

The digitisation of post-production processes can be considered complete. But are pure digital reproductions of non-digital workflows desirable? Does this approach explore all the potential? How would you design a state-of-the-art broadcast post-production from scratch today?

The German broadcaster, SWR, faced precisely this challenge in 2021. It planned to move digital colour correction to new premises and completely overhaul it in the process. The requirement was not only of a technical

nature, but also to bring existing working methods up to date. We at FilmLight, the developer of the Baselight colour correction system, had the pleasure of supporting SWR in realising these ambitions.

In this three-part series of articles, we take a closer look at the different aspects of this project. In this part, part two, we are focusing on the redesign of the IT systems.



Online editorial workplace at SWR in Baden Baden. The creation of ergonomic workplaces was a priority when equipping the rooms.

Emergent systems

The radial structure of IT systems allows each system to interact directly with every other system. The central question for any multi-component system is: how can it be designed to be capable of more than the individual parts it consists of? As well as increasing functionality, you also want to increase resilience.

The pursuit of emergent systems - or systems where the components achieve more together than separately - is by no means new, but it is less understood as a design criterion in the context of post-production. The monolithic, linear workflow is still the most popular design in our post-production industry, but we can create novel systems that do not work linearly through better programmable integration and adaptability of the software components.

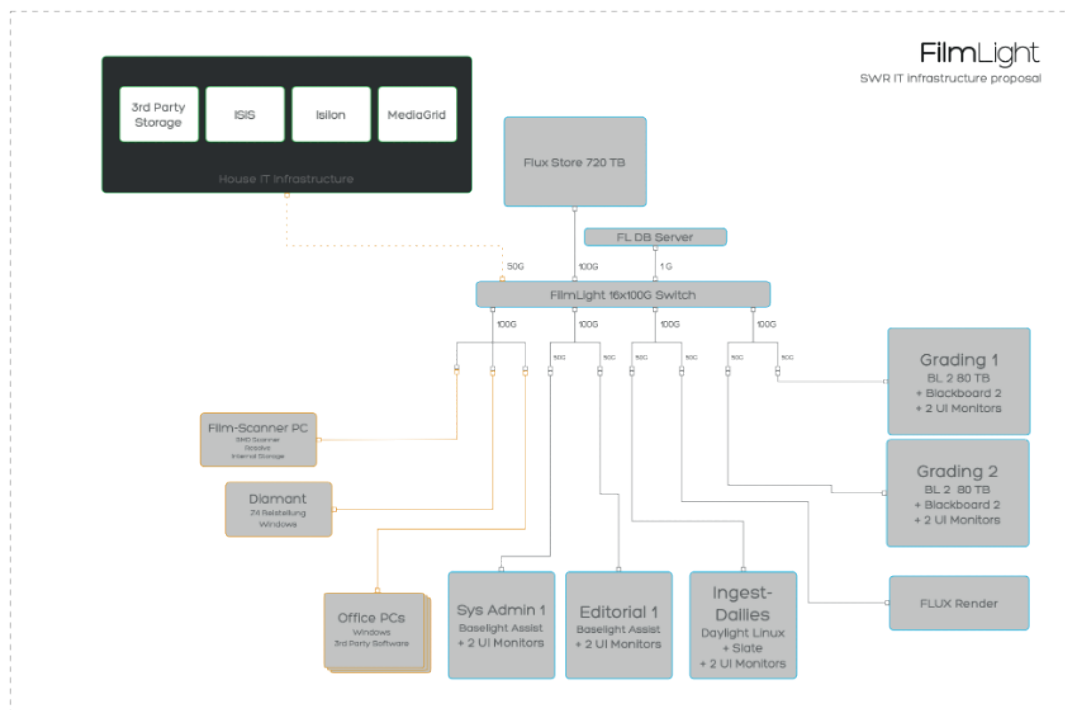
In the following, we provide some examples which illustrate how the integrated post-production system at SWR has the ingredients of this philosophy, and how further emergent effects can be created with existing hardware in the future.¹

IT hardware

Before we examine the application level, let's review the IT hardware. The hardware infrastructure is the backbone of modern post-production and the requirements at SWR were clearly defined. Roughly speaking, the requirement was to play slightly compressed UHD 50p material from central storage on all five systems simultaneously. In addition, each individual connection had to provide the bandwidth for UHD 50p 16bit uncompressed.

It was essential to SWR that we did not focus on nominal or lab measurements but on real performance at the application level.²

A performance-based specification of the IT hardware made sense to correctly measure the required components, such as the central storage and network infrastructure.



The central storage has a 100 GbE uplink. High-bandwidth clients are connected via 50 GbE. Low bandwidth clients are integrated into the network with 10 GbE and metadata clients with 1 GbE.

We won't go into a detailed description of the storage systems, IT infrastructure and network protocols in this article, as they are only a particular manifestation of the performance requirements. The question of what to do with such an infrastructure is much more exciting.

¹ We skip techniques relevant to non-post-production, like central user management or backup systems.

² Not by assessing the performance of the dd command from /dev/null to /dev/zero, but testing by playing the desired format in the actual application.

FilmLight Cloud

Each of the eight main workstations and servers at SWR are members of just such an emergent IT network, the FilmLight Cloud. The FilmLight Cloud is to be regarded as a complete system. Each member is aware of every other member's exact technological and service characteristics. This collective cluster provides new convenient services to the end users.

For example:

- » All FilmLight Cloud members see all storage volumes of the other clients. This is important to ensure the smooth running of distributed rendering and other processes. For example, hard drives attached to a FilmLight Cloud computer (e.g. for ingest) are immediately and automatically visible and usable as a volume on all FilmLight systems. The external hard drive is not mounted on a single computer, but on the collective cluster. This paradigm can often lead to irritation in internal IT departments as they have established a user-centred way of thinking. However, this user-centred view of the entire system is not necessarily efficient or practicable in the high-performance post-production world.³
- » If you want to copy media from one volume to another, the FilmLight Cloud dynamically decides the most efficient route for the copy operation. The copy operation is then scheduled and processed in the processing queue of the cloud member.
- » All media files of all cloud members are indexed, and metadata is extracted (more on this below).
- » Render jobs can be transferred from one FilmLight system to another, with the FilmLight Cloud automatically managing all volumes, paths and configurations.
- » All clients regularly run diagnostics to warn the user in good time of any problems that may occur. The diagnostics always cover the entire FilmLight Cloud, so every user on every system knows the current condition of the network.

Metadata indexing

Metadata has become almost as important as the actual media data. Metadata in post-production can automate many processes and reduce errors. For example, the user can trigger rule-based automation via metadata.⁴

Users can also use metadata to group or filter shots that belong together in a film scene. This can significantly speed up the operational process of colour correction. Generally, metadata provides important context, but due to the ongoing digitisation of film it is becoming more extensive. In addition to camera and lens data, increasingly more metadata from 'Script & Continuity' is reaching post-production without any additional effort.

There are many ways to store metadata and storing it in the file headers has become the established method. This makes it possible to assign metadata unambiguously. This is true for camera files and modern mezzanine formats such as OpenEXR.

However, storing metadata in this way does have some disadvantages. For single files, the metadata retrieval overhead is negligible. However, when several terabytes or petabytes of data is involved, reading out the metadata on demand becomes inefficient. Conforming an AAF/XML/EDL with 1000 events using metadata from a pool of several million files can take several hours, as the software must read each file partially.

This situation is considerably improved by the central storage of metadata in a database, for example, by a metadata indexer. The indexer opens all media that enters the FilmLight Cloud once (during ingest or copying) and stores the metadata in a database. Sophisticated algorithms then ensure that metadata and its data are always correctly referenced, even if the media files are moved. The metadata index can reduce the time for conforming to a few seconds.

Selection	
1 sequence (294 frames), 592.1 MB selected	
ColorGammaSxS	LOG-C
Dst	00:00
DynamicMetadataVersion	3
ExposureIndexAsa	800
EyeIndex	SINGLE
FramelineFileName1	ARRI 1.85
FramelineFileName2	Off
Height	1038
HorLengh	1920
ImageDenoising	0d
ImageOrientation	0d
LdsEncFocusMax	12982
LdsEncFocusMin	4990
LdsEnclrisMax	4685
LdsEnclrisMin	2985
LdsEncZoomMax	0
LdsEncZoomMin	0
Left	0
LensSerialNumber	8943872
LensType	Zeiss MP135 T1.3
MainVoltage	241
MasterSlave	OFF
MediaType	QuickTime Movie
Mirror/ShutterRunning	0d
Name	1.85:1_scaling100
NdFilterDensity	0
NdFilterType	0
OnBoardVoltage	165
Product	2
ProductionInfoCinematographer	MR_B
ProductionInfoCompany	ARRITEST
ProductionInfoDirector	MR_A
ProductionInfoLocation	BOTMUC
ProductionInfoOperator	MR_C
ProductionInfoProduction	ARRI
ProjectFps	24000
ReelName	S001R62J
SensorFps	24000
ShutterAngle	112
SoundReel	MUSIC
SubProduct	33
SupVersion	AlexaSXT_1.0:362
SxsSerialNumber	X512GJ1F370030
TempAcomBoard	300
TempAcomFpga	528
TempAlogBoard	496
TempAlogFpga	556
TempApicBoard	453

Metadata stored by a camera in the header of a media file. In modern workflows, processes can be automated using metadata. However, access to metadata can be slow when dealing with a large number of media files.

³ It should be mentioned that security aspects of user-specific permission management can be combined with emergent systems.

⁴ Such as via the Media Import Rules in Baselight.

The idea of a metadata index is not new. There are many systems that transfer metadata into their own format to allow faster access. What is unique about the FilmLight Cloud is that there is no direct 'translation' of the metadata, in the true sense of the word. In the FilmLight Cloud, the native metadata of the various camera and mezzanine files is read and referenced. The index is automatically updated if a user moves the original file to another location. As a result, this technology vanishes into the background of everyday work, and metadata-driven workflows establish themselves without effort.

However, the real production landscape does not only consist of the FilmLight Cloud. External storage with petabytes of data and metadata are distributed across various platforms. This plays a smaller role at SWR because the central storage is a native resident of the FilmLight Cloud. This means it is always aware when files are moved, or new files are copied to the storage.

A dedicated indexing service integrates the data and metadata of external storage into the FilmLight Cloud. The FilmLight Indexer's sole task is to index data and metadata from external storage and to serve metadata queries from FilmLight Cloud members. This happens completely transparently to the user. With the FilmLight Indexer, for example, the conforming time of a show with media spread across five Avid Nexis volumes⁵ could be reduced to less than five seconds.

Furthermore, since you can use the AvidUID as a reference key in the AAF, errors during conforming are eliminated. This means that if Avid can see the media, Baselight can see it too.

Adaptive systems

In addition to the positive emergence effects, modern IT systems should enable additional efficiency gains through adaptability. Below are a few concrete examples.

Template-based system

Setting up a project in post-production for the first time is a technically complex undertaking. Users must identify and document several hundred settings - colour spaces, scaling algorithms, media import rules, render pre-sets, metadata-based render paths for the output for different deliverables, filter groups for media management and much more.

These settings must also be maintained and updated over time. And, as new formats are introduced, new tasks require new filters, etc. In daily work, such complexity is a recipe for errors and inefficiency.

In this instance, holistic templates can help. All project-relevant settings must be stored in a single template. Manufacturers can help transport best practices to users with pre-configured templates. Or, even better, manufacturers can develop individual and customised templates in dialogue with users. This was the case with SWR, who can then refine and expand these templates as they need to. When they create a new project or timeline, they can select the template from which all settings are to be inherited.

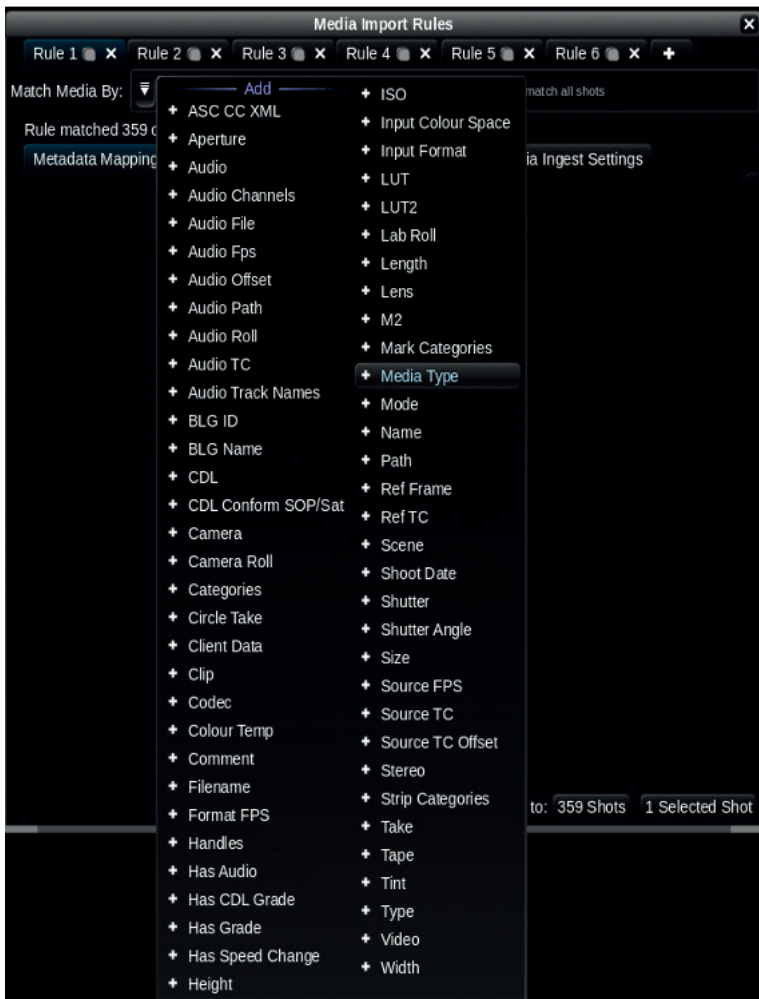
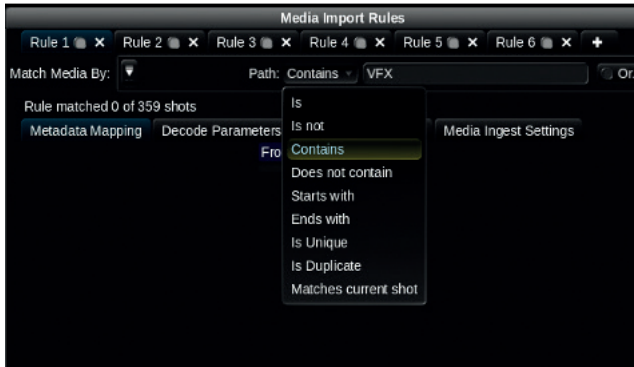
This way, the user does not always start from scratch, but instead builds on existing configurations. It is important to reiterate the completeness of these templates. A template must contain all production-relevant configurations. Conform settings, import and render parameters, file paths, media import rules, folder structures, and much more must be pre-set and, through the use of metadata, should also be designed to be valid generally.

In practice, this reduces the administrative overhead around colour correction and avoids sloppy mistakes. For example, at the end of the day, users do not have to worry about the location and naming of their renders, they can simply select from a list of deliverables. Everything else is done automatically. Users can concentrate on the creative aspects, and the 'pipeline' takes care of the rest.

Media import rules

Metadata-driven workflows can reduce or even eliminate manual work for various tasks. Importing media into a timeline is a good opportunity to initiate such processes. First, the user must define a rule (condition), e.g.:

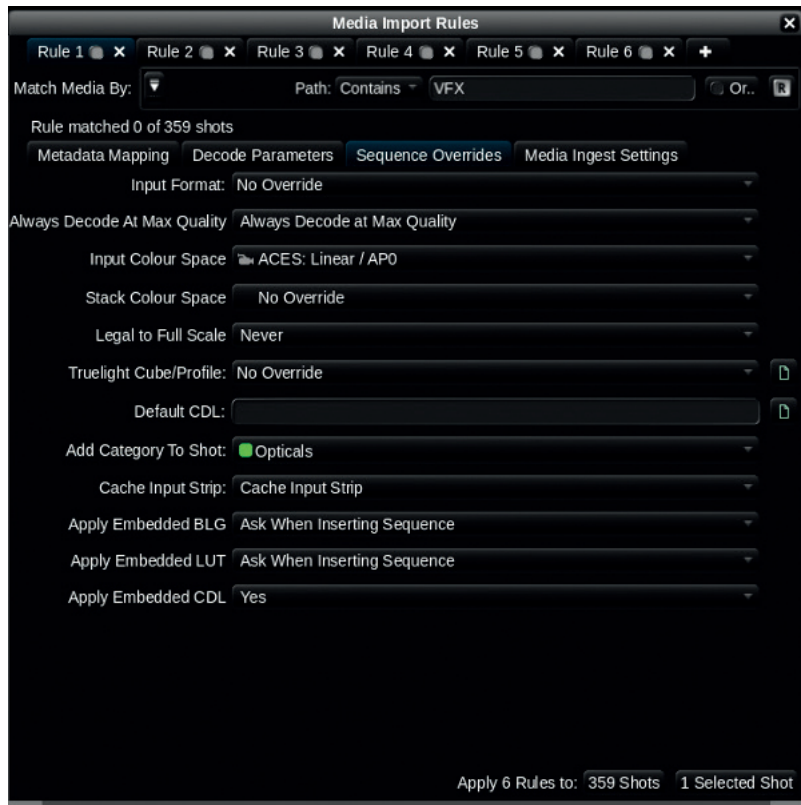
- » if the string 'VFX' occurs in the file path
- » if the ISO of the footage is > 800
- » if the LDS lens metadata serial number contains the letters 'xyz'
- » if the encoding is AVC-Intra
- » if the horizontal resolution is smaller than 1280



Media Import Rules - rules applied to footage coming into the timeline are a simple way of automation. The screenshot shows the criteria by which media can be selected. A combination of criteria is also possible.

If one or more conditions are met, many actions can be triggered, e.g.:

- » override the input colour space with S-Log3/S-Gamut3
- » apply a 'Legal to Full' scaling
- » tag the clip with the category 'VFX'
- » set the ARRIRAW sharpness parameter to -10%
- » kick off an external script



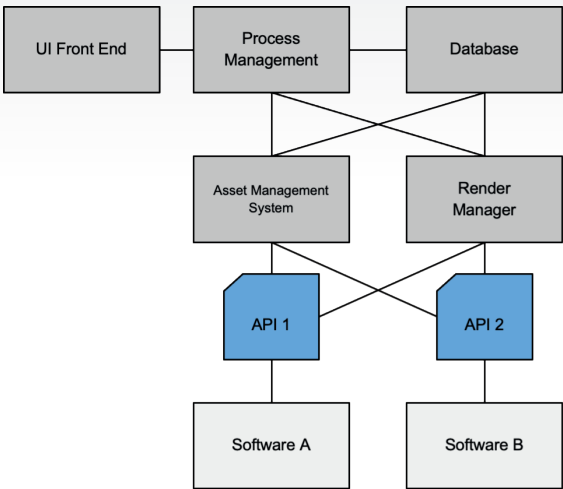
Media Import Rules. After selecting the relevant material based on criteria, the desired operation can be applied to these clips. Marking shots with categories has proven to be particularly useful.

These simple rule-based processes help to implement metadata-driven smart workflows in the field. The Media Import Rules are part of the scene template system which guarantee that the rules are applied on a project-specific basis. Again, the operational user does not have to configure anything, which makes it easier for the technical manager to roll out and maintain powerful, easy-to-use workflows.

Customised programming

Native digital workflows have great potential for automation, and the associated error prevention and increases in efficiency, if the software products in use provide Application Programming Interfaces (APIs). With APIs, software features can be used without manually controlling the software via the user interface. For this reason, FilmLight offers programming interfaces in the most common languages and protocols, such as Java, JavaScript, NodeJS and Python, via the FilmLight API.

While APIs provide the ability to execute software functions, you need a lot more to achieve automation success in a production environment. Asset management, process monitoring, persistent data storage and failure recovery are just some of the issues that need to be implemented when considering automation in the post-production process. This usually leads to the use of meta software that takes over these control tasks. Quickly, most of the development time is spent on asset management systems, render managers and suitable web front ends, etc.



Minimal software architecture for conventional post-production automation (greatly simplified). If you want two applications to work together via their APIs, you quickly have a pile of development tasks in front of you, such as front-ends and asset and process management.

In practice, however, this is often a significant challenge. Unifying the different levels of abstraction in one system is very complex. This is the right approach if a post-production department has sufficient internal software development resources. However, if only few development resources are available, this hurdle is often hard to overcome.

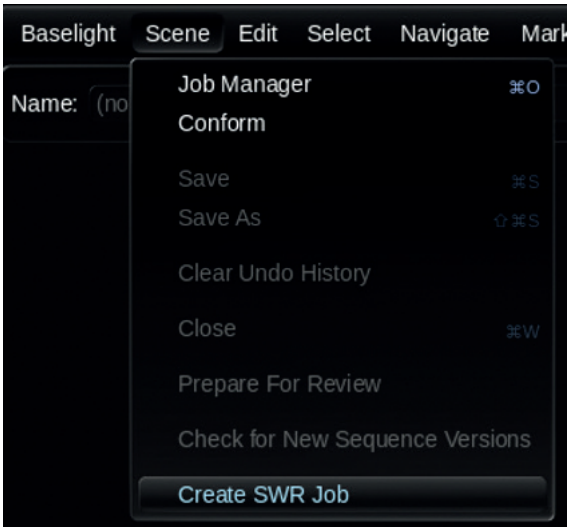
For this reason, the FilmLight API has onboard tools that programmers can employ to perform automation on their own. For example, simple graphical user interfaces can be created with little effort and clip-related information can be stored directly in the Baselight database. The FilmLight queue provides a simple but effective framework for process control to monitor long-running operations. As a result, application-based automation can be rolled out into the production environment with considerably less effort.

Below are two examples from SWR:

1) Project creation

Creating a new project (a movie, a season of a show or a commercial) is perhaps the best start to custom coding. It is simple and offers a lot of potential but also sets the initial course for the project.

Creating a new project usually comes with a series of tasks that must be completed stringently. Errors at this stage are often carried through the project's entire lifespan.



Project creation: At SWR, a function was added to a menu with the help of the FilmLight API to create new jobs, including all relevant administrative data. The script also creates the appropriate folder structure on the central storage.

At SWR, a new project is not created using the standard software features. There is an SWR-specific menu option in the user interface called 'Create SWR Job'.

If this option is selected, a window opens in which production-relevant information is collected. This window was tailored to the internal processes at SWR. With the data entered by the user, a script is started that executes various actions. Among other things, it includes:

- » Creating a job in the FilmLight job database
- » Creating a blank timeline in the proper format with the right template
- » Creating a predefined folder structure on the central storage

This concept can be expanded as desired over time. For example, the project could also be created in third-party databases. Project-specific watch folders, cronjobs, services or daemons can be registered.

Since all FilmLight-specific tasks are also carried out via the FilmLight API, creating the project does not necessarily have to be completed in the FilmLight system.

A higher-level PAM⁶ could handle the initialisation and remotely control the FilmLight-specific tasks.

2) Baselight and Diamant - integration of two applications

Audiences today have even higher expectations of the quality of archive material. Spoiled by the pristine images of modern UHD cameras, analogue image artefacts such as dirt, scratches and tape dropouts are disturbing.


Users can directly minimise defects in the colour correction system, such as interlaced moire patterns or cross-colour effects. However, as archive material from film or video tapes are still being used regularly, a high-end restoration solution should be part of the infrastructure. For broadcast environments, it is crucial that the workflow is not unnecessarily slowed down and tedious manual exporting and re-importing of clips is avoided.

For this reason, in the SWR installation, Diamant by HS-Art was integrated into the infrastructure via an API. The task is to send clips from the colour correction system directly to the restoration tool. The user can select from a list of predefined templates of the restoration software in the colour correction system and monitor the render progress.


After the successful completion of the restoration, the clips in the timeline are exchanged with the new version - completely automatically.

To carry out this automation, a few core components are required:

- » Asset management/UI (selection of clips)
- » Render management (transcode of the clips into the correct format)
- » UI (user selection of HS-ART Diamant pre-set)
- » Render management (Diamant rendering)
- » Asset management (detection and swapping of new versions)



User interface for Baselight-Diamant integration. Users can select shots directly in the Baselight timeline and have them treated in Diamant using one of the predefined templates. After the rendering is complete, the treated shots are automatically sent back to the Baselight timeline.



Software architecture of the Baselight-Diamant automation

This integration was achieved thanks to FilmLight’s Python API and HS-ART’s Rest API. In addition, the FilmLight API provides important core functions such as user input, process control and media asset management.

This solution scales very well with the different tasks, from high-speed processing that delivers improvements easily, to detailed processing with high demands on the result.

Conclusion

The digitisation of post-production has matured over the last decade. The goal is no longer to replicate non-digital processes digitally. Instead, new ways are being explored to optimise the post-production environment operationally. Through emergent and adaptable systems, it is possible to find new ways of working that gradually erode the linearity of the ‘workflow’.

Thanks to APIs and using the right triggers in the right places, workflows can be simplified or even fully automated with little effort. This is made easier by the simple and robust integration of metadata. In the next few years, we will see a similar development in post-production as in VFX production over the last decade. A pipeline will largely take over all technical processes, and the user will concentrate on creative tasks.

About the authors

Andy Minuth started a career as a colourist at CinePostproduction in Munich after graduating with a Bachelor of Engineering degree in Audiovisual Media in 2008. After also working as head of the colour department at 1000 Volt in Istanbul, he joined FilmLight in 2017. As a colour workflow specialist, he is responsible for training, consulting and supporting customers and is also involved in the development of new features in Baselight.

After completing his Master of Arts in Electronic Media, *Daniele Siragusano* worked for almost five years at CinePostproduction in Munich, where he advanced to become the technical director of post-production. In 2014, Daniele joined FilmLight as a workflow specialist and image engineer. Since then, he has been deeply involved in developing HDR grading and colour management tools within Baselight.

Head Office & EMEA

London, UK
t: +44.20.7292.0400
info@filmlight.ltd.uk

China

Beijing
t: +86.139.1073.7940

Germany

Munich
t: 49.89.323.09485

India

Mumbai
t: +91.9819.426.677

Japan

Tokyo
t: +81.3.6801.6280

Korea

Seoul
t: +82.10.7244.6122

Mexico

Mexico City
t: +52(1)55.5165.2132

Singapore

Singapore
t: +65.9670.3283

Thailand

Bangkok
t: +66.891.259.009

USA

Los Angeles
t: +1.323.785.1630

www.filmlight.ltd.uk



Northlight, Baselight, Truelight, Daylight, Prelight, FLIP, FLUX, Blackboard and Slate are trademarks of FilmLight Ltd.

Other products that are referred to in this document may be either trademarks and/or registered trademarks of the respective owners. The publisher and the author make no claim to these trademarks.

© FilmLight 2023

