Digital Intermediates and VFX for Modern Independent Film Production

A Production Handbook

Tashi Trieu

Version 2.0 · 2026

Getting Started

Where to Start

Use the path that matches your responsibility. Read the Production Workflow at a Glance first if the project has not established a post plan.

Producer or Post Supervisor

1. Review the production workflow.
2. Choose a color-management approach.
3. Complete the pre-production checklist.
4. Confirm distribution deliverables before budgeting.
5. Assign ownership for VFX tracking, DI turnover, delivery, and archive.

Editor or Assistant Editor

1. Confirm frame rate, timecode, reel structure, and source-name metadata.
2. Follow the VFX plate-pull turnover for effects work.
3. Follow the DI editorial turnover for conform and grading.
4. Use the DI turnover specification.

VFX Supervisor or Producer

1. Define the image pipeline.
2. Issue the vendor-onboarding template.
3. Set naming, handles, plate format, review format, and final-delivery rules.
4. Track every version and approval in one production database.

VFX Artist or Vendor

1. Read the supplied show and shot specifications.
2. Confirm working format, resolution, frame range, and color encoding.
3. Work through the VFX final-delivery checklist.
4. Do not infer missing color, naming, or version information. Ask production.

Colorist or Finishing Artist

1. Review the DI turnover and DI overview.
2. Confirm the color-management approach.
3. Test camera, VFX, title, HDR, and delivery round trips.
4. Define archival masters before the final render.

Delivering to Theaters or Platforms

Start with Distribution Deliverables, then use the relevant package:

- DCP for theatrical and festival exhibition
- IMF for component-based platform and international delivery
- HDR Mastering for HDR10, Dolby Vision, or HLG
- Archival for preservation masters and project reconstruction

Quick Lookup

- Production Checklists
- Technical Specifications
- Reusable Production Templates
- Glossary
- Standards, Tools, and Further Reading

Introduction

Digital intermediate finishing and visual effects now appear at every production scale. Independent filmmakers can execute work on laptops that once required a dedicated facility.

Why this handbook exists

Independent productions rarely have dedicated VFX producers, coordinators, production managers, or post supervisors. Freelance artists and picture editorial often absorb those responsibilities.

Poorly managed post-production spends limited resources correcting preventable errors. A documented workflow and clear departmental handoffs protect the budget.

Bring the colorist into planning before VFX begins. Lock the color space, resolution, format, and delivery requirements while changes remain inexpensive. Spend the budget improving images rather than rebuilding them.

Intended audience

The handbook is for independent producers, editors, colorists, and freelance VFX artists adapting studio practices to smaller productions. Basic familiarity with post-production and computer graphics is helpful.

About the Author

Tashi Trieu is a digital intermediate colorist and finishing artist with over fifteen years of post-production experience, and a co-founder of **MOSAIC**, an artist-driven picture-finishing boutique in Burbank, California. His work spans color grading, digital intermediate finishing, and end-to-end digital workflows for studio and independent features and premium OTT content: including multi-format HDR, Dolby Cinema, and IMAX delivery.

His color and finishing credits include *Avatar: The Way of Water* and *Avatar: Fire and Ash*, *Star Wars: The Force Awakens*, *Thor: Ragnarok*, *Bombshell*, *Titans*, and Netflix's *Stranger Things* and *Sense8*.

He co-authored *Modern Post: Workflows and Techniques for Digital Filmmakers*, contributed the Digital Intermediates chapter to the *VES Handbook*, has contributed to the Academy's ACES color standards, and has taught at Chapman University and Taipei National University of the Arts. He also builds web-based tools for color and post-production professionals.

Additional Resources

- **Supplemental Reading & Resources**: cited standards, white papers, and books.
- *Cinematic Color: From Your Monitor to the Big Screen* by Jeremy Selan, cinematiccolor.org

Special Thanks

- Billy Brooks
- Brian Gaffney
- Joshua Pines
- Christian Zak

Website

Updates: divfx.tashitrieu.com Questions and errata: divfx@tashitrieu.com

Generalized Production Workflow

The image path from pre-production through acquisition, post, delivery, and archive.

Pre-Production

Pre-production is not previs. It turns creative intent into a tested production and post-production plan. Production, editorial, VFX, dailies, color, sound, and the DI must agree on how images, sound, metadata, and approvals will move between them before photography begins.

Define the Workflow

Document the decisions that every department and vendor must follow:

- **Ownership:** pipeline owner, handoff owners, contacts, approval authority, and escalation path.
- **Acquisition:** cameras, recording formats, frame rate, timecode, raster, framing, audio, and required camera and lens metadata.
- **Data management:** folder structure, file and shot naming, versioning, checksums, backups, transfer methods, security, and retention.
- **Dailies and editorial:** proxy format, sync method, burn-ins, color metadata, review system, and the metadata required to relink camera originals.
- **VFX:** bid and turnover process, plate and final-render specifications, handles, mattes, references, review media, version tracking, and delivery schedule.
- **Finishing:** editorial turnover, conform method, titles, audio references, color pipeline, masters, distribution deliverables, and archive.

Publish these decisions as one show specification and one vendor-onboarding package. Use a shared tracking system so production can locate the current asset, version, owner, and approval state without reconstructing the history from email.

For every handoff, define who publishes the material, where it goes, how the recipient confirms receipt and QC, when it is due, and how missing or rejected material is corrected. The tool can vary. The method cannot.

Start with the Production Checklists, Technical Specifications, and Reusable Production Templates.

Develop the Look and Vet the Color Pipeline

Include the colorist in camera, lens, lighting, wardrobe, makeup, and production-design tests. Use those tests to develop the show look with the director and cinematographer, establish exposure and monitoring practices, and identify where the intended look creates limits for VFX or delivery.

Document each camera input transform, working space, viewing transform, Show LUT or creative transform, and target display. Distribute the same versioned package to dailies, editorial, VFX, and finishing.

Prove the pipeline with a round trip:

1. Shoot representative material with every camera and recording mode.
2. Process dailies and editorial proxies with the approved look and metadata.

3. Pull a VFX plate using the production specification.
4. Return a test composite through the approved interchange encoding.
5. Conform and grade the returned shot in the DI.
6. Compare it with the source outside the composited region and review it on the target display.

Resolve mismatches before production. A visually plausible test is not enough. The round trip must preserve framing, frame range, metadata, dynamic range, gamut, and unaffected pixels.

Plan Lenses and VFX Data Capture

During camera prep, the DP and VFX supervisor should review how the proposed lenses affect VFX. Photograph distortion grids and lens charts for every production lens. Repeat the tests at the focus and zoom settings that materially change the lens.

Decide whether VFX plates and elements should use the principal-photography lenses. Green-screen elements may be easier to extract and track when photographed with cleaner, less distorted spherical lenses instead of vintage or anamorphic glass. The DP and VFX supervisor should make that choice together. The goal is to simplify VFX without losing the film's creative intent or look.

Test the lens-metadata path during prep. Systems such as Cooke /i Technology, ARRI Lens Data System (LDS), and ZEISS eXtended Data can record lens identification, focal length, focus, iris, distortion, shading, and other measurements. Some cameras and lenses also record inertial data. Confirm that the camera records the required fields, that dailies and VFX plate pulls preserve or export them, and that the receiving vendor can read them. Automated metadata supports the VFX camera report. It does not replace it.

Previsualization

Previsualization (**previs**) is a separate creative-planning tool. Use it to test complex blocking, staging, camera movement, lenses, equipment, coverage, and pacing. It complements storyboards. It does not replace workflow planning.

If a full previs team is outside the budget, assign a freelance artist to the shots with the greatest technical or scheduling risk.

Production / Acquisition

Camera and Workflow

Choose a camera that records high-dynamic-range, scene-referred images in a documented gamut and transfer function. The label (cinema, mirrorless, drone, or action camera) matters less than what the file preserves and whether post can build a correct input transform.

Scene-referred recording either manifests itself as a recorded logarithmic (log) raster image, a scene-linear raster image, or a Camera RAW file. These topics are explored further in An Overview of Digital Intermediates.

On-Set Visual Effects

Poor plates increase cost, schedule, and risk. Put the VFX supervisor on set for VFX photography, ideally the same person who will supervise post. If unavailable, assign a trusted senior artist or another supervisor.

The on-set supervisor coordinates with the director, DP and camera department, production design, and practical effects. They decide what to capture practically, flag conditions that complicate post, and collect the measurements and references artists will need.

Detailed camera notes aid in reproducing accurate scenes in CGI. Note that these are not the camera assistants' notes. The AC's report is a camera-department document covering rolls, takes, lenses, and stops. It exists for editorial and the lab. **The VFX team owns the VFX camera report.** It records the information a matchmove artist needs, including camera height and tilt, lens and filtration, focus distance, sensor mode, and survey data. Ask for it explicitly. The AC's paperwork is not a substitute.

See Lens distortion grids and camera metadata.

Common Visual Effects Production Mistakes

Do not assume a difficult setup will be easy to fix later. Have the supervisor document each risk and solve it in camera where practical. For example, production design can incorporate trackable patterns into a set instead of adding markers that must later be removed.

Frame Rate: 23.976 or 24.000

Choose the rate before photography and use it across camera, sound, editorial, VFX, and finishing. 23.976 and 24.000 are both normal. Mixing them creates progressive picture and sound drift.

Why both exist. 24.000 is true 24 frames per second. **23.976** is `24000 ÷ 1001` and exists due to legacy broadcast compatibility with interlaced 29.97.[1]

Modern displays are not region-locked by mains frequency. Current TVs commonly play 23.976, 24, 25, 29.97, 30, 50, and 60 fps regardless of whether the local power system is 50 or 60 Hz. Display capability is not a good reason to choose between 23.976 and 24.000.

Broadcast is the important exception. Existing broadcast plants, schedules, and delivery systems often require a specific rate. Streaming apps do not have the same infrastructure constraint, and there is little legitimate technical reason they cannot support both 23.976 and 24.000. Delivery specifications can still impose business or workflow requirements, so confirm them before production.

Scanning or telecining film at the agreed frame rate is not a problem. Re-clocking a progressive master between 23.976 and 24.000 also preserves every frame. Avoid conversions that blend, interpolate, drop, or repeat frames. Avoid delivering 2:3 pulldown as if it were a native progressive master. Those are the damaging conversions that specifications reject.

What a mismatch costs. Run material at the wrong one of these two rates and nothing drops or repeats. Each frame stays on screen for slightly the wrong duration, so the error accumulates:

Elapsed	Drift
1 minute	~1.44 frames
42 seconds	~1 frame
1 hour	3.6 seconds (~86 frames)
90 minutes	~5.4 seconds
120 minutes	~7.2 seconds

The signature is **progressive** drift, not a fixed offset. Everything lines up at the head and degrades from there. If each reel is head-synced separately, the drift resets at every reel head and then climbs again.

Sound is where a mismatch becomes obvious. At 48 kHz, one 23.976 frame is exactly 2002 samples. One 24.000 frame is exactly 2000 samples. Production recorders can record at **48048 Hz** so the file plays at the intended speed when a 48000 Hz clock reads it. Recorder labels and NLE behavior vary.

!!! warning “Specify the operation” **Pull-up** and **pull-down** can describe the sample rate, playback speed, or other unrelated processes. Write the actual instruction: “Record at 48048 and stamp it 48000,” or “Slow the audio by 1000/1001.”

A DCP cannot be 23.976. Digital cinema uses integer frame rates. A 23.976 finish destined for theatrical must be conformed to 24.000.[2] The picture keeps the same frames and plays them 0.1% faster. Every audio element must change speed to match.

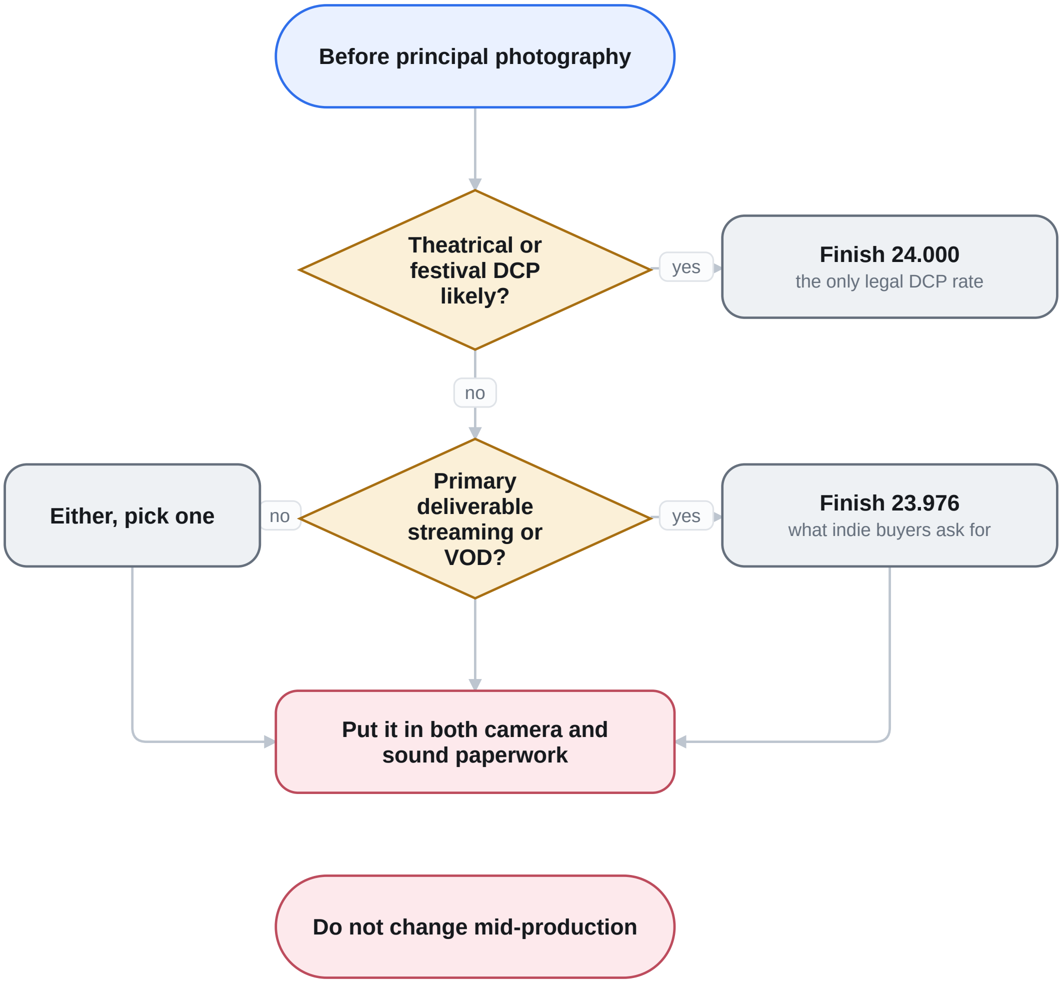

The frame-rate decision. Converting a finished master between the two rates is routine. Fixing mismatched camera and sound rates is not.

Choosing the Rate

Converting between 23.976 and 24.000 is a small speed change. The frame count stays the same. Choose the rate that avoids conversion on the primary delivery:

Leaning	Finish at	Why
Independent VOD/TVOD is the realistic destination	**23.976**	Common in independent distribution schedules. Confirm the actual platform specification.
Theatrical or festival exhibition is the priority	**24.000**	23.976 is not a legal DCP rate, so a DCP always forces the conform. Festival acceptances arrive late and their deadlines are short.
Broadcast is primary	The broadcaster's required rate	Broadcast infrastructure creates real regional and technical limits. Get the specification before photography.

Consistency matters more than the choice. A show shot consistently at either rate is fine. Re-clocking a finished master is routine. Mismatched camera and sound rates create a sync problem that compounds through every reel.

Pick one number, put it in both departments' paperwork before you shoot, and do not change it mid-production.

Digital Dailies

In a typical studio production, camera original media is ingested on set by a Digital Imaging Technician (DIT) or a Digital Loader (also called a data wrangler or data manager) and archived to hard drives, LTO tape, or both. A copy of those assets is delivered to the dailies team, either on-set, near-set, or at the lab or post-production facility. There, a dailies colorist or dailies technician will apply looks created on-set by the DIT or the cinematographer to relevant shots in scene context and make any necessary balancing or aesthetic color corrections. They will sync the production audio tracks to picture and then process the camera original media into an offline proxy format suitable for editorial, including relevant metadata necessary to reference the original media. These offline proxies are rendered with the dailies color grade and the Show LUT baked in, and are sized and formatted in accordance with the framing chart. Dailies are often processed unmatted in order to provide editorial with creative reframing flexibility.

In the independent production world, this level of complexity is rarely observed. Typically, there is a single DIT or a Digital Loader, but rarely both. They may also be responsible for syncing and processing dailies, but that task is often delegated to the production editorial team.

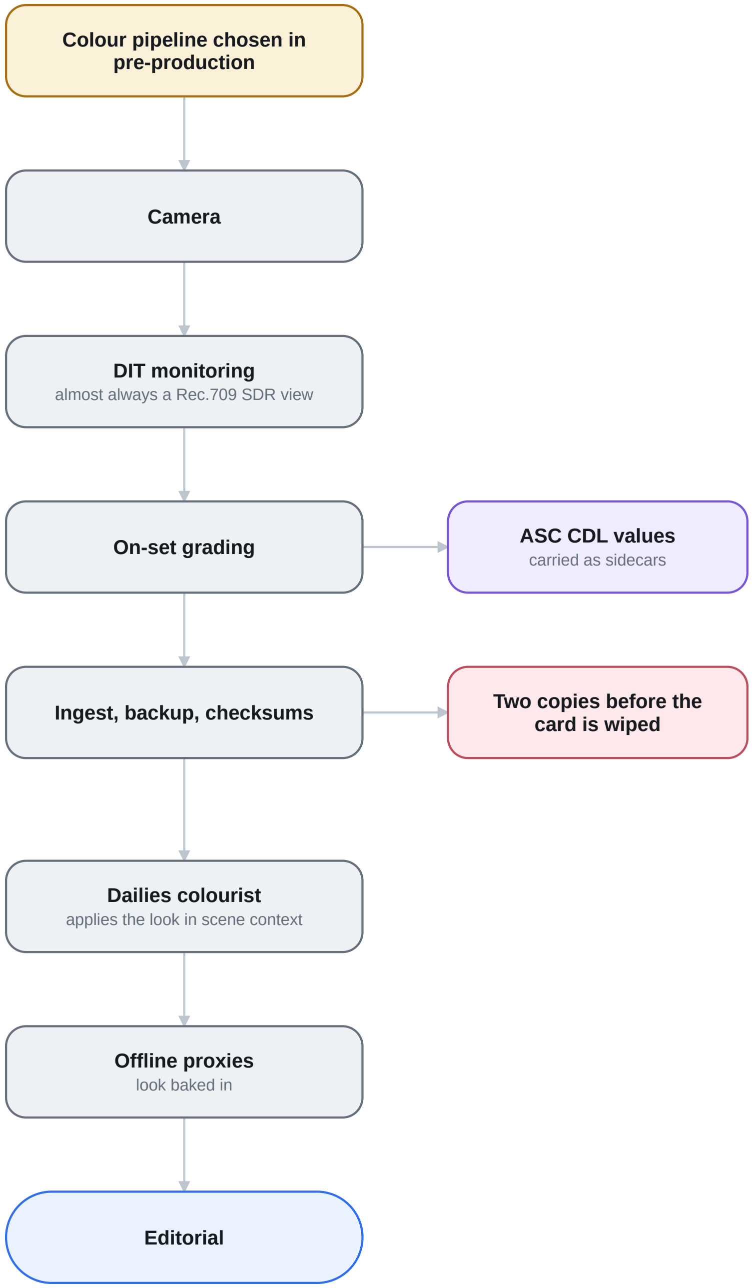

The colour pipeline from set to editorial. On-set grading reaches dailies as CDL values, so the decision made in pre-production has to hold all the way along.

Dailies Must Execute the Pre-Production Color Decision

Dailies is the first production use of the approved color pipeline, not the place to invent it. The dailies team applies the documented camera input, color metadata, Show LUT or creative transform, and viewing transform when producing editorial and review media.

If the show is going to be **ACES**, dailies is where that starts: an Input Transform per camera rather than a per-camera LUT of unclear provenance, and an Output Transform generating the viewing render. The practical benefit is that multiple vendors can implement the same pipeline without negotiating a bespoke workflow per show.

Choosing ACES, RCM, TCS, or a manual pipeline affects the entire post chain. Changing course after dailies begins means re-deriving looks, re-issuing plates, and re-testing vendor round trips. The colorist, cinematographer, VFX supervisor, dailies team, and post supervisor must approve the pipeline before the first day's material is processed.

The alternatives are equally legitimate. An undocumented one is not. See Choosing a Color Management Approach and confirm the production settings against the approved pre-production test.

Processed dailies are delivered to production editorial as Avid DNx LB (formerly DNxHD 36 / DNxHR LB) QuickTimes or MXFs, or Apple ProRes QuickTimes. Circled takes, and occasionally all takes, are transcoded to H.264s with audio and uploaded to a secure server and provided to the filmmakers for creative review. PIX (now Autodesk Flow Capture) and DAX (now part of CLEAR) are the predominant studio-approved review systems, but there are many more out there that can benefit productions of all scales. **Frame.io** (an Adobe product since 2021) belongs in the same tier rather than below it. It is TPN-assessed, carries SOC 2 and ISO 27001, and is one of the platforms that actually implements camera-to-cloud on studio productions. Its range from a free tier to studio deployments makes it practical for an independent production. Vimeo and Dropbox can serve as ad-hoc review in a pinch, but they offer no shot-level review, no burn-in control, and nothing comparable in security posture.

Efficient tools become even more important as the decrease in staffing brings the consolidation of responsibilities onto a select few individuals. Blackmagic Design's DaVinci Resolve has become a prominent resource for many DITs, loaders, and editors, as it supports all major Camera RAW formats, syncing audio, color grading, applying LUTs, rendering metadata burn-in, and transcoding to various offline and online codecs.

Creative Editorial

On studio films the picture editor and their assistants (who comprise the production editorial team) are on-set or near-set assembling cuts as the dailies arrive each day. Having nearly-immediate feedback is critical not only for the director, but also for the VFX supervisor. The VFX supervisor and VFX producer will be intimately involved with the timing and execution of complex sequences. As soon as the best takes are selected, they will have those plates processed and delivered to the visual effects vendor, whether that is a studio or a single artist. It is common for early visual effects work to begin while a film is still in production.

As the production progresses and visual effects shots are added spontaneously or reduced through alternative practices, the VFX supervisor and their team maintain an up-to-date assessment of the visual effects needs of the film through continual dialogue with the editorial team. On an independent feature with fewer complex sequences or shots, having early edits and reviews of dailies provides additional opportunities to identify any potential issues that could still be addressed during principal production.

During creative editorial it is common to use temporary graphics, rough composites (called slap comps), and previs animatics to convey story beats that will be produced and refined at a later date. This is important in estimating timing and determining whether key story elements are adequately conveyed. On a studio film, this is often the job of the visual effects editor, who works alongside the picture editor and the assistants. On an independent production without an expansive editorial team, those duties fall to the picture editor or their assistant(s).

The modern picture editor is a jack of all trades. Not only are they an expert in performance timing and an experienced storyteller, but they are also creative problem solvers looking for effective solutions to issues that arise during production or editorial post-production. Sometimes a line is cut, but there is insufficient coverage to cut away to, so they may creatively blend two halves of a shot together to remove a line of dialogue. This is accomplished with FluidMorph in Avid or similar optical flow morphing tools found in other software. Similarly, the desired dialogue of two actors in a two-shot may require the splitting and compositing of two separate takes over each other. Split-screen composites made with Avid's AniMatte or similar garbage mask tools are common and are a great example of seamless visual effects artistry.

The standard practice is to involve the VFX supervisor. On projects where the editor is the de facto VFX supervisor, VFX experience helps them judge what can be produced reasonably. They are not expected to produce the final result.

Visual Effects Plate Pulls

Once a sequence edit is locked, or very close to locked, the editorial team will produce a pull list of the visual effects shot elements so that they can be prepared and delivered to the vendors or artists who will be working on them. This list is often as simple as an EDL or XML of just the necessary visual effects shots, excluding any non-VFX footage. This list also correlates the original camera shot names with their intended VFX shot names.

It is important to pull these plates in a way that is efficient for the visual effects artists. It is not advisable to give the artists a hard drive containing the Camera RAW footage and a copy of the edit and expect them to accurately pull the shots they need. It is best to have pulls originate from a single third party or from production directly.

Pulls include **handles**, which are extra frames at the head and tail of the cut. This gives the visual effects artists the frames they need plus margin in case the edit is creatively rolled once the shots are complete.

Plate handles and comp handles are not the same number, and it is worth specifying them separately. **Pull plates with 16 to 32 frames at the head and tail.** Re-pulling a plate months later costs far more than pulling extra frames now. **Delivered composites usually carry 8 frames** head and tail. The VES Handbook's rule of thumb is a minimum of 8 at each end. Facility specs commonly ask for 32 on plates and 8 on returned comps.[3] The exact counts are a per-show agreement between production and the vendors, because every handle frame is also a frame the artists must finish and the production must pay for.

Visual effects plate pulls are commonly produced as image sequences (DPX or EXR), where each frame is an individual file rather than a single movie container. This is a common practice in visual effects and is highly recommended.

Visual effects plates are not provided to artists in the Camera RAW format, but instead debayered with show-specific settings for resolution, color space, and transfer function. These plate pulls are performed with a scene-referred transfer function to maintain all of the camera's native dynamic range and gamut. It is important that all plate pulls are produced by the same technician or company in order to maintain consistency. Otherwise individual visual effects artists or vendors are expected to coordinate clip pull methods together, and more often than not, unintentional deviations will occur.

Camera RAW is an inefficient working format for visual effects, as it must be debayered on the fly every time a frame is viewed or referenced. This is computationally expensive and redundant compared to pre-debayering to a rasterized working format. Moreover, standardizing on a single plate format makes working with footage from multiple camera types or sources much easier for the visual effects artists and maximizes software compatibility. While many compositing applications provide support for Camera RAW formats, debayering quality and algorithms can differ system to system, and many support applications used in asset management, 2D/3D tracking, and paint do not support them.

In studio productions, plate pulls are often handled by the digital intermediate vendor, as they have intimate knowledge and experience working with Camera RAW formats and color-managed workflows. At the independent production level, involving the colorist, or someone else with digital intermediate experience, is advisable. This can be someone like a DIT or a very technically savvy editor. When in doubt, consult with a post-production vendor or the final colorist to make sure you are following the correct procedures.

There are services and products designed for managing production assets on large projects and producing visual effects clip pulls. Codex Backbone and Technicolor PULSE were two such products used to decrease the time for visual effects plate pulls and turnovers to vendors, via hardware/software and cloud-based solutions that minimize the time and staffing required. Both have since left the market. Codex Backbone is no longer in Codex's product lineup, and Technicolor PULSE is defunct. Media-management and production-tracking platforms now serve the same need.

Visual Effects Plate Pre-Grading (Optional)

In studio productions it is common to perform a pre-grading pass on visual effects plates prior to their delivery to the visual effects vendors. This is usually performed by the final film colorist, or an associate of theirs, in collaboration with the director, cinematographer, VFX supervisor, and VFX producer. This process is done to give the plates a grade close to the final intended look of the film, as best as it can be approximated at this early stage in post-production. Naturally, this occurs long before the final color grading takes place, but is done to provide an approximation of how the shot will eventually be graded.

This is critical when the eventual grade will have a significant impact on CGI rendering and lighting. For example, a day-for-night look would likely influence the creative lighting decisions the CGI artists would make.

These grades usually use the ten ASC Color Decision List values: slope, offset, power, and saturation.[4] These controls are analogous to, but not identical with, lift, gamma, and gain. Apply them to the scene-referred plate under the display-referred Show LUT.

Deliver the unaltered plate and the CDL as metadata in an EDL, ALE, CC, or CCC file. This lets VFX preview the pre-grade without baking color changes into the image.

Alternatively, more complex grades can be baked in, but it is important to provide the ungraded plate as well, as grades can often reduce scene detail needed for tracking or degrade green/blue screen separations.

The Show LUT, shot-specific CDL metadata, and ungraded plates are delivered to the visual effects artists or vendors. If they apply the CDL values and the Show LUT, they will preview the exact same grade that you performed during the pre-grade. By accurately previewing the eventual grade of a shot, visual effects artists can composite their effects to look best and most accurate under that anticipated grade. This is particularly important when dealing with product shots or with characters that have a well-established color palette.

Figure 1: CDL grading operations are applied prior to the viewing transform. Ungraded log plate → ASC CDL → Show LUT (display reference transform) → Rec.709 preview.

The grade and Show LUT are baked in to any reference QuickTimes delivered to editorial in order to match the original dailies grade, but ultimately not included in the final VFX render, as the colorist will apply those looks in the DI session.

Visual Effects Production

Here is where all the magic happens. During this process the visual effects artists and vendors bring the shots to life. This is an iterative process that relies on continual support and collaborative feedback from the filmmakers. Depending on the complexity of a shot there may be more than one artist assigned to it, and in some cases multiple VFX companies. There are many different departments within visual effects that artists belong to. Many are compositors, specializing in integrating elements and solving problems using 2D or 2.5D tools. Others are CGI generalists or specialists, creating something from nothing in 3D and working with compositors and lighters to integrate it into a live-action scene in a photorealistic way.

The more shots you have and the more artists working on them, the greater the need for efficient shot management. This falls under the responsibility of the VFX producer and the VFX coordinators, often grouped under VFX production management. On a small independent production this may, once again, fall back on the editor. Managing visual effects work in a flexible, organized way is critical to your success.

Different VFX producers and coordinators have different management styles and tools that they use for tracking assets, artists, client feedback, and progress. Sometimes this is as simple as shared spreadsheets either from Excel or Google Sheets. Other times it requires more complicated software specific to the visual effects industry like **ftrack Studio** (Backlight) or **Autodesk Flow Production Tracking** (formerly Shotgun, then ShotGrid). Open-source options such as Kitsu and AYON have also become realistic choices for small studios and freelance teams. These production management systems are purposely designed to maintain the flow of visual effects production and help the producers stay on track. Nothing is more worrisome than not having a reliable estimate of when you can complete your work. VFX production management is explored in greater detail in Visual Effects Production Management.

One of the core tenets of visual effects production is that regions of an image without visual effects (i.e. elements of the camera original plate that are retained in the final comp) should undergo zero net change. Unless specifically mandated by the visual effects workflow, there should be zero net change in color space or color encoding. It is very common for material to be transformed into a more convenient working space for compositing (e.g. camera log to scene-linear). It must return to its original encoding before delivery and without clamping, quantization, or other artifacts.

A best practice to avoid issues is to double-check the final render against the original plate. If there is a measurable color shift or degradation that is not part of the intended effect, there is a pipeline issue that needs to be addressed prior to delivery to the client.

It is common for the VFX production team to produce a confidence package consisting of a few test images in the project's working format and provide it to the various vendors and artists they are hiring. The vendors and artists will pass the test images through their entire pipeline and return them back to the VFX production team for review. This allows production and the vendors to identify issues in the VFX pipeline before expensive work begins on actual shots.

Visual Effects Editorial

During the visual effects process the filmmakers will receive preview renders from artists and vendors. These renders are usually lower quality QuickTime proxies and have the temp color correction (from the pre-grade) and the Show LUT baked in so that the picture editors can cut it into their timelines and review in the context of the scene or sequence dailies. These are regularly delivered as Avid DNx LB or Apple ProRes 422 LT. Here the goal is to review for content and design, not technical quality control. These offline references are named and contain timecode to match against the full quality, ungraded DPX or EXR renders that will ultimately be delivered to the DI.

It is important that the imported QuickTimes retain their timecode metadata throughout offline editorial. Without timecode, the editorial conform turnover to the DI will not accurately reference the full quality renders that VFX has delivered to the DI. It is standard practice to conduct workflow meetings and tests early on between the VFX vendors, creative editorial, and the DI to ensure that conform and color management issues are identified and resolved early.

Visual Effects Reviews

Later in the visual effects process, as shots are nearing completion, the filmmakers (director, producers, VFX supervisor, VFX producer, and their respective teams) review full quality renders of the shots in a color accurate context, often in a digital cinema theater at a post-production facility. At this stage it is important to review in a calibrated environment accurate to the intended exhibition (theater for theatrical projects, broadcast monitor for broadcast or streaming projects). Here they can identify technical issues that are more apparent on the big screen and can experiment with color grading decisions with the colorist or an associate colorist. Here they will review in the context of the Show LUT and VFX shot CDL values.

Digital Intermediate

As visual effects near completion, they are delivered both to the filmmakers and to the digital intermediate vendor/facility either simultaneously or when requested by the production editorial team. Production editorial (often the VFX editor) will cut in the latest and greatest visual effects shots and the assistants will turn over edit lists (EDLs, XMLs, or AAFs) to the DI, where the conform editor(s) will conform the DPX or EXR renders and update the colorist's timeline. It is necessary for the file naming and frame numbering of visual effects shots to consistently match the offline reference QuickTimes in order to facilitate a fast and accurate conform.

In the grading theater, filmmakers will view the shots in the full, highest quality with the colorist and make color grading adjustments. This again is an iterative process. It is prudent to provide temp visual effects shots or works in progress (WIPs) to the colorist so that they can do some of the preliminary grading work, rather than waiting idly until a shot is completely final. As the deadline approaches, it becomes more and more common for the colorist to tackle issues that are too time consuming to send back to visual effects, such as lighting continuity and color matching.

Visual effects typically fall into three categories of completion: temp (WIP), final, or could-be-better (CBB). CBBs are any shots that are deemed passable as is, but have improvements to be made if time permits. Those shots are only a priority once all of the temp shots are finalized. It is not uncommon for CBBs to make it into early deliveries of a film, such as international theatrical releases which often ship earlier than the domestic versions. On a studio picture there is usually a home-video mastering pass where outstanding CBBs get improved, on the reasoning that home video is where the film will be scrutinised most closely over time. **Do not assume that pass exists on an independent film.** More often the theatrical version is the only version anyone pays to finish, and a CBB shipped is a CBB shipped for good. Treat the list as work to schedule rather than work that will find its own opportunity later.

Visual Effects Version Checking

Before finalizing the DI on a reel or episode, the editorial assistants and VFX production coordinators will work with the conform editor to verify that all the correct versions of visual effects are in the film or show. This usually involves checking version numbers, but can also require visual inspection of shots to verify that the intended revisions are present. It is not uncommon for a production to roll back to an earlier version of a visual effects shot if they like it more than subsequent revisions, so it is never safe to assume the latest available version is the correct one.

In addition to manually verifying the shots, an EDL of the VFX shots, exported from the DI to production editorial, can aid in performing this critical check.

Theatrical Distribution Masters

Feature films destined for digital cinema release are mastered and delivered to theaters in Digital Cinema Packages (DCPs). DCPs are mastered from a DCI-P3 referred graded master. They are often assembled and mastered on separate systems apart from the DI, and oftentimes by facilities separate from the colorist's. The various picture deliverables required to produce a DCP are:

DSM (Digital Source Master)

An uncompressed, unencrypted, graded render of the film. Typically DCI-P3 (RGB) color space. Sized per digital cinema specifications. Often 10-bit or 16-bit DPX.

!!! note This element is often skipped, as the colorist or DI facility will produce a DCDM directly from the color grading system, rather than from a DSM.

DCDM (Digital Cinema Distribution Master)

An uncompressed, unencrypted, graded DCI-P3 render of the film in DCI-X'Y'Z' color space. Sized the same as the DSM. 16-bit TIFF. The DCDM is only useful for producing a DCP and is typically not utilized in any other deliverables.

Because those 16-bit TIFF sequences can reach about 10 TB for a feature, exchanging a DCDM between facilities is slow and costly. **SMPTE ST 428-24 (2024)** addresses this with what it calls a **packed image**, known in practice as **pDCDM**. SMPTE's own drafting group and ISDCF use the shorthand, though it does not appear in the published standard. A pDCDM maps the DCDM into a **losslessly (reversibly) coded JPEG 2000** file sequence and packs the image's 12-bit code values directly instead of padding them into a 16-bit container.

The baseband image reconstructs exactly. This is *not* a lossy proxy. File size drops by roughly half from the reversible coding, plus a further ~25% from the 12-bit packing. It may optionally use the high-throughput JPEG 2000 (HTJ2K) block coder for much faster encode and decode. Frames carry the required suffix **`.pdc.j2c`** to distinguish these reversible frames from the lossy `.j2c` codestreams destined for a DCP. Think of it as a transport-efficient equivalent of the TIFF DCDM, not a substitute for the DCP's lossy essence.

DCP (Digital Cinema Package)

The DCP is a compressed, encrypted package muxed from the DCDM picture render, audio printmasters, subpictures, and timed text (subtitles). Encrypted DCPs can be played back on servers that have had valid KDMs (Key Delivery Messages) generated for them by the DCP creator.

Features being commercially exhibited in mainstream theaters are encrypted and the various theaters are given digital keys that allow them to play the film for a certain range of dates, often limiting the hours of the day that the film can be played. For festivals and private screenings of independent films, it is most common to produce unencrypted DCPs to avoid potential security issues in the event that a venue changes last minute or there is not enough time to test and validate the keys prior to the screening.

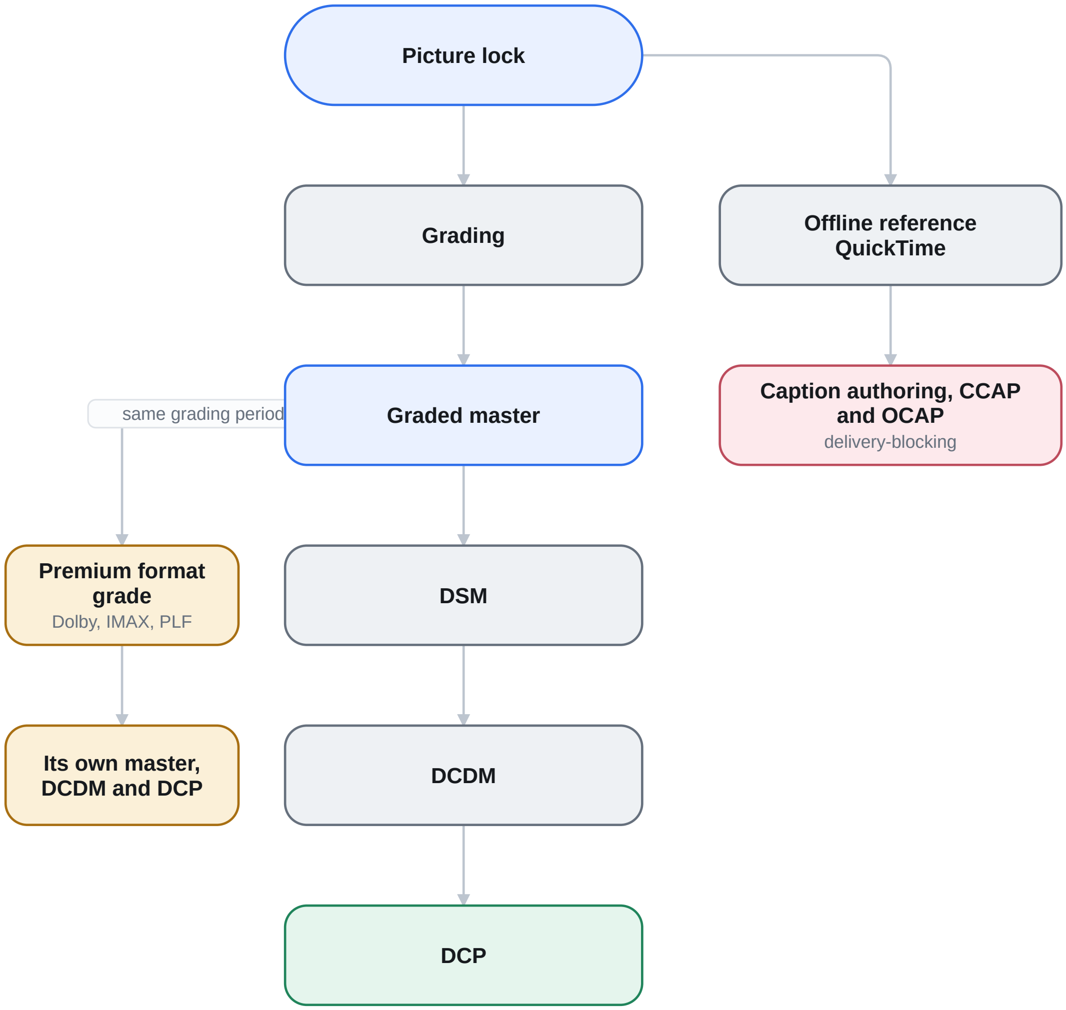

Theatrical masters. Captions are authored from the offline reference rather than the graded master, and a premium format is a parallel grade producing its own DCDM and DCP.

Accessibility Tracks Can Block Delivery

Captions are not an optional extra on a theatrical deliverable. Distributors and an increasing number of festivals require them, and a package that arrives without them can be rejected outright. The 2025 Utopia delivery schedule is one of several that treats accessibility as a delivery condition rather than a nice-to-have, and **SXSW requires closed captions on every film, TV project and music video it programmes.**

Two distinct things get confused:

- **CCAP (closed captions)** are delivered to the viewer's own seat-back display or handheld unit. They carry dialogue *plus* non-speech audio information: sound effects, speaker identification, musical scoring cues.
- **OCAP (open captions)** are burned into the picture and visible to the whole auditorium. These are subtitles by another name and are a picture decision, not a metadata one.

Practical constraints that catch people out: caption files have size limits (SXSW caps theirs at 256 kB, with reel-splitting as the prescribed remedy), and line-length conventions cluster around **30 characters per line, three lines**. This limit appears to come from seat-back display hardware rather than from any individual festival's preference. Non-English content generally needs English subtitles regardless of whether captions are also supplied.

Budget the authoring time. Captioning a feature properly is days of work, not an export.

Premium large format

If the show has a premium-format release, those are **separate deliverables**, not crops of the DCP. **Dolby Cinema** takes a Dolby Vision cinema grade. This is a PQ-graded DCDM at a far higher peak luminance than the SDR master, flagged in the CPL rather than carried as an enhancement layer. **IMAX** runs its own DMR process and asks for elements in its own specifications, including a final-cropped master, since every pixel delivered reaches the screen. Other **PLF** chains sit somewhere between.

Each means an additional grade, additional QC, and a delivery specification you obtain from the vendor rather than infer from the general-release rules. See DCP. Almost no independent film will need this, but if a premium release is on the table it belongs in the post budget from the start, because it is a second finishing pass rather than a re-export.

Video Distribution Masters

For projects destined for home video or over-the-top (Netflix, Amazon, Apple TV, etc.) release, there are some typical delivery formats:

Uncompressed Video Master

A complete, color graded master with titles, graphics, and visual effects, with accompanying textless elements. Textless elements are often delivered with handles (the entire shot before and after the textless element). Video masters are rendered in the mastering display's color space (i.e. Rec.709 or Rec.2020) and at agreed-upon delivery resolutions, typically HD (1920x1080) and UHD (3840x2160).

IMF (Interoperable Master Format)

The studio and streamer interchange master, mastered per the SMPTE ST 2067 (IMF) specifications. Unlike a flat video master, an IMF package is not one married file: picture (JPEG 2000), each audio configuration (linear PCM), and timed text are carried as **separate track files**, referenced by a Composition Playlist that describes how to assemble them. That structure is what makes territory and edit versions possible from one package via supplemental packages. It is why **textless elements are a core requirement rather than an omission**: the alternate versions are built from them. See IMF.

ProRes 422 HQ

For independent distribution this is the dominant deliverable. Not just the Apple storefront but the great majority of TVOD platforms and independent distributors want a compressed **ProRes 422 HQ** variant of the video master. ProRes 4444 or 4444 XQ is increasingly specified outright, particularly for HDR. Very few independent distributors ask for an IMF package at all. IMF is a studio-and-streamer interchange format. The label releasing an independent feature almost always wants ProRes and handles the platform-specific encoding downstream. See Distribution Deliverables for Independent Film for what real indie delivery schedules actually require.

Archival Elements

The workflow does not end at the distribution masters. The last stage is often deferred until the money is gone. Preserve the elements a future version would have to be rebuilt from: the graded digital source master, the un-graded conformed master, camera original, project files and the VFX elements. Distributors increasingly specify these explicitly, and some require a preservation master that is deliberately *un*-colour-corrected and uncropped so that a future remaster is not locked to today's grade or framing.

Archival is a chapter in its own right. See Archival and Preservation for storage media, migration intervals, and what actually belongs in a preservation package.

Choosing a Color Management Approach

Choose one of five color-management approaches:

1. **Display Referred, un-color-managed**: the look is baked into the grade, with no recoverable scene-referred master.
2. **Manually Color Managed**: a hand-made pipeline with a consistent scene-referred working space, color-space transforms in, node- or layer-based grading, and a standard or bespoke display transform out. This is typically what people mean when they talk about “node-based color management” in Resolve.
3. **ACES**: an open, standardized, scene-referred pipeline (SMPTE ST 2065).
4. **DaVinci Resolve Color Management (RCM)**: Blackmagic’s in-application managed workflow on DaVinci Wide Gamut / Intermediate.
5. **FilmLight Truelight Color Spaces (TCS) + T-CAM**: FilmLight’s managed workflow on E-Gamut / T-Log with a color-appearance display transform.

In every managed approach, track whether the image is scene-referred or display-referred and apply the display transform last. All but the first approach retain a scene-referred working space and recoverable master.

!!! note “This is a choice about interchange, not image quality” Managed systems enforce consistency across cameras, vendors, applications, and deliverables. They do not guarantee better images. Choose based on interchange, workflow, and archival needs. ### The five approaches

Display Referred, un-color-managed

You grade the footage more or less as the camera and monitor present it: a camera-native log or even a display-space signal, and bake the look directly into the grade. There is no defined working space, no separately applied display transform, and **no route back to a graded assembly master.** The master *is* the look. There is no scene-referred version to re-render from.

- **What it is not.** A simple job with a consistent working space and the display transform applied last is Manually Color Managed. An un-color-managed finish leaves only a baked image for future deliverables.
- **Why it fails.** Technical and creative decisions are entangled. It does not scale to multi-camera, multi-vendor, or multi-deliverable work, and it fails quietly rather than loudly: you discover the cost only when someone needs the scene-referred master that cannot be produced.

!!! warning “Avoid an un-color-managed finish” A second deliverable, re-grade, show LUT, VFX preview, or archival master requires a recoverable scene-referred source. Establish the working space and deliverables with the colorist before engagement. #### Manually Color Managed

The colorist maps each source into a common scene-referred working space, grades with node- or layer-based tools, and renders through a standard or custom display transform. CSTs, LUTs, and DCTLs replace an automatic vendor framework.

- **Benefits.** Full control over the working space, tone mapping, and look while retaining a recoverable graded master.
- **Drawbacks.** Requires an experienced colorist to design, document, test, and distribute every transform. An undocumented manual pipeline is as fragile as an unmanaged one.

ACES

A scene-referred interchange encoding (ACES2065-1, SMPTE ST 2065-1), working spaces derived from it (#encodings) for comp, ACEScct for grading), an Input Transform at the front and an Output Transform at the back, delivered in practice as an OCIO config.

- **Benefits.** Open, vendor-neutral interchange through OCIO and a documented scene-referred archival encoding in ACES2065-1.
- **Drawbacks.** Requires setup, testing, and version discipline. Its interchange benefits may go unused on a single-camera, single-deliverable project. Use another approach when the desired look conflicts with the standard Output Transform.

DaVinci Resolve Color Management

Blackmagic's DaVinci YRGB Color Managed mode: per-clip Input Color Space → a timeline working space → an Output Color Space per deliverable, with the display transform applied last.

- **Benefits.** Low-friction Resolve workflow. Documented DWG/Intermediate working space. Straightforward SDR, HDR, Rec.709, P3, and Rec.2020 outputs.
- **Drawbacks.** The display rendering is internal to Resolve and is not an archival interchange standard. External vendors need supplied transforms and guidance.

FilmLight TCS + T-CAM

FilmLight's Truelight Color Spaces: an **E-Gamut / T-Log** working space and the **T-CAM v2** color-appearance display transform, in Baselight and Daylight.

- **Benefits.** Scene-referred Baselight workflow with a documented CAM-based DRT and portable FilmLight transform files, OCIO config, and Flame policy.
- **Drawbacks.** FilmLight controls the implementation. Cross-application support is narrower than ACES. ### Side-by-side[5]

	Display Referred, un-color-managed	Manually Color Managed	ACES	Resolve RCM	FilmLight TCS
Working space	None (camera-native / display)	Consistent scene-referred (colorist's choice)	ACEScg / ACEScct	DaVinci Wide Gamut + Intermediate	E-Gamut + T-Log
Display transform	Baked into grade	Chosen standard or bespoke DRT	ACES Output Transform	Resolve DRT (internal)	T-CAM v2 (CAM-based)
Enforces the discipline	No	By the colorist	**Automatically**	**Automatically**	**Automatically**
Route back to a scene-referred master	**No**	Yes, if documented	**Yes** (ACES2065-1)	Yes (working-space master)	Yes (working-space master)
Open standard?	n/a	No, hand-built	**Yes**, SMPTE ST 2065	No, single vendor	No, single vendor
Published?	n/a	As the colorist documents it	Fully (free SMPTE library + open CTL)	Encoding yes (white paper). Rendering no	Color-space/DRT files + OCIO config. Not a spec
Cross-app native support	n/a	Only via the colorist's exported transforms	**Broad** via OCIO	In-application (Resolve)	Via published files / OCIO config / Flame policy
In-app setup (single facility)	None	Moderate–high (build and test)	A little config	Low (if in Resolve)	Low–moderate (if in Baselight)
Multi-camera normalization	Poor	Strong (manual CSTs)	Strong	Strong	Strong, *equal labor to ACES*
Multi-vendor VFX (total effort)	Fails	High, colorist supplies transforms + guidance	**Low, vendors already speak it**	High, colorist authors transforms + guidance per vendor	High, same despite portable files
Interoperates with ACES	n/a	Yes, if built on ACES	Is ACES	Yes (Resolve has ACES modes)	Yes (v7 ACES 2.0)

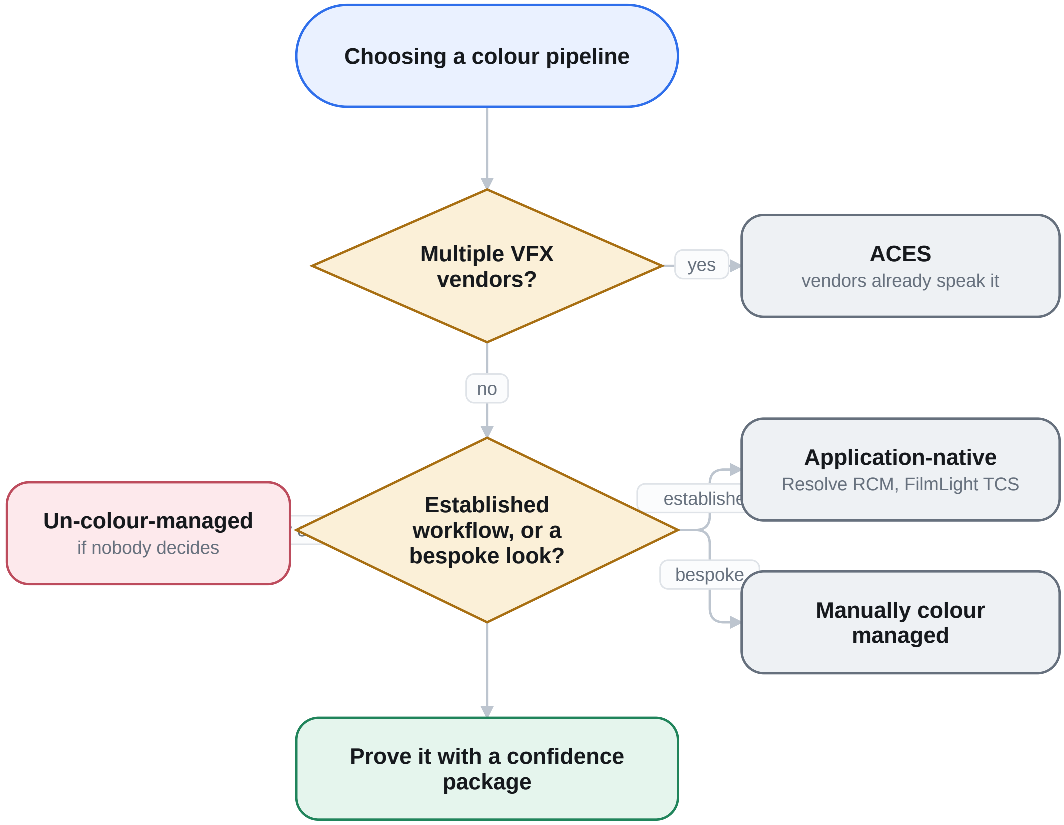

Choosing an approach. The expensive mistake is failing to choose.

Which should an independent production choose?

Rule out truly un-color-managed. For a production that will deliver professionally, a baked, unrecoverable master is not a reasonable option. This is particularly true for a VFX production.

The baseline is a **manually color-managed** pipeline: a consistent scene-referred working space, the display transform(s) applied last, and a documented master you can re-render from. That floor costs almost nothing over "just grading," and it is the difference between a show you can repurpose and one you cannot.

Reach for a framework (ACES / RCM / TCS) when:

- You have **multiple camera systems or VFX vendors**. Prefer **ACES** when work crosses applications and vendors.
- You **lack the color-science resources to build and vet a bespoke workflow.** A documented, standardized pipeline you can adopt is safer than a homemade one you cannot fully test: a framework gives you the discipline of a manually color-managed pipeline without having to author it.

Between the frameworks:

- Choose **ACES** for the widest, most vendor-neutral interchange and for a **scene-referred archival master**: the Graded Archival Master case.

- Choose **Resolve RCM** for a single-facility Resolve finish with scene-referred VFX deliveries in an agreed working space.
- Choose **FilmLight TCS / T-CAM** when you finish at a Baselight house or specifically want T-CAM's rendering. Its published files and OCIO config make it more portable than Resolve's DRT, and it interoperates with ACES.

A manually color-managed pipeline earns its place when:

- You have the **color-science resources** to design, document, and test it, plus a specific reason to control the DRT or look.
- The job is small enough that a full framework is overkill but you still want a consistent working space and a recoverable master: the responsible version of "keep it simple."
- Bear in mind it costs the same as a vendor framework to push across VFX vendors: the colorist has to supply the transforms and written guidance, just as with RCM or TCS.

Even a single-deliverable commercial benefits from a consistent working space: controls behave predictably and grades and looks remain reusable across camera sources.

!!! tip "Use interchange and archival as the tie-breaker" More sources, tools, vendors, deliverables, or uncertainty favor a framework and often ACES. A small, fixed pipeline can remain manual. In either case, retain a consistent working space and a recoverable master.

Choose the pipeline in **pre-production** with the colorist. Record it in the format specifications and prove it with a confidence package before shot work. Preserve a defined working space and re-renderable master.

Image Pipeline Foundations

Scene-Referred vs. Display-Referred (Output-Referred) Imagery

Most production images are either scene-referred or display-referred.

!!! quote "… categorize color spaces by the 'direction' of this relationship to real world quantities, which we refer to as image state. Color spaces which are defined in relation to display characteristic are called display-referred, while color spaces which are defined in relation to input devices (scenes) are scene-referred. While there are other flavors of images states (intermediate-referred, focal-plane referred) display-referred and scene-referred colorimetry are most commonly used in motion-picture color management …"

```
Jeremy Selan, Sony Pictures Imageworks, *Cinematic Color*, VES, 2012
```

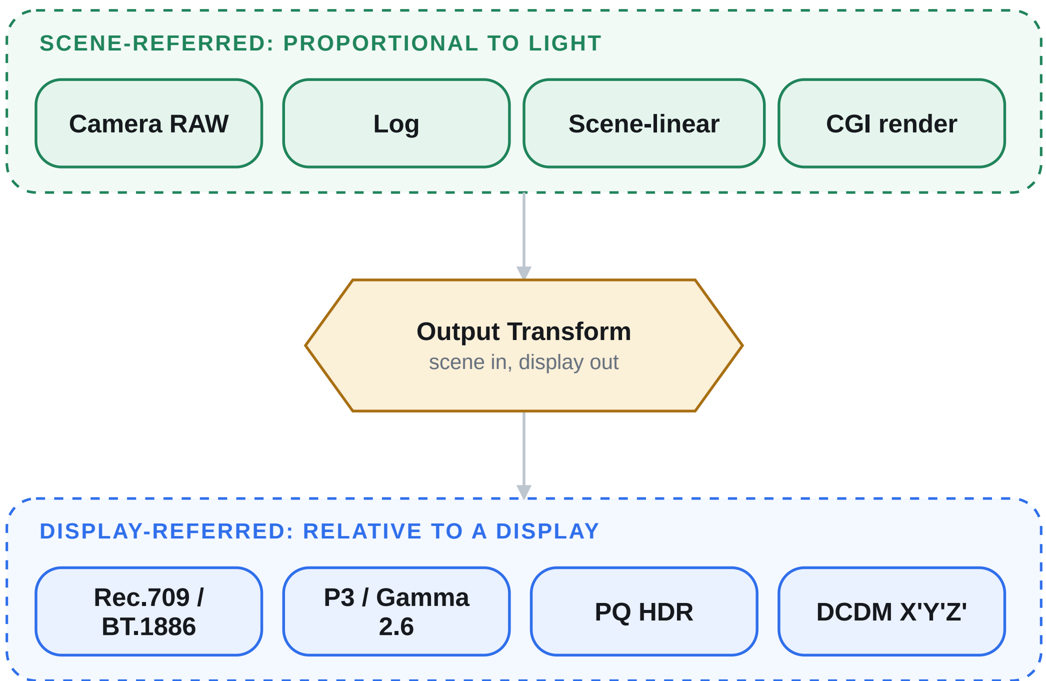

The image-state boundary. Everything on the left is proportional to scene light. Everything on the right is relative to a display.

Display-Referred Imagery

Display-referred values describe output on a target display, not real-world exposure. Web, phone, television, and cinema images are display-referred and require no additional viewing transform on their intended display.

Almost all reference displays have a built-in gamma encoding. These encodings are different depending on the display, its usage and intention, and the assumed viewing environment. For example, an sRGB image on a computer monitor or phone display is viewed at roughly gamma 2.2, home video Rec.709 material through a display EOTF of about gamma 2.4 (BT.1886), and standard dynamic range digital cinema through gamma 2.6.

In analog terms, a display-referred image is our film print.

An image intended for exhibition will ultimately be display-referred.

Scene-Referred Imagery

Scene-referred values preserve relative exposure relationships in a photographed or CG scene. They do not describe final display output or absolute light.

Scene-referred images in the context of digital cinema workflow are any camera original images (or visual effects shots encoded in camera original encoding) that retain the camera's native dynamic range and tonal response.

Camera RAW images are inherently scene-referred, while traditional video images are display-referred.

Scene-referred images can come in two major flavors: scene-linear (linear light space) or any number of camera log encoding schemes (such as ARRI LogC3 / LogC4, RED Log3G10, Sony S-Log3, Canon Log 2 / Log 3, Blackmagic Film Gen 5, or Apple Log). Always record the version number. “LogC” and “Canon Log” each name more than one curve, and they are not interchangeable.

Scene-Linear

Scene-linear encoding is a purely linear (no gamma) encoding of relative light values, expressed in floating point values. Middle gray is mapped to 0.18, each stop doubles the value, and highlights may exceed 1.0. Scene-linear images are commonly written to **half-float** OpenEXR files, 16-bit *floating point*, IEEE 754 binary16, rather than historically integer-based formats like DPX or QuickTime. The distinction matters: “16-bit” on its own is ambiguous in post, where it can mean either half-float or 16-bit integer, and only the float type carries values above 1.0.

Scene-linear is not to be confused with normalized linear, or camera linear, which are float or integer encoded values ranging from 0.0 to 1.0 and correspond directly to a camera's sensor raw analog-to-digital output values prior to encoding.

Camera Log

Log encoding expresses scene exposure in stops, preserving wide dynamic range in 10- or 12-bit integer files. Scene-linear work generally requires 16-bit floating point.

As the dynamic range of new cameras increases, the need for 16-bit image encoding increases. 10-bit log is insufficient for incredibly high dynamic range encodings, as it can produce quantization artifacts (or banding) if too much dynamic range is compressed into a small range of code values. This is particularly problematic when mastering for HDR, and for theatrical delivery, where the projected image carries more range than a Rec.709 video master. In those situations 12-bit log should really be the minimum bit depth of camera original images, with 16-bit log or 16-bit scene-linear being ideal.

The major reason we still use logarithmically encoded images is that it allows us to capture scene-referred imagery to a variety of convenient formats without the overhead of Camera RAW. The ARRI ALEXA 35 and RED V-RAPTOR allow for native ProRes capture, while Sony's VENICE 2 and BURANO allow native capture to other raster image formats like XAVC. Not every camera offers this. Blackmagic's URSA Cine line, for instance, records only Blackmagic RAW internally (alongside H.264 proxies), so it remains a RAW workflow whether you want one or not. External recorders provide further options for capturing scene-referred imagery without Camera RAW.

As Camera RAW typically exhibits higher data rates than ProRes and necessitates more intensive dailies processing, Camera RAW may be an expensive luxury for some productions. In particular, broadcast television productions rarely utilize the full benefits of RAW and instead record in a camera log based format like ProRes at a convenient resolution and format.

Camera log images are analogous to log film scans and fit in with established workflows very easily.

Video ≠ Linear

Do not use **linear** as shorthand for **video**. Rec.709 video is display-referred and gamma-encoded. It is not linear.

The distinction matters in VFX specifications. “Linear plates” normally means scene-linear OpenEXR, not display-referred video.

!!! example When a client tells you they are providing “linear plates” you would naturally expect a scene-linear, scene-referred workflow centered around OpenEXRs. But if they really mean video, not linear, there are some serious ramifications for your visual effects and color grading workflow.

Color Spaces and Color Transfer Functions

An RGB value is meaningless without a color space and transfer function. The color space defines the gamut. The transfer function defines how tonal values are encoded.

Always state both. P3 primaries, for example, may use gamma 2.6 for theatrical projection, the sRGB curve for Display P3, or PQ for HDR mastering. Naming only “P3” leaves the encoding ambiguous.

For a more detailed explanation of color science and the processes involved, see *Cinematic Color* by Jeremy Selan.

Apply these concepts through a documented color-management system. See ACES and Color Management and OpenColorIO.

Color Spaces

Color spaces define the gamut a camera can encode or a display can reproduce.

Display Color Spaces

Color space	Description
sRGB	The standard color space used by most computer and mobile phone displays, with an implied approximate gamma of 2.2.
Rec.709	The standard HDTV color space and the reference target for home video and broadcast mastering. The same primaries and white point as sRGB. It specifies a camera-side OETF, not a display gamma: the reference display EOTF is BT.1886, approximately 2.4.
Rec.2020	The wide-gamut UHDTV standard and, through BT.2100, the colorimetric container for essentially all current HDR delivery, including HDR10, Dolby Vision, HLG broadcast, UHD Blu-ray, and HDR streaming. No shipping display reproduces the full gamut because its primaries are monochromatic. Content is delivered *in* the Rec.2020 container but almost always mastered to a P3-D65 volume inside it.
DCI-P3	The standard gamut and white point used in theatrical digital cinema projection: standardized as **P3DCI** (see below). Substantially wider than Rec.709 (about 36% more area in CIE 1931 *xy*, mostly in the greens and reds), with an implicit gamma of 2.6 and the DCI white point (0.3140, 0.3510).
Display P3	P3 primaries and a D65 white point with the **sRGB** transfer function: a registered ICC/CSS encoding and the default display color space on Apple hardware since 2015. Not to be confused with DCI-P3 (gamma 2.6, DCI white) or P3-D65 PQ (HDR mastering). They share primaries and nothing else.

Three P3 color spaces are standardized by **SMPTE ST 2113**, sharing one primary set and differing only in white point. The standard also fixes the naming convention: “P3” followed by the white-point name, with no separators, so `P3D65` is the conformant spelling and `D65P3` is not:

Color space	White point	Description
P3D65	0.3127, 0.3290	P3 primaries on a D65 white. Common for HDR home video mastering.
P3D60	0.32168, 0.33767	P3 primaries on the ACES white point (ST 2065-1).
P3DCI	0.3140, 0.3510	P3 primaries on the DCI reference projector white: the “DCI-P3” of the table above.

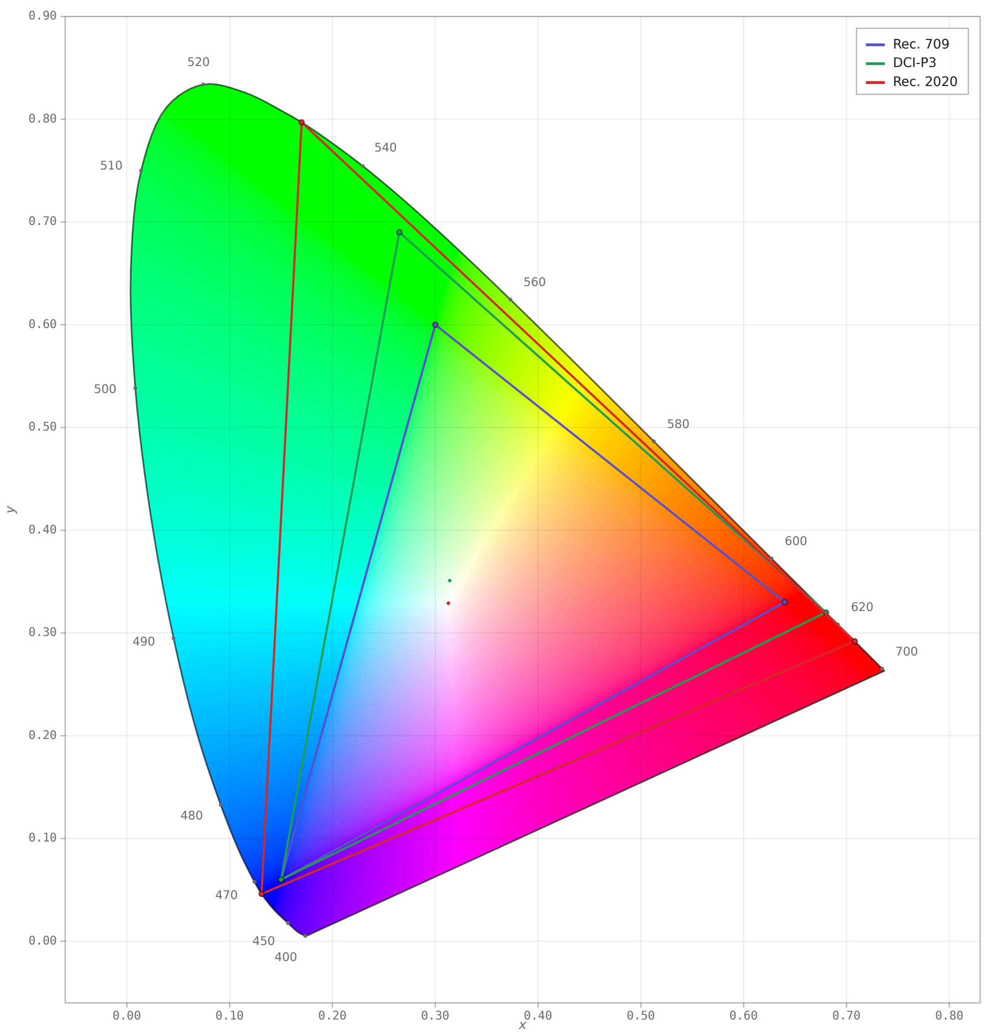

A comparison of display color spaces

Figure 18: A comparison of display color spaces. Interactive: drag to pan, scroll to zoom. Static version.

Camera Color Spaces

Camera gamuts are defined as “virtual primaries” as they are not necessarily indicative of color ranges a camera sensor can register. The primaries often exceed the human vision spectrum as plotted on a CIE 1931 chromaticity diagram. Comparing the virtual primaries of different cameras does not indicate whether a camera actually has broader sensitivity than another, but rather serves as a comparison of the color encoding models used by each camera.

Color space	Description
ARRI Wide Gamut 3 (AWG3)	The acquisition gamut of AMIRA and ALEXA models through ALEXA Mini LF, paired with LogC3.
ARRI Wide Gamut 4 (AWG4)	The wider acquisition gamut introduced with the ALEXA 35 under REVEAL color science, paired with LogC4. Also used by ALEXA 35 Xtreme and ALEXA 265.
REDWideGamutRGB	A camera color space designed to contain all colors a RED camera can capture without clipping. It has replaced the earlier DRAGONcolor / DRAGONcolor2 encodings.
Sony S-Gamut3.Cine	The practical de facto working space for Sony professional cinema cameras: a slightly smaller, more grading-friendly variant of the fuller S-Gamut3.

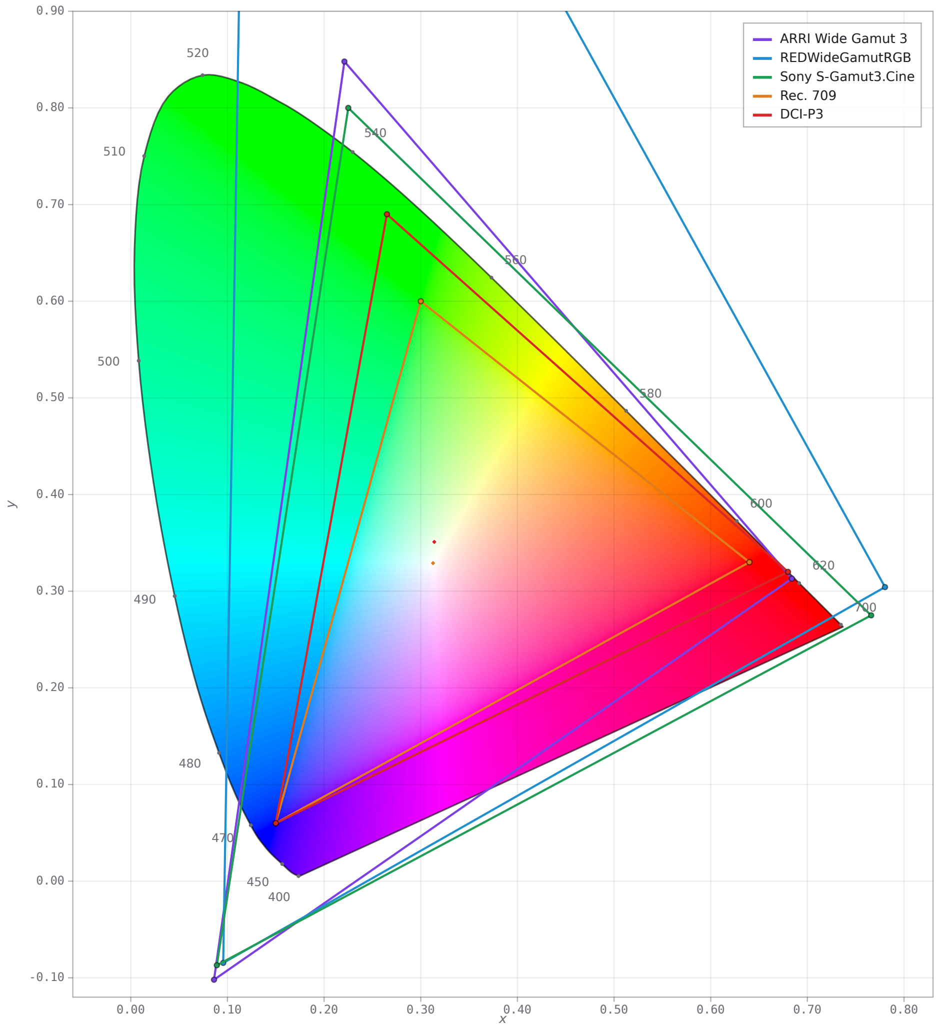

A comparison of camera color spaces and delivery color spaces

Figure 19: A comparison of camera color spaces and delivery color spaces. Interactive: drag to pan, scroll to zoom. Static version.

Transfer Functions

Transfer functions define the relationship between RGB tonal values and their equivalents either in the “scene” if scene-referred, or on a display if display-referred. Transfer functions are often called color encodings, encoding spaces, and gamma spaces.

Camera Transfer Functions

Transfer function	Description
Cineon	The standard log encoding of film scans.
ARRI LogC3	The exposure-index-dependent log encoding used by ARRIRAW and ProRes from AMIRA and ALEXA models through Mini LF.
ARRI LogC4	The EI-*independent* successor introduced with the ALEXA 35. Different mid-gray placement (0.2784 vs LogC3's ~0.391), and ARRI specifies 12-bit as the minimum quantization: never 10-bit. Always state the version number.
REDlogFilm	RED's Cineon-style log curve from the legacy (pre-IPP2) pipeline. RED describes it as "RED's version of the Cineon transfer function" rather than an exact match. In IPP2 it survives only as a selectable output gamma, not as the acquisition encoding.
RED Log3G10	An alternate log curve optimized for HDR grading applications. (Also see Log3G12.)
Sony S-Log3	The latest logarithmic encoding utilized by Sony professional cinema cameras.

Display Transfer Functions

Transfer function	Description
Gamma 2.2	Material mastered for sRGB computer displays. Viewed in a normal, bright surround (office or home interior).
Gamma 2.4	A target for material mastered for Rec.709 displays. Viewed in a dim surround (dim room with low ambient light).
Gamma 2.6	Digital cinema projection gamma. Viewed in a dark room (theater, little or no ambient light).

Linear Transfer Functions

Transfer function	Description
Scene-Linear	Gamma 1.0 with middle gray encoded at 0.18. Floating point values exceed 1.0.
Normalized Linear	Gamma 1.0 with arbitrary mapping of signal bounded between 0.0 and 1.0.

DCI-X'Y'Z'

DCI-X'Y'Z' is a device-independent container rather than a conventional RGB color space. Digital Cinema Initiatives[6] selected it so a DCP can map directly to compliant projectors without a manufacturer-specific color conversion.

DCI-X’Y’Z’ uses a 2.6 power encoding[7] with headroom for multiple white points.

Use DCI-X’Y’Z’ only as a display-referred digital-cinema delivery format, not as a working space. DCDMs and DCPs use this encoding.

!!! note “Why the normalization constant is 52.37” SMPTE ST 428-1 normalizes X’Y’Z’ by 52.37 rather than the 48 cd/m² reference luminance. This preserves 48 cd/m² while making room for D55, D61, and D65 white points. At the DCI projector white, Y’ reaches code value 3960, leaving 135 values before the 12-bit maximum.[8]

Look Up Tables (LUTs)

LUTs are pre-computed color transformations that can contain technical color space conversions, or creative transformations from scene-referred to display-referred states.

LUTs vs. Transforms

A LUT and a transform can describe the same color operation two different ways, and the difference matters when you decide which to use.

A **transform** is the operation expressed as *math*: a matrix, a curve or formula, or a chain of them, evaluated live on every pixel at full precision. An ACES Output Transform, an OCIO config's color-space conversions, a DCTL or CTL program, and DaVinci Resolve or Baselight's native color-management nodes are all transforms. Because they are computed, they are **exact, resolution-independent, and invertible** wherever the underlying math allows.

A **LUT** is that operation *sampled and baked* into a table of input→output values on a fixed grid. Anything between grid points is interpolated. A LUT is **fast, portable, and self-contained**. It needs no knowledge of the math to apply, which is exactly why it is the lingua franca for exchanging a look between systems, devices, and vendors. The trade-off is **precision**: a LUT is only an approximation of the transform, bounded by its grid resolution, the input encoding it was built for, and how it treats values outside that range.

Use a transform when you can and a LUT when you must. Prefer a live transform inside a single application where accuracy matters. Fall back to a baked LUT when a look has to travel to something that cannot run the transform: an on-set monitor, a DIT box, an editor or vendor on a different platform. The Show LUT exists precisely because everyone from the camera cart to the VFX vendor needs the *same* look, and not all of them can run the same transform engine.

Creative Transformations

The umbrella of creative transformations includes any s-shaped curves used to tone map a scene-referred image into a desirable display-referred image. These transforms are subjectively designed to give the image an aesthetic look. Any time you are viewing a camera image in a display-referred setting (e.g. ALEXA camera monitoring in Rec.709 on a studio monitor or viewfinder) you are viewing through a creative transformation.

Film print emulations and other stylistic color correction looks can be baked into LUTs for convenient use and exchange between various software and hardware.

We commonly reference a scene-to-display transform as a "forward" function, while we refer to any display-to-scene transform as an "inverse" function. While it is possible to derive inverses of many functions, once an image undergoes an s-shaped transformation in which values are clipped or multiple input values mapped to a single output value, the inverse will not yield the full dynamic range or a true scene reference.

Inverses are commonly used for converting display-referred graphics so they appear correct when viewed under a global Show LUT, or to convert a display-referred image to log for film-out.

The Show LUT

A Show LUT is the project’s primary scene-to-display creative transform. Grades occur beneath it, and the same look drives on-set monitoring, editorial, VFX review, and finishing. Choose or design it with the cinematographer and colorist during pre-production. It may combine tone mapping, gamut shaping, and film-print emulation.

Show LUTs are display-referred, so they are intended for specific mastering displays and color spaces. Variations of the Show LUT will be tailored for the various display methods and deliverables applicable to a given project (e.g. Rec.709, DCI-P3, Rec.2020).

Examples of common creative transformations:

- ARRI LogC to Rec.709 (K1S1)
- `Slog3SGamut3.CineToLC-709TypeA`
- REDlogFilm to REDgamma4

Technical Transformations

Typically, transformations that do not cross the barrier between scene-referred and display-referred image states are “technical” transformations. Technical transformations rely on published math formulas and defined standards, so there is no subjectivity involved.

Technical transformations on display-referred images are necessary to convert between different display types (e.g. digital cinema to home video).

A common display-referred technical transformation: converting one display standard to another, such as home video to digital cinema:

Similarly, transformations on scene-referred data can convert from camera log to scene-linear, for compositing, and back again without compromising meaningful colorimetric data. Additionally, it is possible to convert from one camera’s native log encoding and color space into another.

A common scene-referred technical transformation: matching one camera’s native log encoding and gamut to another’s:

Use technical transforms to normalize different cameras before intercutting or combining plates. Apply the same transforms to VFX pulls and the corresponding DI media.

When intercutting Camera RAW with scene-linear OpenEXR visual effects shots, it is common to transform the scene-linear visual effects into the camera log encoding (or vice versa) to maintain continuity in the working environment. This makes the visual effects shots match the original camera media and allows for grades to be copied between VFX shots and non-VFX shots with predictable results.

Types of LUT: 1D, 3D, and shaper

Not all LUTs are equal, and the differences decide which one is right for a job.

1D LUTs map each channel independently: one input value to one output value, per channel. They are ideal for anything that acts on a single axis: transfer-function and gamma changes, exposure and contrast curves, log-to-linear conversions. What a 1D LUT *cannot* do is change one channel based on another, so it cannot remap a **gamut** or carry a cross-channel creative look.

3D LUTs sample the full RGB cube: a grid of input RGB triplets mapped to output RGB triplets, so they *can* mix channels and therefore carry gamut conversions and complete creative looks. (A Show LUT is a 3D LUT.) The catch is resolution. The cube is sampled on a coarse grid, commonly **17^3, 33^3, 64^3 or 65^3** points, and everything between grid points is **interpolated** (tetrahedral interpolation is more accurate than trilinear along the diagonals a color transform stresses). A larger cube is more accurate but bigger and slower.

Common practice is 17^3 for on-set boxes and 33^3 or larger in post, but treat those as conventions rather than rules. A sufficiently smooth transform is well served by 17^3 at any stage, some on-set hardware handles larger cubes perfectly well, and what actually matters is whether the cube resolves the transform you are asking it to carry: a LUT with a hard tonal break or a saturated gamut clip needs more grid points than a gentle film emulation does.

Shaper LUTs (pre-LUTs). A 3D LUT's grid is evenly spaced in its input encoding, so feeding it wide-dynamic-range **log** or **linear** data wastes most of the grid on values you don't have and starves the range you do and produces banding. A **1D shaper LUT** applied first redistributes the input into a more perceptually even space before the 3D LUT samples it. This is how log and HDR material is fed through 3D LUTs cleanly, and several exchange formats bundle a shaper with the 3D LUT for exactly this reason.

Formats and precision. LUTs travel in a zoo of formats: `.cube` (Resolve/Adobe), `.3dl` (Autodesk/Lustre), `.csp` (cineSpace, which carries a shaper), and others. Few are universal, which is why the LUT-management tools below exist. The **Academy/ASC Common LUT Format (CLF, `.clf`)** is the modern standard-track answer: an XML format that chains 1D and 3D LUTs *and* parametric operators with defined bit depth and range, and is the interchange format used by ACES. Whatever the format, a LUT is only valid for the **input encoding and range it was built for**. Apply a Rec.709 LUT to log material, or push values outside its domain, and it will clamp or misbehave.

LUT Tools

Grading systems generate LUTs in vendor-specific formats that are not universally compatible. Distribute the Show LUT in a format supported by every application and monitoring device.

Lattice and ColourSpace convert LUT formats and perform precise color-space transformations.

Common Working Formats

Movie-Based Formats

Use frame-based sequences for VFX plates, intermediate renders, and film scans. ProRes 4444 QuickTimes are common on smaller productions and may be acceptable for final VFX delivery, but avoid movie files for intermediate CGI and compositing renders.

Test a full round trip between production, VFX, and the DI before turnover. Check for unintended color or transfer-function changes during transcode.

Aside from visual effects, various flavors of ProRes, DNxHD, and DNxHR are suitable for offline editorial and mastering.

Apple ProRes

Apple ProRes codecs are resolution-independent, professional video codecs most often seen in QuickTime (`.mov`) wrappers. The bitstream itself is published as **SMPTE RDD 36**, with carriage in MXF standardized as **RDD 44** and the IMF application as **RDD 45**. Avid Media Composer manages ProRes as MXF OP-Atom, and Final Cut Pro has exported ProRes-in-MXF since version 10.3, so ProRes inside an IMF is a realistic delivery route, not an exotic one.

Variant	Encoding	Typical use
ProRes 4444 XQ	12-bit 4:4:4 (RGBA / Y′CbCrA), with alpha	Highest quality: camera acquisition, intermediate grading, VFX plates and delivery
ProRes 4444	12-bit 4:4:4, with alpha	Camera acquisition, intermediate grading, VFX plates and delivery
ProRes 422 HQ	10-bit 4:2:2	High-quality mastering and delivery
ProRes 422	10-bit 4:2:2	Mastering and delivery
ProRes 422 LT	10-bit 4:2:2	Lighter editorial
ProRes 422 Proxy	10-bit 4:2:2	Offline dailies and editorial proxies

The **4444** variants sustain multiple generations of encoding without meaningful image degradation and are visually indistinguishable from equivalent uncompressed alternatives, which is why they serve as camera acquisition,[9] intermediate grading, and VFX plate and delivery formats. See the Apple ProRes white paper for details.

ProRes also has a RAW variant (**ProRes RAW**, 2018), but it is a camera-acquisition format only, like ARRIRAW or R3D, and is not applicable to transcoded VFX plate pulls.

!!! warning “Don’t encode delivery ProRes with FFmpeg” FFmpeg’s ProRes encoders are reverse-engineered and unlicensed. A file may look correct but fail delivery QC on metadata or encoder identification. Encode deliverables with an Apple-approved implementation in Resolve, Premiere, or another professional application.

```
FFmpeg remains suitable for disposable dailies, references, and internal proxies.
```

Avid DNx

At IBC 2025 Avid rebranded the DNxHD and DNxHR families to a single name, **Avid DNx**. It dropped the `HD`, `HR`, and `GX` acronyms, which were misleading once the codec became resolution-independent, in favor of five quality levels. The rebrand also made **8- to 16-bit encoding and an optional alpha channel available at every Avid DNx level**. The level indicates the chroma sampling and bandwidth (compression), not a fixed bit depth:[10]

Level	Chroma	Typical use
444	4:4:4 / RGB	Finishing, VFX delivery, cinema masters
HQX	4:2:2	Color and mastering. Multi-generation work
HQ	4:4:4 / 4:2:2 / 4:2:0 / RGB	High-quality editorial and production
SQ	4:4:4 / 4:2:2 / 4:2:0 / RGB	Editorial and broadcast delivery
LB	4:2:2 / 4:2:0, low bandwidth	Offline editorial and proxies

Data rates scale with resolution and frame rate. From Avid's published figures, at **UHD (3840×2160) / 23.976 fps** they run roughly **167 MB/s** (444), **83 MB/s** (HQX and HQ), **55 MB/s** (SQ), and **17 MB/s** (LB): about 1.3 Gb/s down to 137 Mb/s. At HD they are roughly a quarter of that (444 ≈ 42 MB/s, LB ≈ 4 MB/s). Avid Media Composer manages DNx media natively in **MXF OP-Atom**, and Avid also ships DNx **QuickTime** codecs. The codec is standardized as **SMPTE ST 2019-1 (VC-3)**.[11]

H.264 / H.265 (HEVC)

Treat H.264 and H.265 as distribution formats, not VFX or DI working formats. Their long-GOP compression stores many frames as differences from neighboring reference frames, making frame-accurate access slower and less reliable for post.

They are, however, perfectly good **creative-review proxies**. Much VFX review happens in a browser on platforms like Frame.io or cineSync. Whatever the quality of the file that was uploaded, the player streams an H.264 or H.265 transcode anyway. For a director or supervisor giving creative direction, that is plenty. They can judge staging, timing, and the look. It is not a substitute for final QC sign-off, which belongs on the graded shot in the DI, but for broad creative approval an H.264 / H.265 proxy is entirely sufficient.

In practice, consumer H.264 and H.265 deliverables are 8-bit and 10-bit 4:2:0, respectively, because that is what hardware decoders and streaming platforms support, but both standards define higher-bit-depth, higher-chroma professional profiles (H.264 High 10 / Hi444PP up to 14-bit, and HEVC Main 12 / Main 4:4:4 12) used in mezzanine and archival contexts. Neither commonly supports an embedded alpha channel.

These codecs are commonly found in `.mov`, `.mp4`, `.m4v`, and `.mkv` file wrappers.

AV1 and H.266/VVC

AV1 and H.266/VVC are downstream distribution codecs, not post-production working formats.

AV1 (AOMedia, royalty-free) is the significant one for streaming. It has broad browser and hardware support, though Safari depends on AV1 support in the underlying device. By late 2025, a large share of Netflix's streaming hours were served in AV1. For an independent production it matters as a *destination* your master gets encoded into by the platform, not something you deliver, edit, or master in.

H.266/VVC is far more niche, with little consumer hardware decoding and no browser support. Its real traction is in broadcast. Brazil's TV 3.0 / DTV+ mandates VVC and began 4K HDR over-the-air in 2025). An independent film will not deliver in it.

Do not acquire, work, or master in either format.

Frame-Based Formats

OpenEXR (.exr)

OpenEXR (more commonly referred to as EXR) is an open-source format developed by Industrial Light and Magic (ILM) as a replacement for legacy image formats commonly used in visual effects. EXR is designed for multichannel 16-bit and 32-bit floating point: almost always in scene-linear encoding.

EXR supports alpha, multiple layers, and arbitrary channels in one file. It offers uncompressed, lossless, and lossy compression schemes.

The choice of compression scheme matters in a digital intermediate: some are cheap to decode and play back in real time, others are not, and the **ACES2065-1 archival container (SMPTE ST 2065-4) permits only uncompressed data** regardless. For final EXR renders delivered to the DI, uncompressed or a fast lossless scheme (ZIP / ZIP1) is the safe choice. Confirm with your DI facility or colorist before using anything heavier.

Compression	Lossless	Real-time DI playback	Notes / when to use
None (uncompressed)	✓	✓	Largest files, no decode cost. **The only option permitted in the ACES2065-1 archival container** (ST 2065-4).
ZIP / ZIP1 (ZIPS)	✓	✓	zlib. Very efficient to decode and no real problem in modern DI systems: a safe default for VFX renders to the DI.
RLE	✓	✓	Run-length. Light, fast, modest savings.
PIZ	✓	✗	Wavelet + Huffman. Excellent on grainy / noisy plates, but slower to decode and can impact realtime playback on some systems.
PXR24	✓ half / ✗ 32-bit	✓	Good ratio. Lossless for half-float and integer, lossy for 32-bit float.
B44 / B44A	✗ (lossy)	✓	Fixed-rate, designed for fast real-time playback.
DWAA / DWAB	✗ (lossy)	✓	DCT, JPEG-like. The smallest files and fast to decode. DWAA uses 32-scanline blocks (better for partial reads, as Nuke does). DWAB uses 256-scanline blocks and decodes full frames faster. Lossy: avoid for final DI renders.
HTJ2K	✓	✓	High-Throughput JPEG 2000 (ITU-T T.814 / ISO/IEC 15444-15), added in OpenEXR 3.4. Two variants: `htj2k256` (smallest files) and `htj2k32` (much faster encode/decode). Among the fastest schemes in OpenEXR, but new, so confirm your facility's OpenEXR version supports it before delivering.

When in doubt for a DI delivery: uncompressed, or ZIP / ZIP1.

EXRs with multi-channel render passes (AOVs: arbitrary output variables) such as normal maps and z-depth are especially advantageous in visual effects workflows. However, they are rarely included in final renders delivered to the DI, as those additional render passes are intended for lighting and compositing and may serve little benefit to the colorist and will only increase the size of the files.

EXRs, unlike most image formats, may contain image data (the "data window") outside the visible frame area ("display window"). Conversely, a data window smaller than the display window is also supported. This feature is designed for CGI compositing, particularly when producing large-kernel blurs or 2D camera shake effects during compositing. Most DI systems are incompatible with EXRs containing mismatched data and display window sizes. VFX vendors must be informed that their final renders delivered to the DI must have matching data and display window sizes. In some DI systems, mismatched window sizes can lead to image artifacts and in some cases prevent the software from reading the files at all.

OpenEXR display window vs. data window

Figure 9: OpenEXR display window vs. data window.

DPX (.dpx)

DPX has been the standard file format for digital intermediate scanning, color grading, and finishing for over thirty years. SMPTE 268M dates to 1994, and DPX remains a common working format for digital intermediates as well as an intermediate format for transcoded Camera RAW media. Note that SMPTE has since designated ST 268-1 a **Stable** document: still in force, but formally flagged as no longer representing current technology and not subject to further review. In practice DPX still owns film-scan delivery and log-encoded conform, while ACES/OpenEXR (ST 2065-4) has taken over scene-referred interchange and archival mastering. It is the successor to the legacy Cineon format (`.cin`) which preceded it as the standard film scanning format.

DPX file headers can store embedded timecode and tape name for assisting in conform.

Image data is routinely stored as fully uncompressed RGB 10-bit or 16-bit integer. Although the DPX specification defines support for floating point data types, many software and digital intermediate systems do not support floating point DPX. Alpha channel support is defined in the specification, but many systems do not support DPX with embedded alpha.

DPX files are advantageous because in practice they are written uncompressed and, unlike lossy compressed media, retain image fidelity across successive passes. (SMPTE ST 268-1 does define an optional run-length encoding, but essentially nothing in a post chain writes it, so treat DPX as uncompressed in planning.) File sizes are easily predicted, as all files of a given raster dimension, bit depth, and header length will be the same size, regardless of the image content.

TIFF (.tif, .tiff)

TIFF is an older image format that, within the context of visual effects, is largely a legacy format. DPX and EXR displaced it for plates and comps. It is emphatically not obsolete in mastering, though: **16-bit TIFF is the standard carrier for X’Y’Z’ DCDMs** (Digital Cinema Distribution Masters), which is about as current a use as a format can have. Aside from DCDMs, some studios prefer 16-bit DCI-P3 RGB TIFFs over DPX as an archival master, but this is quickly being replaced by EXR.

TIFFs and other legacy graphics formats[12] are often used for simple title graphics that are comped over picture during the DI. In these cases, TIFF is chosen for its simplicity, as it supports an alpha channel and is readily compatible with Photoshop, where most of those graphics are generated.

TIFFs should not be used as an intermediate working format for either visual effects or digital intermediates, as DPX and EXR both provide advantages over TIFF.

TIFFs are commonly either RGB 8-bit or 16-bit integer, while 24-bit and 32-bit floating point variants exist, with support for embedded alpha channel, but without support for tape name or timecode metadata.

TIFFs can be uncompressed, losslessly compressed, or lossy compressed in a variety of ways, and support multiple layers of image data.

Common Working and Delivery Resolutions

A project may combine camera rasters, aspect-ratio guides, working resolutions, and delivery containers. Document each one. Do not use “2K” or “4K” without exact pixel dimensions.

Camera Original Media

Cameras record many native, cropped, and downscaled rasters. There is no universal acquisition resolution or aspect ratio. Record the exact raster, sensor mode, pixel aspect ratio, and framing guide for every camera.

Delivery Formats for Video

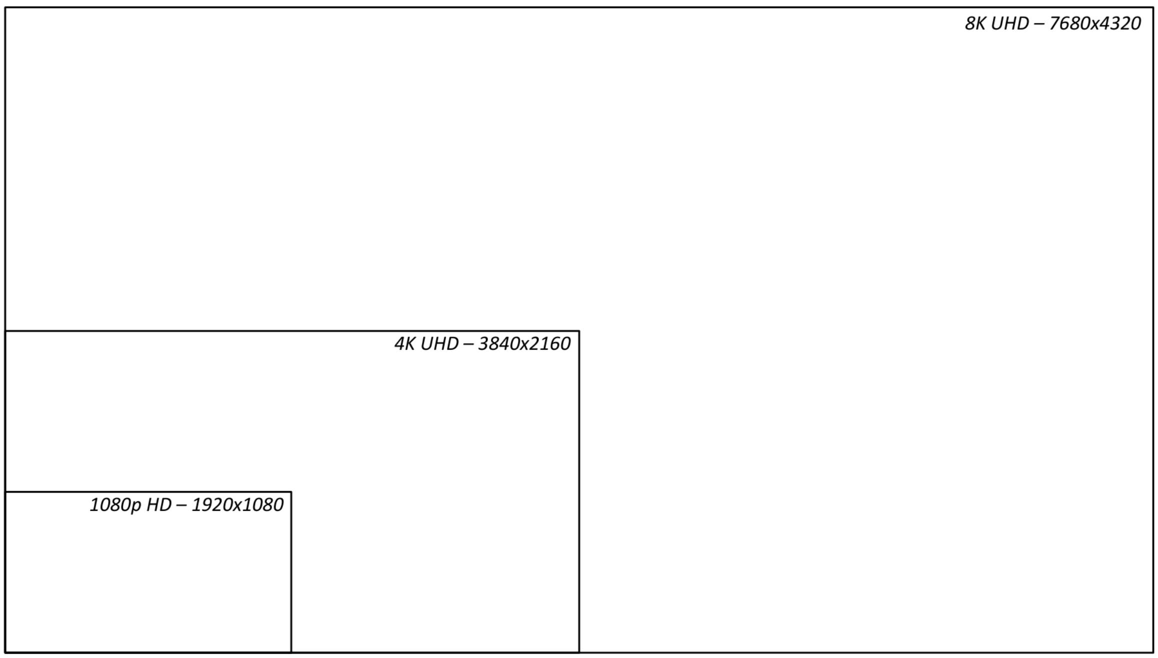

Comparison of HD and UHD resolutions

Figure 10: Comparison of HD and UHD resolutions: 1080p HD (1920x1080), 4K UHD (3840x2160), and 8K UHD (7680x4320).

High-Definition

Most high-definition deliverables are made at 1080p (1920x1080) at 23.976 frames per second, though 24.000 is equally legitimate, and the choice between them is worth making deliberately before you shoot. Rarely is feature film content produced at 720p, but some broadcast or cable networks may require it. In most cases the 720p sub-masters are derived from 1080p masters.

1080p HD is 1.78:1 (16:9). Letterbox wider compositions inside it, commonly 2.39:1 (Scope) or 2.00:1.

Ultra-High-Definition

The UHD (2160p) standard is prevalent in new consumer television sets and a variety of computer displays: equivalently four times the resolution of 1080p. Movie studios as well as OTT (over-the-top) distributors and content creators are producing films and television content with the goal of concurrent or eventual UHD distribution. Upscaling a 2K-acquired finish to UHD was common a decade ago, but the major streaming platforms now require true-4K acquisition for originals: Netflix, for example, mandates a sensor at least 3840 photosites wide for 90% of a programme's runtime. Plan acquisition at or above your UHD deliverable rather than expecting to upscale. The UHD equivalent of 4K is 3840x2160, not truthfully "4K", but it is referred to by many manufacturers as "4K UHD".

8K UHD (7680x4320) remains a niche *distribution* format: NHK's BS8K, launched in December 2018, is still essentially the only regular 8K broadcast service, and no major streaming platform delivers 8K. 8K-class *acquisition*, by contrast, is now ordinary: RED V-RAPTOR, Blackmagic URSA Cine and ARRI's 65mm bodies all record above 6K, typically to be oversampled down to a 4K or UHD finish.

Delivery Resolutions for Digital Cinema

Digital cinema distribution in theaters follows the Digital Cinema Initiatives' guidelines and specifications for content preparation and delivery.

The DCI 2K image container is a 2048x1080 (1.89:1) signal raster. This aspect ratio and frame size was created to accommodate both Flat (1.85) and Scope (2.39) films.

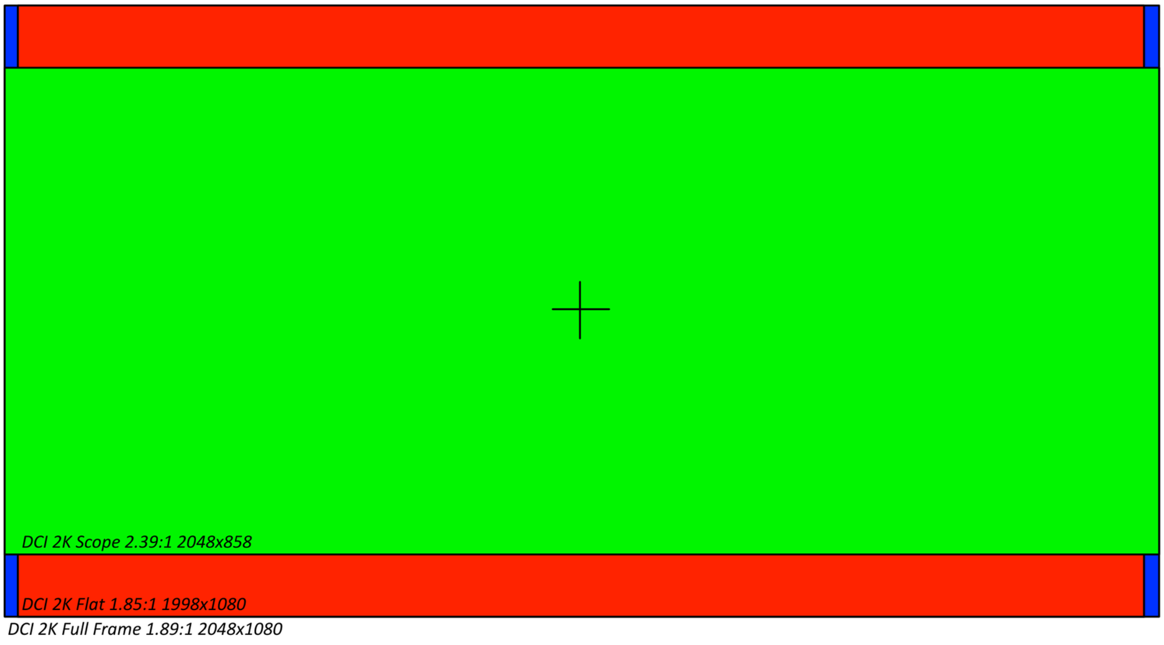

Comparison of digital cinema image rasters

Figure 11: Comparison of digital cinema image rasters: DCI 2K Full Frame 1.89:1 (2048x1080), DCI 2K Flat 1.85:1 (1998x1080), and DCI 2K Scope 2.39:1 (2048x858).

The DCI 4K equivalents are 4096x2160 (Full), 3996x2160 (Flat), and 4096x1716 (Scope).

The vast majority of films are delivered to theaters at either DCI Scope or DCI Flat resolutions. The exceptions are premium large-format screens: IMAX Digital and various PLF and Dolby Cinema installations, which may use the entire DCI Full container. Note that IMAX is not one format: IMAX laser GT projection runs at **1.43:1**, taller than anything in the DCI container set, and is a separate deliverable rather than a crop of the DCP.

Scope

There are some common misconceptions about the “Scope” (CinemaScope) aspect ratio. Depending on who you ask, they may say it is 2.35:1, 2.39:1, or 2.40:1.

Each of these was, at one point, the standardized **projectable image area**: the region of the release print that actually reaches the screen (the negative usually holds a little more picture than the audience ever sees). Note the distinction: the standards specify the *projectable image area*, not the dimensions of the aperture plate that produces it. The history is a sequence of SMPTE standards, not camera-manufacturer decisions:

Standard	Projectable image area	Unsqueezed ratio = (width × 2) / height
PH22.106-1957	0.839 × 0.715 in	(0.839 × 2) / 0.715 = **2.347** (≈ 2.35:1)
PH22.106-1971	0.838 × 0.700 in	(0.838 × 2) / 0.700 = **2.394** (≈ 2.39:1)
SMPTE ST 195-2000	0.825 × 0.690 in (style B max)	(0.825 × 2) / 0.690 = **2.391** (≈ 2.39:1)

The **1971 revision** left the width essentially unchanged and trimmed the height by about 2%. The reason was mechanical, not aesthetic: four-perf anamorphic uses almost the full height between frames, leaving little frameline in which to make splices, so a splice passing the projector gate could flash as a bright band near the top or bottom of the screen. The shorter aperture formalized the projectionists’ practice of masking that area: a thicker hidden frameline at the cost of a sliver of picture height. That is the real origin of “2.39,” and why “the industry decided 2.39 looked better than 2.35” is misleading.

The **1993 revision** (carried into the current **ST 195-2000**, reaffirmed 2012) was an interoperability cleanup rather than a creative change: it set a common **0.825-inch projected width** shared by Flat and Scope, so both formats used the same lateral boundaries while Scope’s ratio stayed at 2.39. Worth knowing that the standard gives Flat’s height as a **range**: style A is 0.446 in minimum to 0.500 in maximum, rather than a single figure, so “1.85 is 0.825 × 0.446” is the tightest case, not the only legal one.

So is it 2.39 or 2.40? 2.39:1 is the correct name for the standardized theatrical format. “2.40:1” is a convenient rounding (it also implies an exact 12:5) that spread through framing charts, telecine, and home-video paperwork. It does not name a separate film format. Digital cinema preserves the same shape: DCI 2K / 4K Scope (2048 × 858 / 4096 × 1716) computes to **2.387**, still called 2.39 Scope, whereas an HD raster like 1920 × 800 is *exactly* 2.40 (a genuine 2.40 digital master, but not the historic SMPTE ratio).

And “2.35”? The name stuck as a cultural synonym for Scope and survived the 1971 change. When someone says “we’re framing at 2.35” today, they may mean a literal 2.350 (slightly *taller* than modern Scope) or simply “Scope” while intending 2.39. For restoration and precise delivery the distinction matters: a literal 2.35 extraction is about 1.9% taller than 2.39 (roughly 0.95% more picture top and bottom on a constant-width image). A pre-1971 title was projected around 2.35, a post-1971 anamorphic release around 2.39.[13]

However today, there is a single digital cinema aspect ratio for super widescreen films, known as DCI "Scope", and its aspect ratio is 2.39:1. Arguably one could frame for 2.35:1 or 2.40:1 and inscribe that within the DCI Scope frame. However, they would need to add pillarboxing or letterboxing to do this, and the difference is only a few pixels. It is highly recommended that if your project may eventually have a life on the big screen, you simply frame and finish at a standard aspect ratio unless the difference is significant.

Online content distribution is a wide open sandbox and any resolution you can inscribe within a standard 1.78 container is technically possible and without detriment, as they will all be letterboxed regardless. Wide aspect ratios are now common in episodic television. 2.00:1 became a streaming house style after *House of Cards*, and full 2.39 Scope series such as *The Mandalorian*, *Andor*, and *The Rings of Power* are routine. Confirm the target ratio with the platform rather than assuming 16x9.

Flat

Historically, features framed 1.85:1 were often opened up to 1.78:1 (16x9) for broadcast and DVD, because distributors disliked the thin letterbox. To achieve this, filmmakers would maintain the width of the frame and remove cropping top and bottom. Where acquisition or mastering constraints prevented that, they would hold the vertical framing and crop the sides instead.

That is no longer standard practice. Netflix, Amazon and the other major platforms require the original theatrical aspect ratio, letterboxed inside a 1920x1080 or 3840x2160 container. Confirm the deliverable aspect ratio with your distributor before deciding whether visual effects renders need to be protected above and below frame: protecting for an opened-up 1.78 extraction is work you may no longer need.

16x9 Widescreen

16x9, or 1.78:1, is not a standard DCI aspect ratio. Technically a valid DCP can be authored and will be accepted by most DCI servers and projectors. However, traditional movie theaters regularly only have macro configurations for Flat (1.85) and Scope (2.39). These affect the projector lens zoom, format and image scaling, as well as theater masking. If a film mastered and framed for 1.78 is to undergo a wide theatrical release, it is advisable to package it within the DCI Flat container with pillarboxing on the left and right sides to ensure maximum compatibility. Check with your exhibition theater when in doubt.

Frame Padding

It is very common for independent productions to frame for the full width of a digital sensor, cropping the top and bottom of the frame in post to achieve their desired aspect ratio. However, with the increase in resolution of modern digital cameras, it is often the practice of studio productions to intentionally frame for a region less than the full sensor size, recording additional pixels outside the intended frame lines. There is no standard procedure for this, as it is a specification decided in part by the post-production supervisor, VFX supervisor, cinematographer, and other workflow professionals on the production team.

This process yields an image with additional padding for image stabilization and stereo 3D depth grading during stereo conversion. Since digital image stabilization is typically a last resort, productions prepare for the inevitability of it by recording this extra frame padding. This allows them to stabilize an image in post without dramatically altering the frame composition, which would otherwise be necessary if they had framed for the entire sensor. The padded region is protected (i.e. free of production equipment, crew, or other unwanted objects[14]).

This is not unlike film acquisition, which has always had a degree of padding. Manufacturers have always left a small pad between the ground-glass frame lines and the film gate.

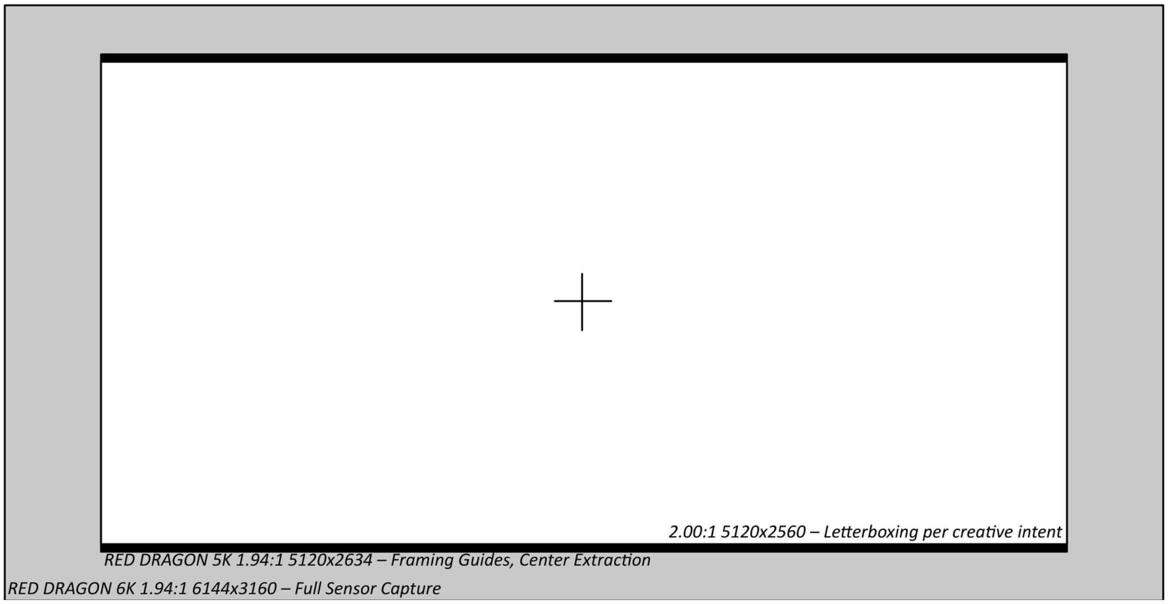

Example frame padding for RED Dragon 6K with 5K center extraction

Figure 12: Example frame padding for RED Dragon 6K with 5K center extraction. Full sensor capture at 1.94:1 (6144x3160), framing guides for a 5K center extraction (5120x2634), and a 2.00:1 letterboxed extraction per creative intent (5120x2560).

The above example depicts a common process used when recording on the RED Dragon in native 6K (6144x3160). The frame guides are set to frame a 2.00:1 center extraction, leaving significant padding in all directions to provide ample resolution for creative reframing, stabilization, and compositing in post-production.

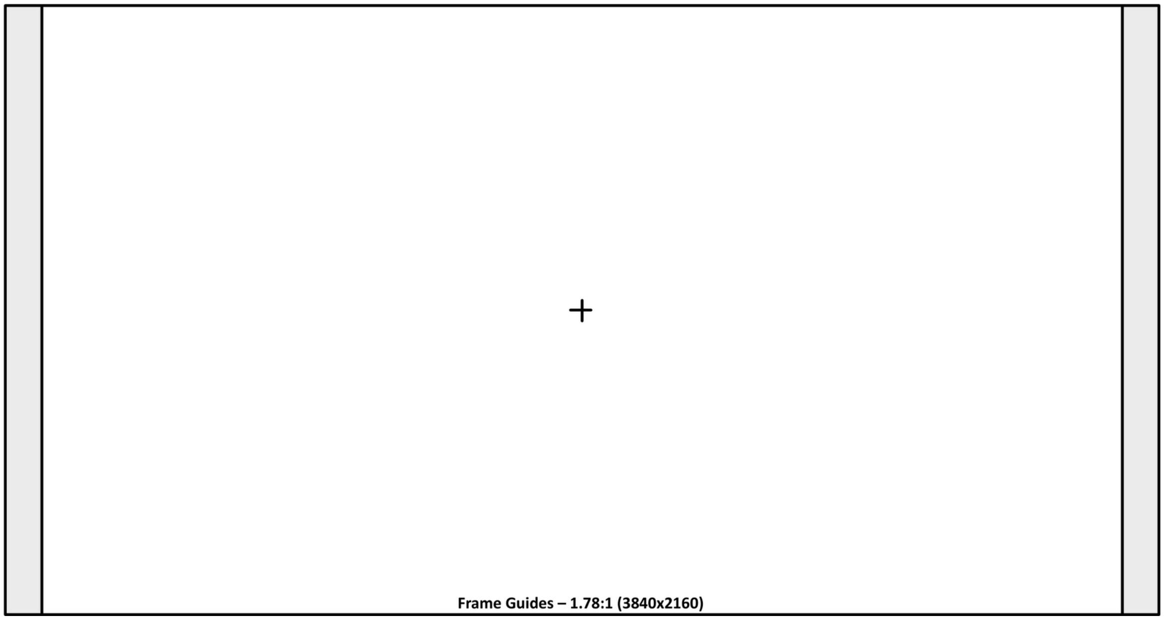

Example frame guides for a UHD extraction from a Sony F55

Figure 13: Example frame guides for a UHD extraction (3840x2160) from a Sony F55, natively recording at 4096x2160.

Productions mastering for 4K and shooting on cameras whose sensor is at or near the delivery raster will regularly frame full height on the sensor. To avoid upscaling, frame padding is usually only employed with cameras that record at resolutions greater than the target delivery resolution.

Image Processing Workflows

Having a precise image processing workflow in place prior to production is critical to producing consistent results throughout production, dailies, editorial, digital intermediate, and final exhibition. Once developed, it requires the cooperation of multiple departments to fulfill, but provides a clear description of expectations from all parties.

The following example process diagram illustrates the various formats and processes an image will undergo from acquisition, to dailies, visual effects, digital intermediate, and distribution.

In this example, the filmmakers are recording 2.8K ProRes (2880x2160) from an ALEXA Mini. This is the camera acquisition format. Due to ProRes recording limitations, the full sensor size is not available for capture in this scenario. They are shooting with anamorphic lenses with a 2.0x anamorphic factor. They are providing a 5% frame padding and framing for a 2.39:1 extraction. This provides extra padding on left and right, as the full sensor yields a 2.67:1 ratio after de-anamorphizing. The eventual goal is to produce 4K deliverables.

The first dailies pass will be produced at 1920x1080 HD, with the de-anamorphized center extraction filling the frame with appropriate letterboxing to 2.39:1. No padded region is viewable in this pass, only the intended framing. This dailies pass is used for all dailies review, distribution to director, producers, and executives, and principal editorial.

The second dailies pass will include the entire recorded image, de-anamorphized, and presented at 2.67:1 in a 1920x1080 HD container. This pass is produced as required by editorial when performing temp stabilizations or creative reframes requiring pixels in the padded area. This pass is not to be distributed in its unedited state.

The VFX plates are pulled at camera acquisition resolution (2880x2160). This is the visual effects scan format. As this project is destined for a 4K finish, the plates are not to be downsampled and VFX are to be completed at native (or best-possible) resolution.

VFX are completed in the camera acquisition (2880x2160) resolution and returned to the DI, where they will be de-anamorphized and scaled to 4K (4096x1716) using the same scaling parameters as any native footage. This is the finishing format. The DI frames for the intended frame guides, producing a 2.39:1 image. The padded regions are available in the event filmmakers request a reframe in the DI.

From the DI, a 4K DCP (4096x1716) is produced as well as UHD and HD letterboxed home video masters.

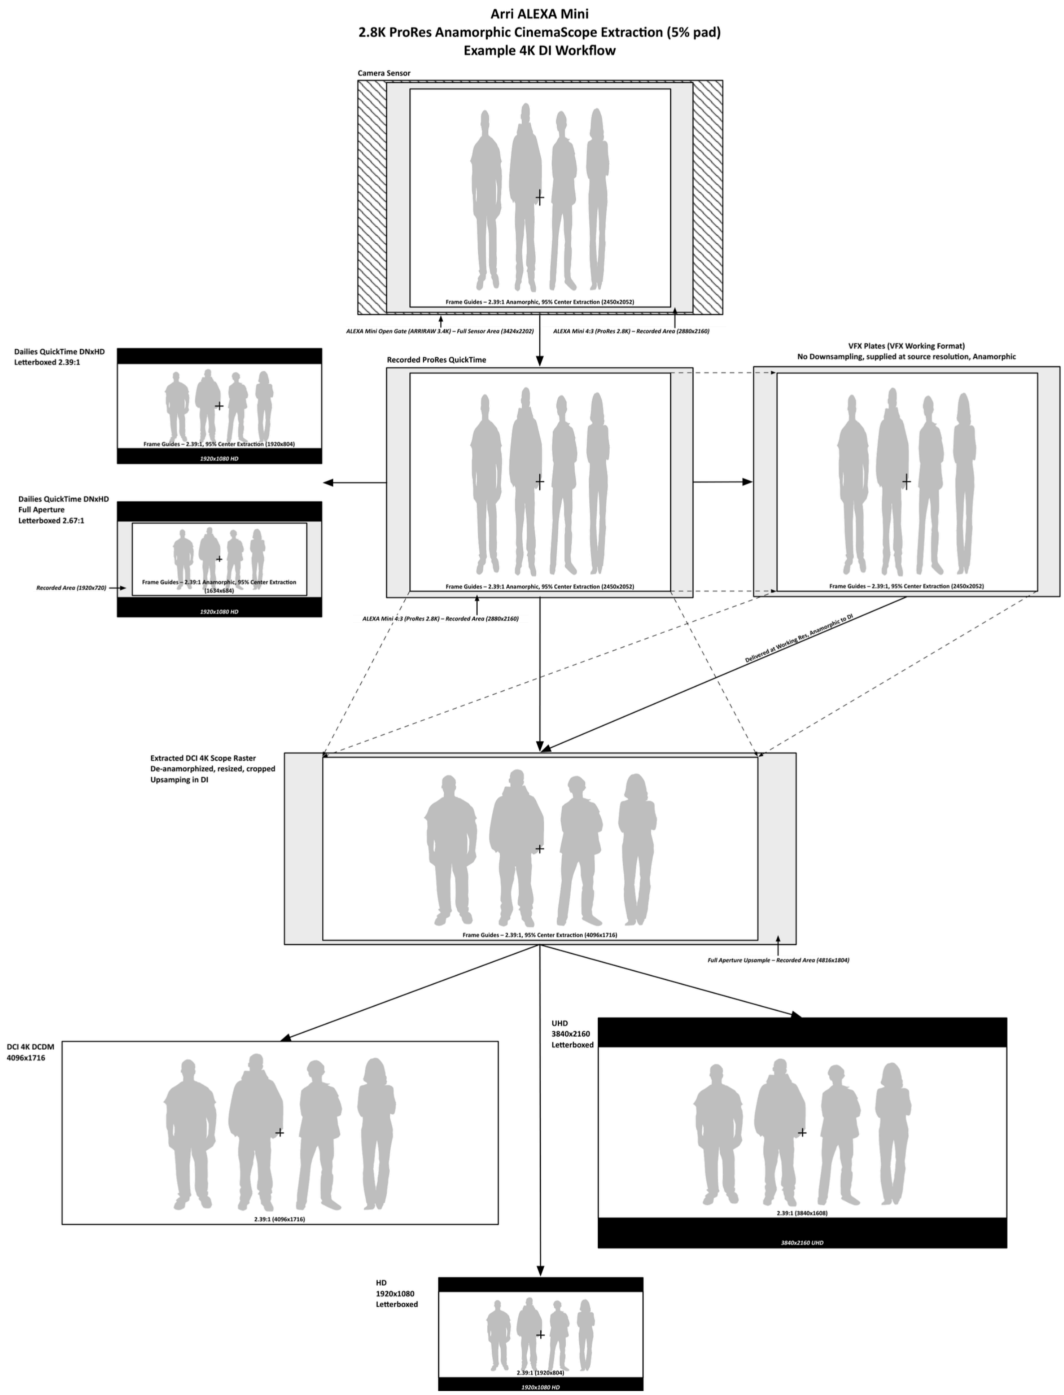

Example ALEXA ProRes anamorphic workflow

Figure 14: Example ALEXA ProRes anamorphic workflow: a 2.8K ProRes anamorphic CinemaScope extraction with 5% padding, carried through dailies, VFX plates, DI, and final DCP/UHD/HD deliverables.

Anamorphic Workflows

When capturing using anamorphic lenses, a de-squeeze scaling factor is necessary to correct the resulting image. Common anamorphic scaling factors, based on the type of lens you shoot with, are 2.0x and 1.3x.

The decision to de-squeeze during a visual effects plate pull or not depends on your available production bandwidth and the implications that will have for your visual effects artists.

Prior to performing visual effects pulls, consult with your visual effects artists, supervisors, vendors, and colorist to decide whether to de-squeeze during the pull, as part of the visual effects final render, or leave it until the DI.

Framing Charts

Framing charts tend to be less common among independent digital productions. However, they are quite common on studio feature films and television shows. Without a framing chart (either photographed in camera or a generated reticle) anyone viewing, working on, or exhibiting the footage is left guessing what the intended framing is. It is not sufficient to simply tell them an aspect ratio (e.g. “1.85” or “Scope”) as that leaves room for interpretation as to how the image is being sized, cropped, and viewed. A picture says a thousand words and a framing chart eliminates the need for words at all. With a framing chart at the head of a reel or as a sidecar reference file, editors, colorists, projectionists, and visual effects artists have a clear guide to the expected framing.

Pixel-accurate framing reticles are often drawn by the DIT or provided by the digital intermediate facility.

During camera prep the camera assistants will shoot a framing chart using each of the production cameras for reference throughout post-production. Furthermore, the DIT or post-production lab should produce a pixel-accurate reticle using Photoshop or other graphics tools, providing information regarding the show’s shooting format, resolution, and intended framing guides. These are better than photographed charts as they are pixel-accurate and not subject to lens distortion.

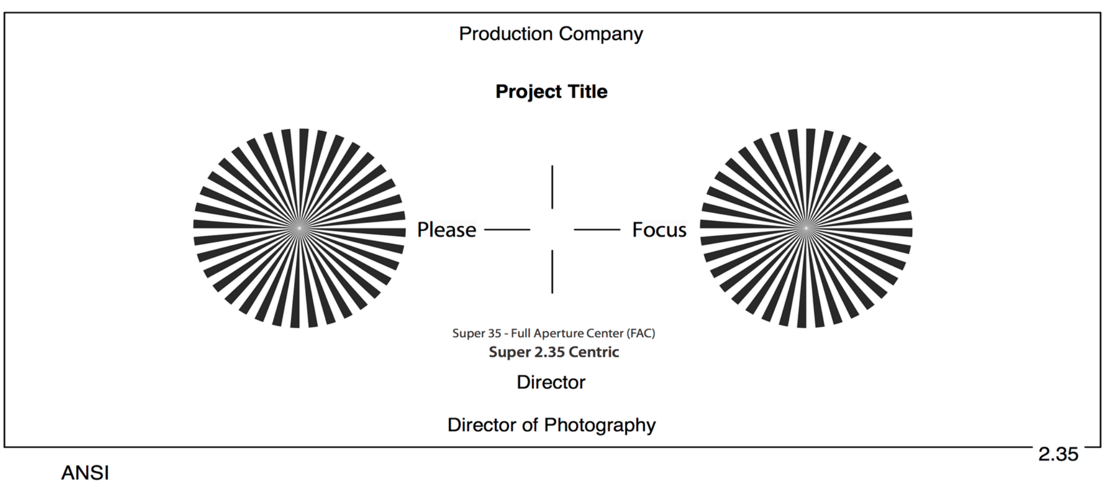

Example camera framing chart (film)

Figure 15: Example camera framing chart (film).

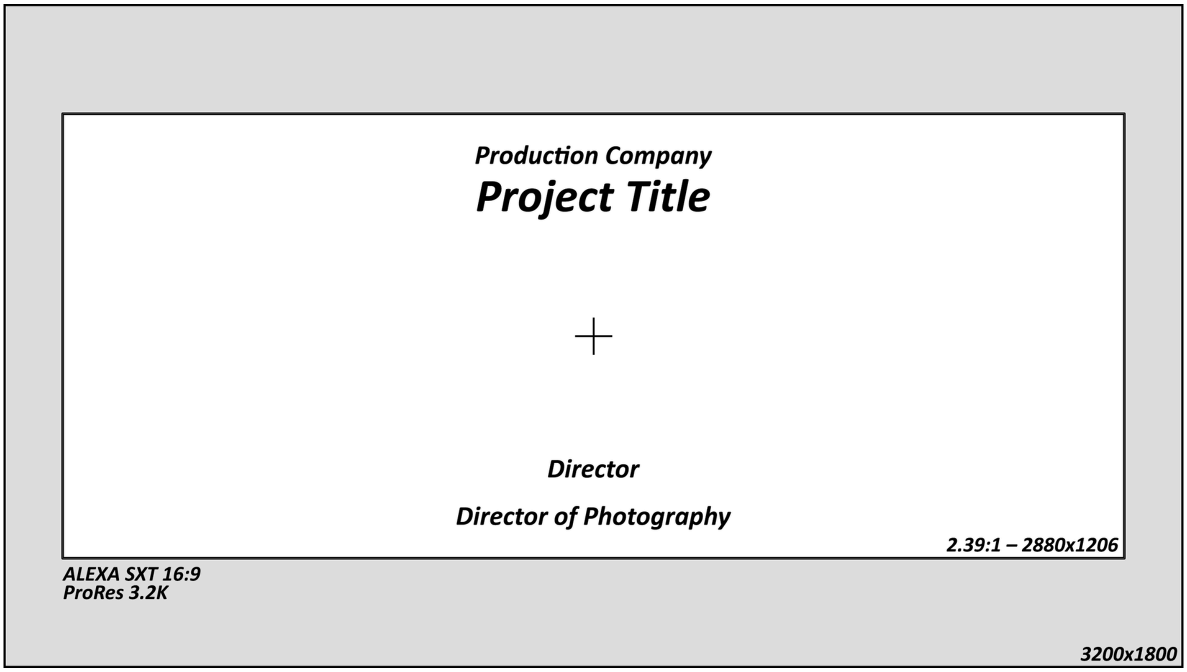

Example camera framing reticle (digital)

Figure 16: Example camera framing reticle (digital), pixel accurate, depicting the image capture area (gray) and the intended framing area (white).

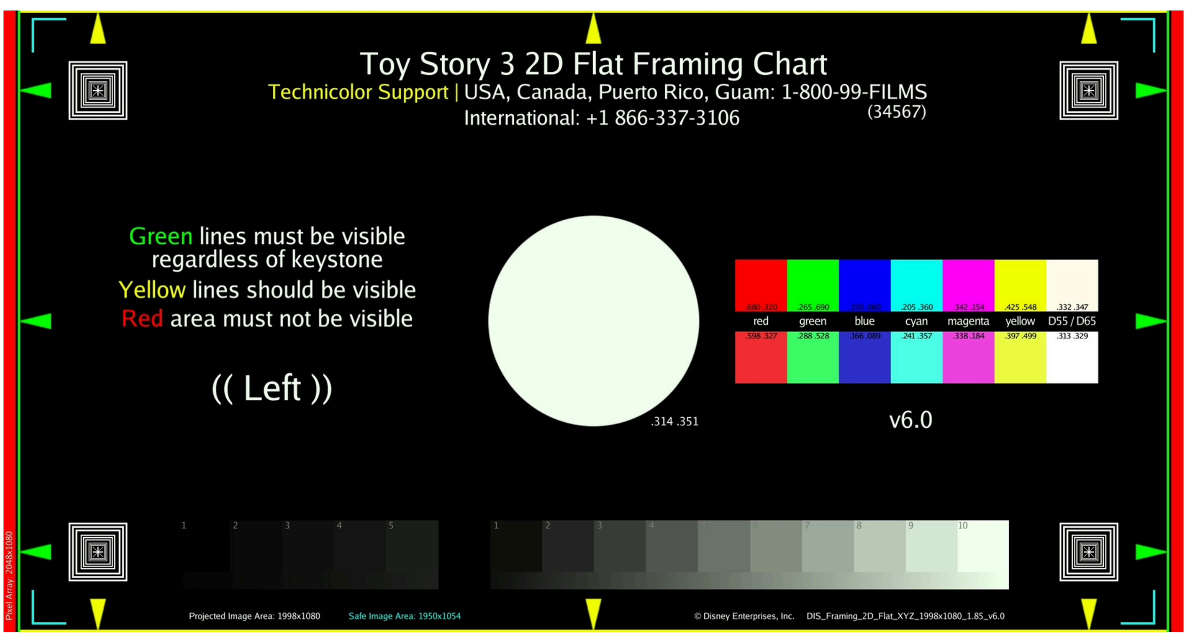

Example digital projection framing chart

Figure 17: Example digital projection framing chart.

Production and VFX

Visual Effects Production Management

VFX production management covers scheduling, assets, progress, vendors, and teams. On large shows, the VFX producer coordinates facilities while each facility manages its artists. On smaller shows, the VFX supervisor may also manage a distributed group of freelancers.

Asset Management

Asset management identifies and names shots, orders plate pulls, organizes media and metadata, and distributes each package to the assigned vendor or artist.

Choose transfer and asset tools around security, bandwidth, file size, and budget. Facilities often use Aspera Faspex, Signiant Media Shuttle, or Sohonet FileRunner. Smaller teams may use Dropbox, Google Drive, or Hightail.

Cloud storage lets production upload once and grant access to multiple artists. This helps teams with limited upstream bandwidth, but does not replace a controlled archive.

Dropbox Business enables granular permissions, allowing production managers to selectively provide artists access to individual plate and render directories on an as-needed basis without granting total access to the production's Dropbox. This provides for higher security and managed control of assets leaving the production. Dropbox also retains file version history and deleted files for a fixed window that depends on the plan: **30 days** on personal tiers, **180 days** on Essentials/Business/Standard, and **365 days** on Business Plus/Advanced/Enterprise, with longer retention available only as a paid add-on.[15] Within that window a file that is deleted or overwritten can be restored without significant downtime or re-uploading from an archive. Treat it as a recovery net for recent accidents, **not as an archive**: a plate overwritten a year ago is gone, which is why the archival copy still has to exist independently.

Many cloud-based file sharing systems include preview and review tools. While visual effects shots are in the creative production phase, a remote artist will upload QuickTime previews to the production's cloud service and the VFX supervisor, director, editor, and other creative executives can review directly in the service's web interface or mobile app and leave creative feedback and notes. This provides for a very fast exchange and review process, often eliminating the need for download and ingest of preview files into a dedicated preview system until it is necessary to perform critical inspection of shots that are presented as potential final versions.

Use **Frame.io** or **cineSync** when review requires timecode-accurate annotations and exported notes. See The Creative Review Cycle.

In addition to uploading renders, artists may also exchange working project files and shared assets via the production's cloud service. This can be useful in collaborations between multiple artists that are not co-located, last minute fixes and tweaks if an artist is unavailable, and facilitating distributed rendering of final shots.

Do not assume a post facility will accept a shared cloud folder. Many require direct, secure transfers and prohibit consumer cloud services. Use the facility's Aspera or Signiant server when available. For unrestricted, one-time deliveries, Hightail or WeTransfer may be sufficient.

For large deliveries, it is sometimes faster to copy large files to a physical drive and send it across town the old fashioned way via courier or production assistant. Overnight shipping through UPS and FedEx can even be faster than transferring large files over the internet.

Production Tracking

Track every shot's owner, assets, versions, notes, milestones, approval state, and current DI version. CGI milestones commonly include layout, animation, lighting, and compositing.

Tools like **Autodesk Flow Production Tracking** (formerly Shotgun, then ShotGrid, and now part of Autodesk Flow) and **ftrack** (acquired by Backlight in 2022) are industry-specific tools for coordinating visual effects projects. They have key features for asset tracking, deadlines, budgets, shot review, annotation, and client feedback, and serve as central collaborative workspaces that production managers and artists alike can leverage to define and maintain their goals.

These tools are especially appealing to large visual effects studios because of their APIs that allow software engineers and technical directors to integrate the management system directly into their automated pipeline. While heavily customizable, they are also now designed to be accessible and friendly to small productions and studios right out of the box, with software tie-ins to common applications like Flame, Nuke, Maya, Cinema 4D, Adobe Premiere, and more.

Many productions elect to build their own shot and resource tracking databases, often using off-the-shelf tools like FileMaker, Microsoft Excel, or Google Sheets. Depending on the size of the production, the number and complexity of shots, and the number of artists and managers involved, these solutions may be the simplest and most cost effective.

These databases include relevant shot information such as shot name, plate(s), original source metadata, shot description, VFX description, notes, production status, assigned artist/vendor, latest version number, approval/review status, and version currently in the DI edit.

The VFX editor and conform editor reconcile the cut against this database. Keep it current.

VFX production coordinates deliveries between vendors, editorial, and the DI. Send editorial previews to the VFX editor for review in context. Record every note and approval against the exact version.

Review proposed finals in a calibrated environment that represents the target display. Apply the correct CDL, Show LUT, or display transform and play at delivery resolution and frame rate.

If the shot is creatively approved, it will be deemed final and the VFX vendor is asked to deliver a final version to DI pending any technical fixes on their end.

On large shows, send selected WIPs to the DI early enough for grading feedback and contextual review. Each extra version adds I/O, editorial, and grading cost. On a limited budget, deliver one approved final instead of a stream of evolving versions.

After the initial conform, production editorial will follow up with change cuts and/or VFX inserts, providing edit lists of the changes and new shots. Editorial must specify the exact version number of a shot they are cutting in (using the offline QuickTime's clip name) so that the DI has clear instructions that cannot be misinterpreted. It is never appropriate to rely on the DI to guess which version is considered final, as the highest version of a shot is not always the version approved by the filmmakers.

Team Management

Distributed teams need self-sufficient generalists who can complete support tasks, follow the pipeline, and deliver clean work with limited supervision. Do not assume studio experience includes pipeline expertise: provide every artist with explicit color-management, naming, versioning, and delivery requirements.

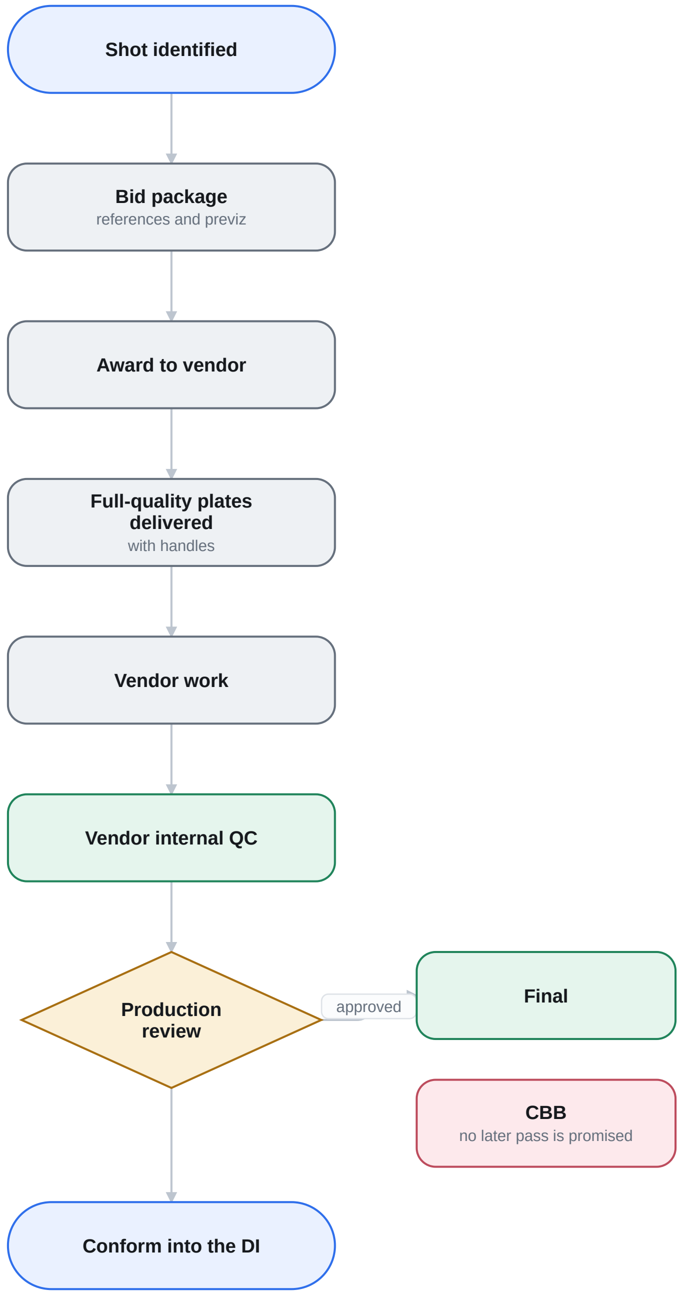

A shot from bid to conform. Bidding can run on references long before the shoot. Full-quality plates follow the award.

Work that arrives from outside the shot list

Not every request that reaches a vendor comes from the VFX schedule. **Trailer campaigns generate their own shot work**, including alternate frame ranges, censored or softened variants, replaced dialogue, and spoiler-free versions. Almost none of it is planned because the campaign develops quickly and mid-cut.

That work is usually funded as an overage against the marketing budget rather than post, so it is rarely a budget crisis. It is a **schedule** problem. Trailer shots are urgent by definition, artists get pulled off their assignments to service them, and the shots they were working on slip. Plan for the interruption rather than the shot list. See Trailer shots are VFX work nobody scheduled.

Editorial Turnover for VFX Plate Pulls

Plate-pull methods vary, but the production must use one documented workflow. The VFX producer and supervisor set the image pipeline, resolution, format, naming, and delivery rules for every artist and vendor.

The Bid Package

The bid package gives a vendor enough information to quote the work. Include plates or representative frames whenever possible so the bid reflects the material rather than a description. The VFX production team owns the package. On a small show, that may be the supervisor or editor.

A complete bid package, covering each shot or group of shots, includes:

- Storyboards or previz that show the intended staging and action.
- Reference photography: on-set stills, cleanplates, texture and lighting reference, and any witness-camera or survey material captured during production.
- Visual examples: temp composites (slap comps), frame grabs from other films, or concept art that communicate the target look.
- The relevant cut or edit, so the vendor can see each shot in the context of the scene and understand its timing and its neighbors.
- The plate or plates the shot will be built from, or at minimum representative frames at the working resolution.
- A detailed written description of the desired effect for every shot: what was shot practically, what needs to be created or removed, how long it runs, and what the shot has to accomplish in the story.

A precise package produces comparable bids and exposes scope before work begins. Ambiguity becomes contingency, underbidding, or change orders.

Timeline Preparation

In order to facilitate an accurate plate pull, edit lists must be derived from simplified timelines. Like a normal DI turnover, lists should not contain any nested sequences. Image effects, resizes, and repositions should all be removed.

Typically pull lists do not include handles. Rather the handles are specified to the vendor performing the pulls and added at that time. The pull lists should reflect the frames in the edit, not the handle frames.

Timewarps should be removed and the clips extended or shortened to precisely match the frames necessary to produce the timewarp.

It is common that separate pull orders are made for each visual effects sequence in a film as those sequences are soft-locked. This often happens out of narrative chronological order, and in smaller subdivisions than reels.

Unique VFX shot names are given to each clip, often designated by a timeline marker or other agreed-upon clip list to communicate that information to the operator performing the pulls.

Visual Effects Shot Naming

Name shots by sequence rather than scene when scenes may be reordered or intercut. Give every shot a sequence prefix and let VFX production define the logical divisions.

Shot numbering is relative to a specific sequence. Shots are often started at 0010 and increment in tens (0020, 0030…) in sequence order. This is done so that if a new shot (or multiple shots) is inserted in between pre-existing VFX shots, there are available numbers to assign to the new shot without breaking the sequence numbering logic.

For specialized effects that do not describe a scene, such as mobile phone message overlays or other forced narrative titles, it may be more convenient to label those as a single sequence name throughout the entire show. That way all of those similar graphics are logically named and grouped together, regardless of their underlying plate or scene content. This is valid whether the graphics are supplied prematted with an alpha channel or pre-composited into the scene.

Choose one naming convention before the first turnover and enforce it across vendors. Consistent names support reliable tracking, automation, and interchange.

File Naming Conventions

Use portable, case-consistent names. Replace spaces with underscores (`_`), and separate the base name, frame number, and extension with periods (`.`). Avoid parentheses and special characters.

VFX producers will provide vendors with documentation prescribing file naming conventions for their project. A common naming convention, used with subtle variations between projects and producers:

```
<SEQ><SHOT#4>_<CLASS>_v<VER#3>_<OPTIONAL>_<EYE>.#4.ext
```

Example: `ABC0130_v002.0012.dpx`, `ABC0150_bg1.0123.dpx`

Token	Meaning
<SEQ>	The sequence prefix is a 3-letter abbreviation of the sequence name.
<SHOT>	The shot number is a 4-digit zero-padded enumerator. Incremented in multiples of 10.
<CLASS>	Descriptor used to denote specific plate classifications.
v<VER#3>	The version number is a 3-digit padded enumerator preceded by a lowercase `v`. Shows running a per-vendor numbering policy use **four** digits instead, so the leading digit identifies the vendor: see Versioning below.
<OPTIONAL>	Additional descriptor.
<EYE>	Stereo 3D eye flag written as a full lowercase word.
`#4`	Frame number is a 4-digit zero-padded enumerator. Preceded and followed by dots (`.`) as divisors, not underscores (`_`).
`ext`	The file extension.

Plate classifications:

Class	Meaning
`bg#`	Background plate
`fg#`	Foreground plate
`comp`	Composite (render)

Optional descriptors:

Descriptor	Meaning
`matte#`	Matte (i.e. `_matte1`, `_matte2`, `_matte123`, `_matte456`)

Stereo 3D eye flags: `left`, `right`.

Example plate pull:

```
./plates/kxn0010_bg1/3072x1280/kxn0010_bg1.1001-1073.dpx
```

Example comp:

```
./renders/kxn0010_comp_v001/3072x1280/kxn0010_comp_v001.1000-1073.dpx
```

Frame Numbering and Handles

Handle frames (padding) exist so that if there is any last-minute need to roll an edit by a frame or two, the picture editor has that flexibility. **Plate handles and comp handles are separate decisions, and they are not the same number.** Pull the plate generously: eight frames each end is a common figure, and the cost of carrying them is disk, not labour. How many of those frames the vendor is asked to *comp* is the expensive question: over an entire feature, comping handle frames nobody uses adds up, so four each end is a common delivery. Ultimately the comp handle count is subject to production budget and schedule.

A common example of how visual effects plates might be pulled and numbered:

It is often mandated that the visual effects artists and vendors return a completed shot with a slate frame containing information about the shot, revisions made, the company name, artist's initials, and a framing chart. In this case, the slate frame is prepended to the file sequence without offsetting the frame numbering.

It is common to start the first frame of comp handle or the first main frame at 1001 and work backwards to the slate frame. This allows for a scenario in which the shot has been started but frames are added to the head during editorial and VFX must add new frames to their work. Starting at 1001 rather than 0001 allows for extension of head frames while preserving file numbering of previously existing frames. For example: even after a change in shot duration, frame 1034 will still be the same frame as in the initial pull.

Alternatively, it may be desired that the first main frame (first frame of the edit) is 1001 with handles working backwards.

In some situations, you may intentionally provide more handle frames in the plate than are required in the final comp. Handles are given at **both** ends: the point of a tail handle is that the cut can be rolled later, so a comp delivered flush to the last edit frame is unusable for the very purpose handles exist to serve.

Here the shot's main action runs 1001–1015. The plate carries eight frames of handle at each end, and the comp is delivered with four:

```
                head handles      main action      tail handles
Plate  8+8:   0993 … 1000       1001 … 1015      1016 … 1023
Comp   4+4:         0997 … 1000  1001 … 1015      1016 … 1019
```

Versioning

A costly mistake in visual effects shot management is losing track of versions or not having version numbers at all. Without version numbers, there is no reliable way to confirm that the intended shots are in the conform.[16] Version numbering practices are generally ubiquitous among artists and vendors, but occasionally there is confusion.

Any revision to a shot, even a simple correction in render settings, or a redelivery because of a corrupt frame, needs to have a new version number. Otherwise it is confusing for the production editorial team and conform editor to know which version should be cut in if they have received multiple deliveries of shots with the same version number.

Sometimes when working with multiple VFX vendors, a producer will institute a version numbering policy in which each vendor starts their version numbers differently. That way anyone involved can tell which vendor did the work just by looking at the file name.

For example Vendor A may start their version numbers at 1001, while Vendor B starts at 2001.

```
 Vendor A    kxn0010_comp_v1001
Vendor B    kxn0020_comp_v2001
```

Format Specification

Many cameras now acquire images in resolutions greater than 4K. What you finish at is increasingly decided by the distribution path rather than the budget. Theatrical DI at 2K remains spec-compliant and still happens, and archival masters are preserved at camera native resolution for future releases. But any project with a streaming distributor is now working to a UHD floor. Netflix requires VFX pulls *and* deliveries at a minimum active image area of UHD, and 20th Century Fox requires pulls at **native camera resolution**.

Production will need to decide on an imaging pipeline that suits their budgetary goals as well as the delivery requirements of their distributor(s). This decision has an impact on how visual effects plates will be pulled, as well as how visual effects artists and vendors will deliver shots back to production and the digital intermediate.

Data transfer time (a billable resource) is also a factor in what specifications your plates and shots fall into.

It is quite common for the camera acquisition format, visual effects plate format, and vendor delivery format to all differ.

For example: you may shoot 8K on a RED V-RAPTOR, pull plates at 6K or native, and receive UHD final VFX back from your artists and vendors. This provides the vendors with a high-resolution plate with enough detail for tracking and reframing, even though the ultimate deliverable is smaller. Note that the older pattern of pulling at 4K and delivering 2K comps is no longer a safe default for a streaming title.

Clearly defining these three critical formats is essential:

1. Acquisition format
2. VFX plate pull format
3. VFX vendor delivery format

Each of these formats is defined by:

- Raster dimensions (resolution)
- Scaling methods used to achieve that raster from the source media (scaled or cropped)
- File type (DPX, EXR, ProRes)
- Bit depth

Bit depth is not a free choice from a menu: the pipeline decides it. In the decade since this handbook's first edition the industry has moved decisively toward **16-bit EXR or DPX** and away from 10-bit, and that shift now reaches all the way back into scanning. Studio and streamer specifications have followed: 10-bit DPX is commonly barred outright for VFX-to-DI exchange, or permitted only where the original capture was itself 10-bit log, and log-encoded EXRs are routinely rejected.

The practical rule is that bit depth follows the encoding, not preference. Scene-linear material belongs in **half-float** EXR. Log material in 16-bit integer DPX or EXR. Reaching for 10-bit because a file is smaller is a false economy once anything is regraded, and it is the kind of decision that is expensive to reverse after a vendor has delivered a hundred shots. - Color space and transfer function (encoding)

Before any plate pulls or visual effects work begins, make sure these formats and requirements are clearly defined to all parties involved.

Support Files

In addition to the plates themselves, it is common to supply the Show LUT and ASC CDL values in sidecar XMLs (`.cc`) and/or LUTs, so they are easily accessible by whichever artists are delivered the plate(s). Alternatively, a Color Correction Collection XML (`.ccc`) "codebook" can be generated containing a lookup of all VFX shots and their corresponding CDL values. This is rarely distributed to individual artists, as it is more convenient to supply them with only the information they need for their specific shots (e.g. a `.cc` file per shot).

Lens distortion grids and camera metadata

Two things belong in the turnover that are easy to forget and impossible to recreate later.

Lens distortion grids. A photographed grid or checkerboard chart, shot on each lens at the apertures and focus distances actually used, lets the matchmove and compositing teams solve the lens's distortion rather than guess it. Without it, a 3D track fights the plate and the artist ends up hand-matching a distortion profile that a five-minute test would have given them exactly.

The time to shoot these is **camera prep**, where the lenses are already on the bench and nobody is waiting on them. On set, mid-shoot, they are expensive and frequently skipped. That matters most in the cases where the grids matter most: **vintage glass, uncommon or rehoused lenses, anything with custom tuning, and any lens introduced partway through a shoot.** A current cinema prime may already be in a vendor's database. A detuned vintage lens with a personality is exactly the one nobody can look up. If specialty lenses join the package after prep, grid them when they arrive rather than assuming the original session covered it.

Camera and lens metadata. Sensor mode and active area, lens make, focal length, T-stop, focus distance, filtration, camera height and tilt, and any in-camera reframing. Matchmove uses these to seed a solve rather than derive everything from the image. Most of this is recoverable from clip metadata *if* the pipeline preserved it, but transcodes routinely strip it, and once the unit has wrapped, nothing is going to reconstruct what filter was on the lens.

Both belong in the bid package as well as the turnover: a vendor bidding a heavy matchmove sequence prices the risk differently when they know grids exist.

Frames are commonly contained in a resolution subdirectory. This is common practice even when only a single plate format is used and is complementary to traditional digital intermediate workflows. It is particularly useful in situations where proxy resolutions are generated to accelerate compositing. Various resolutions can be contained within a single shot directory.

Example plate pull directory structure:

```
./support_files/kxn_RedLogFilm_to_Rec709_FF.cube

./kxn0010_bg1/3072x1280/kxn0010_bg1.1001-1073.dpx
./kxn0010_bg1/support_files/kxn0010_bg1.cc
./kxn0010_bg1/support_files/kxn0010_bg1.cube
```

Turnovers and Count Sheets

The plate pull produces the media. The turnover is how that media, and the information that goes with it, reaches the vendor. Each turnover packages the pulled plates, a locked reference video of the cut, the relevant reference and support files, and a count sheet that tells the vendor exactly what they are receiving.

The count sheet is the master list of every shot in the turnover. It tells the vendor about every shot being requested, what materials they will need, where to find them, and exactly what each shot involves. Per shot it typically carries:[17]

- **Shot code** (see Visual Effects Shot Naming above), and the **sequence, scene, and reel** it belongs to.
- **Master (record) timecode**, where the shot sits in the cut, and the **source timecode** of the plate, plus the frame range, duration, and the number of **handle frames** on head and tail.
- The **plate ID or IDs** that make up the shot and where to find them (path or drive), and a **thumbnail** so the shot is identifiable at a glance.
- A description of the shot and, more usefully, a description of the **effect** required: broken out per element when a shot combines several.
- Technical notes that change the scope of the work: **retimes / speed changes**, **resizes and repositions**, **frame extensions**, split-screens, and any **elements arriving from another vendor**.
- The **assigned vendor**, the current **version**, and the shot's **status** (bid, awarded, in progress, in review, final).

The count sheet travels with the plates so the vendor can reconcile what they were told against what actually landed on their drive, and it is the production's own record of what went out, to whom, and when. Turnovers are usually built per vendor, containing only the shots assigned to that vendor along with their plates, cleanplates, reference, and the Show LUT and CDL sidecars described under Support Files. Keeping each vendor's turnover self-contained avoids sending an artist material for shots they are not working on and keeps the delivery small enough to transmit efficiently.

In practice the count sheet is rarely a separate document at all: it is a view of the same tracking sheet the production is already using to manage the show, filtered down to the shots in that turnover.

Coordination and Tracking

None of this holds together without someone whose job is to keep it organized. On a large show that is the VFX coordinator, working under the VFX producer. On an independent production the role usually collapses into the VFX producer, the VFX editor, or whoever is acting as supervisor. Whatever the title, someone has to own the list of shots, the versions in play, the outstanding notes, and the deadlines, and keep all of it current as work comes and goes.

Full production-tracking platforms such as Autodesk Flow Production Tracking (formerly Shotgun, then ShotGrid) and ftrack are built for exactly this and do it well, with shot databases, review tools, and pipeline APIs. They are also heavy and complex, and the overhead of running one rarely pays off on a show with a few dozen shots and a handful of artists. Most independent productions run the whole thing on a Google Sheet instead: one row per shot, columns for the information that would otherwise live on the count sheet, and a status column that everyone reads from. It is low-tech, but it is shared, it is current, and it is the single source of truth the count sheets and turnovers are drawn from. Production tracking is covered in more depth in Visual Effects Production Management.

Organization is also leverage. Independent productions routinely ask vendors and individual artists for discounted or gifted work, and the ones who get it are the ones who make the work easy to say yes to. Every hour a vendor spends deciphering a disorganized turnover, chasing a missing plate, or reconciling contradictory notes is unpaid overhead, and it is exactly the menial work a favor is not meant to cover. A production that hands over clean, complete, well-labeled material and keeps its notes and versions straight is asking a vendor to donate their craft, not their patience, and is far more likely to get the discount.

The count sheet is a good example. It should be detailed but easy to read, well organized, and lead with the critical information. If it lists the source in and out timecodes of a plate but omits the frame count, every bidder has to do the arithmetic themselves to work out how many frames they are being asked to price. It is a small inconvenience, but a small inconvenience repeated across every shot and every vendor is exactly the kind of friction to remove when you are asking someone to work on your film for less than it is worth.

The Creative Review Cycle

While a shot is being built, the filmmakers review it in progress. These early and in-progress creative reviews are not about technical quality control. They are about content, design, and whether the shot is working. The vendor renders a QuickTime with the Show LUT, or display transform, baked in so the director, VFX supervisor, and editor see the shot in its intended look rather than as a flat log image. The editor can cut that QuickTime into the timeline and review it in the context of the surrounding scene.

Because the artists are usually remote, these reviews happen over the internet. Frame.io (now an Adobe product) and cineSync are the common platforms: the vendor uploads a version, the reviewers play it back and leave frame-accurate notes, and the vendor turns those notes into the next version. That loop repeats until the shot is approved.

The loop is where things fall apart without discipline. Notes get left in an email, a text message, and a review platform all at once. Two people refer to the same shot by two different names. A vendor delivers a new version against notes that have already been superseded. Every one of these is an organizational failure rather than a creative one, and every one of them costs a revision. The defenses are simple and non-negotiable: clear version numbers on every delivery (see Versioning above), one source of truth for notes so nobody is working from a stale list, and unambiguous shot naming so there is never a question of which shot is being discussed. This is what the coordinator and the tracking sheet exist to enforce.

Final Turnover into the DI

Creative approval is not the finish line. When a shot is creatively approved in review, the vendor turns it over in the agreed final format, typically EXR, at the delivery resolution, color space, and handle length settled on at the start of the show (see Format Specification above). That final render is conformed and cut into the DI against the offline reference, graded by the colorist in the context of the finished sequence, and put through quality control on the DI's calibrated pipeline.

That last step matters because a shot can look finished in a review platform and still fail in the DI. Graded up on a theater screen or a reference monitor, matte lines, grain mismatches, and black-point errors that were invisible in a compressed QuickTime become obvious. These are the issues covered under Visual Effects Quality Control Practices. A shot is only truly final once it has been conformed, graded, and QC'd in the DI and confirmed to work in the finished picture. Until then it is a creatively approved shot that has not yet proven itself in context.

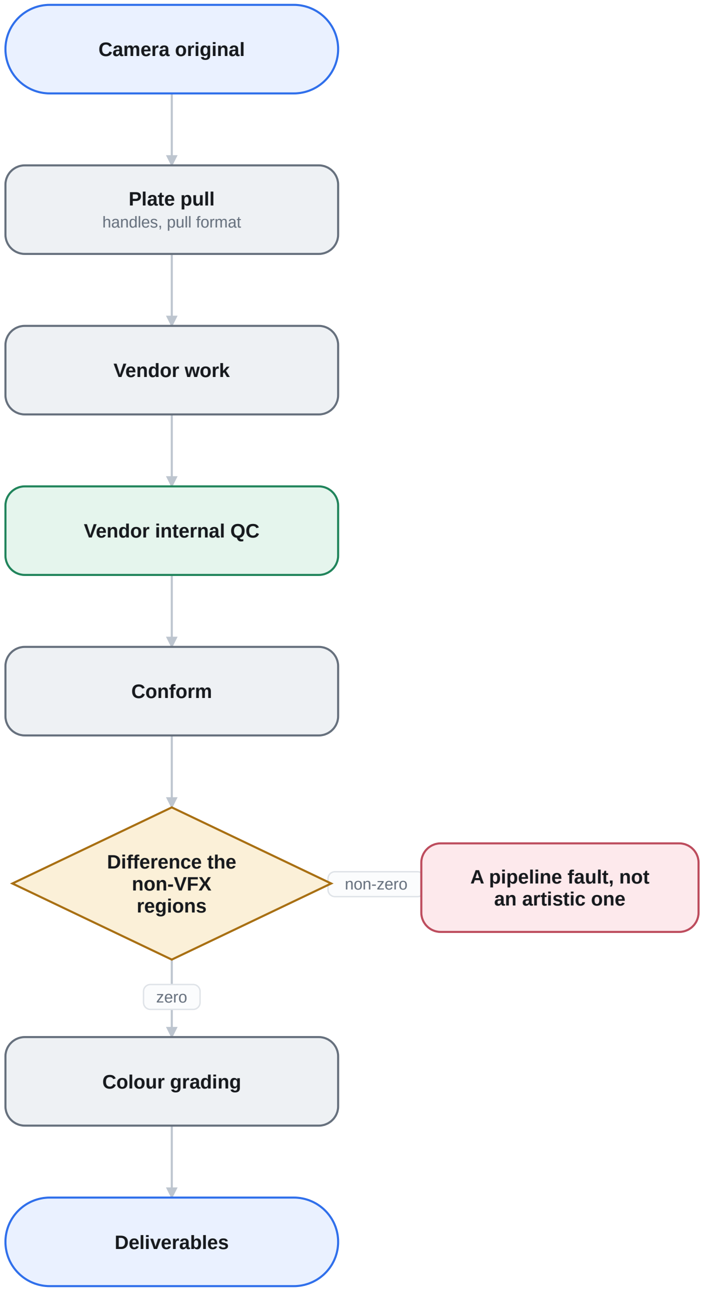

The plate round trip. Differencing the non-VFX regions is what proves the pipeline, not the artistry.

Visual Effects Quality Control Practices

!!! quote “One unfortunate reality in film production is that the typical artist display is not as high fidelity as the eventual theatrical viewing environment. Most often the theatrical audience will see more color detail than the original artist.”

```
Jeremy Selan, Sony Pictures Imageworks, *Cinematic Color VES*, 2012
```

The first rule of VFX is *do no harm*: the work must not reduce the plate’s technical quality. Define and test the imaging pipeline before production, then document the expected result for every vendor and artist.

The facility technical director normally owns this work. On smaller productions, assign it to the VFX supervisor or lead artist. Escalate pipeline problems as soon as they appear.

Image Processing Practices

Scaling Algorithms

Scaling remaps pixels. Filter choice controls interpolation, sharpness, aliasing, and ringing.

Impulse scaling, also called nearest-neighbor, is prone to producing noticeable aliasing and distortion, but is computationally inexpensive.

Better, smoother scaling algorithms result in softer edges with fewer scaling artifacts, but are more computationally intensive.

Many scaling algorithms introduce a degree of sharpening to produce apparent resolution gain when positively scaling an image. These are **negative-lobed** filters, including Rifman and Lanczos. The filter kernel has regions of negative weight. Those negative weights are what produce **ringing**, the halo of overshoot on the bright side of a high-contrast edge and undershoot on the dark side.

Negative filter lobes cause ringing. Negative pixels are the resulting undershoot. Clamping removes the undershoot, not the cause.

Scaling in **scene-linear is physically correct.** Resampling averages light, and light is linear. Averaging code values in a non-linear encoding is not the same operation and does not conserve energy. The difficulty is that linear encoding is also what makes the ringing so visible. A specular highlight sitting at 60.0 next to a shadow at 0.02 gives the filter an enormous local contrast to overshoot on, so the halo that would be mild in a display encoding becomes severe, and the undershoot swings well below zero.

There are several ways to manage this, and the right one depends on the shot:

- **Use a filter without negative lobes** where the softness is acceptable: no negative weights, no ringing.
- **Scale as close to 1:1 as possible, and resample as few times as possible.** Every resample is another opportunity. A chain of transforms that concatenates is one resample instead of four.

- **Round-trip through log around the resize.** This bounds the overshoot by compressing the range the filter sees. It is a genuine mitigation, but not a free one: filtering log values no longer averages light, so it trades physical correctness for artifact suppression. It suits alpha-free images at close to 1:1 best, and is least appropriate where energy accuracy matters.

Compositing applications do not manage this automatically. Nuke filters in the current working space. A log round trip must be explicit. Its Transform and Tracker `clamp` control discards negative intermediate values but leaves highlight overshoot intact.[18]

When a scene-linear resize produces severe ringing, test a log round trip around the resize. This is common for large camera-original-to-VFX-pull resizes with strong highlights. It trades energy accuracy for artifact control, so use it only when the ringing is more damaging.

The same principles apply to titles and graphics compositing in a digital intermediate.

Concatenated Image Transforms

Every resample can soften an image. Concatenate successive transforms so the application performs one net resample. Color corrections, filters, distortions, merges, and other non-concatenating nodes between transforms can break the chain.

In the example below, `Transform1` is scaling by a factor of 0.10 and `Transform2` by 10.0: a net-unity scaling transform, since the two factors multiply to 1.0. In the example on the left the transforms are concatenated before the `ColorCorrect` node, resulting in a color corrected image of the original, identity-transformed image. On the right, the `ColorCorrect` node bisects the transforms and breaks the concatenation. The result is an image with substantial loss of sharpness.

Colour correction is only the most common way to break a chain. **Any** node that cannot concatenate will do it: a filter, a blur, a distortion or lens-warp node, a merge against an element that was itself transformed. If sharpness is being lost across a chain that ought to be an identity, look for whatever sits between the transforms rather than assuming the transforms are at fault.

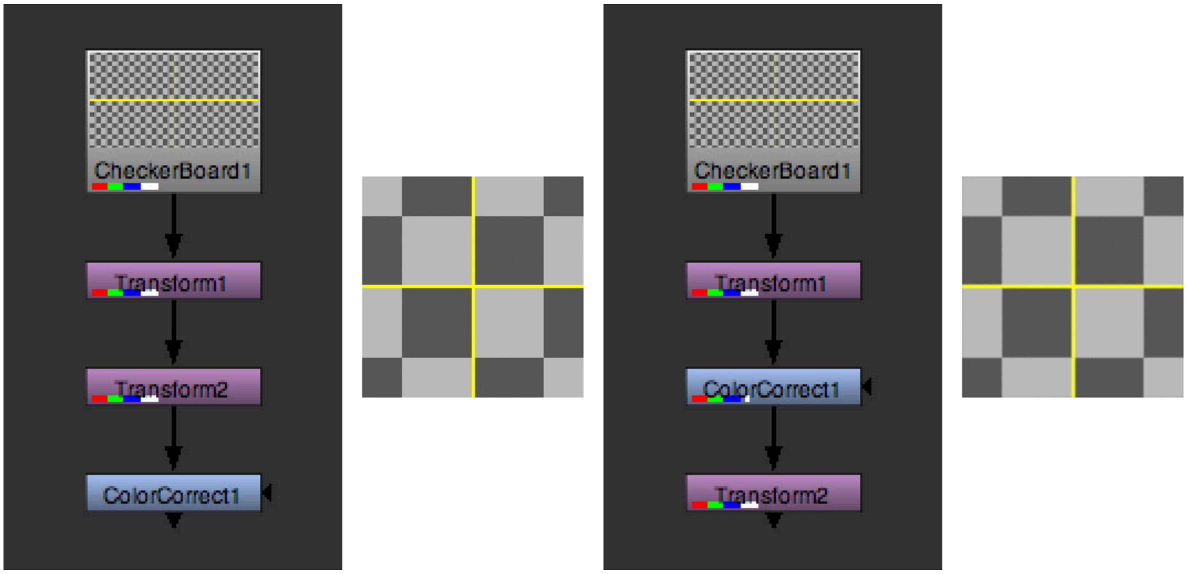

Concatenated versus broken transform chains

Figure 20: Left: both Transform1 and Transform2 occur before color correction. Right: Transform2 occurs after color correction, breaking concatenation.

Transfer Functions

It is common for compositing to be performed in a working space other than the original plate or delivery-specified encoding(s). Strict color management throughout the compositing process is necessary to ensure that the final delivery is in the desired color space and that no loss of dynamic range has occurred. Compositing is often best performed in a scene-linear working space where math operations function predictably and linearly, rather than in a logarithmic or gamma-referred working environment.

Compositing packages like Foundry's Nuke are designed around compositing in a linear working space by applying floating-point transfer functions to pre-convert material from camera-native logarithmic encoding to scene-linear and then back to a chosen destination encoding on render. This all works because the transfer functions are mathematically reversible, are not "s-shaped" display-referred transforms, and are performed in high precision 32-bit float.

Other compositing packages[19] coming from different pedigrees have different philosophies about desired working spaces and color management practices. Conducting end-to-end workflow tests in pre-production is paramount to identifying and resolving potential issues and deciding on a workflow that yields the best results.

Quality Control Practices

QC every shot before delivery. Check the complete frame for dropouts, matte edges, color shifts, formatting, naming, and frame-number errors. Compare against the plate and prior version, then review the sequence for continuity. Automated checks catch missing or malformed frames. Visual inspection catches compositing errors.

Many composited images look great when viewed in their target viewing space on the display they were composited on. However, those shots will ultimately undergo significant color grading and be viewed on consumer displays under less than ideal conditions.[20] Under those conditions, masks with unfeathered edges, differences in noise and grain patterns, and changes in shadow density and black points become significantly more apparent.

Gamma Checking

Raise and lower the viewport gamma aggressively to reveal matte lines and mismatches in noise, grain, or black level.

Gamma check revealing hard matte lines

Figure 21: Left: composited image under normal viewing conditions. Right: gamma increase reveals hard matte lines.

Exposure Checking

Additionally, linear increases in exposure can provide artists with a sense of the real dynamic range of their image. Do artificial (composited or CGI) highlights clip or read at the same level as comparable highlights in the original plate? Do artificial shadows match the black point (color and value) of comparable shadows in the plate? How do these values compare within the shot and to other shots nearby?

Exposure check revealing black point mismatch

Figure 22: Left: composited image under normal viewing conditions. Middle: exposure increase reveals black point mismatch in composite. Right: exposure decrease demonstrates highlight detail.

Difference Checking

Performing a high-gain mathematical difference operation between a visual effects render and an original plate can reveal a variety of compositing issues that may not be visible to the naked eye, including matte lines, noise/grain patterns, tracking issues, and highlight/shadow compression or clamping.

Difference check revealing film grain mismatch

Figure 23: Left: original plate. Middle: composite. Right: difference check reveals a mismatch in film grain between plate and comp.

Editorial and Finishing

Editorial Turnover for Digital Intermediates

Before grading, the conform editor rebuilds the locked timeline from camera originals and final VFX. A clean editorial turnover reduces relinking, interpretation, and billable conform time, regardless of the NLE.[21]

Timeline Preparation

Working in Reels

Studio features are still completed in reels averaging around 20 minutes each. The number of reels is determined by the total run time of the film. This practice dates back to negative assembly and film prints. A 2,000-foot projection reel holds roughly 22 minutes at 24 fps. It is a **projection reel**, the spool a print was mounted on for exhibition, not the length of a roll of raw print stock. That is the constraint that set the length, and it is why “reel” means about twenty minutes rather than some other number.

Use reels for theatrical features when picture, sound, DCP composition, or archival benefits from parallel work and shared reel boundaries. Use a continuous longplay for streaming delivery or a single-operator finish. Episodic projects normally remain one sequence. Confirm the structure with picture and sound vendors before turnover.

In the case of episodic television or web content, an individual episode is not typically subdivided and is instead treated as a single sequence.

Leaders and Tail-Pop

Start each reel on its numbered hour, for example, Reel 1 at `01:00:00:00`, and provide common picture and sound sync references.

The standard eight-second Academy leader starts each reel or sequence and counts down to a two-pop which will visually correspond with an audio pop in the temp and final mixes. The two-pop gets its name because it is followed by two seconds of black before the first frame of picture (FFOA, also written FFOP).

At the end of the reel or sequence, two seconds after the last frame of action (LFOA, also written LFOP) is a tail-pop. This is a way of visually confirming that the picture was intended to end there and did not end prematurely, and that the audio is indeed still in sync.

Do not use dialogue, foley, or effects as the primary sync reference. They may be intentionally or accidentally offset by one or more frames.

Figure 5: Example reel-based composition of a feature length film. Each reel begins on its own hour with an eight-second head leader (192 frames) and closes with a four-second tail leader (96 frames).

Figure 6: Example longplay composition of a feature length film: the same reels assembled into one continuous sequence, with no inter-reel leaders.

Figure 7: Example composition of a premium cable or OTT commercial-free television episode, with a 240-frame slate/leader and a 120-frame tail leader. Note that the first frame of program falls on the hour, at 01:00:00:00, and the head build runs before it: a program that starts after the hour is a common QC rejection.

Simplifying Timelines

After lock, flatten the timeline to the fewest tracks that preserve the edit. Remove redundant clips and tracks, and join through-edits.

Decompose all nested sequences and multicam clips. Many conform systems do not handle nested clip metadata correctly, and this can be cause for a conform editor to reject the turnover and request revisions from the production editorial team.

Creative resize and reposition effects should be completed using the native tools within the editing software whenever possible, as those can often be translated accurately across XMLs or AAFs, saving the conform editor time in recreating optical effects.

Creative timelines often contain overlapping layers. Remove any clip that is completely obscured by a higher track.

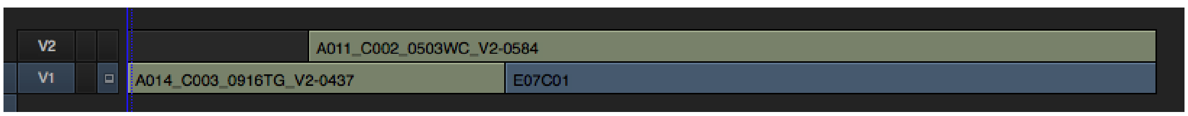

A multi-layer composition

Figure 8: A multi-layer composition.

An edit list may still request obscured clips, forcing the DI to restore, copy, or transcode media that never appears on screen.

Organize Timeline Layers by Type

Depending on the preferences of the conform editor and the particular characteristics of the project, the way you organize the timeline layers in the turnover can vary. It is good to collaborate with the conform editor and come to an agreement on an organizational strategy that benefits you both the most.

A common, simple example for timeline layout is as follows:

Track	Contents
V5	Titles and graphics
V4	Visual effects provided by external vendors
V3	Additional optical plates if necessary (B-sides for split comps)
V2	Opticals (to be completed by the conform editor)
V1	Scans (Original Camera Negative [ONEG], drama, unaltered camera footage)

If multiple plates are necessary to complete a complex optical, and it is intended to be handled in conform editorial rather than with an external VFX vendor, layering those necessary plates on top of each other makes sense. In this example you would have two or more layers dedicated to opticals before the visual effects and title layers.

Edit List Generation

Ask the conform editor which interchange format and export settings they require.

Regardless of the edit list flavor you choose, an important universal requirement is that the clips in your project and timeline correlate to the camera original media through a tape name and timecode relationship. In the case of major digital cinema cameras, this is relatively straightforward and has been planned for by the manufacturers. Each take is named using a unique identifier based on its camera roll, clip number, and sometimes the date and specific camera serial number it was recorded with. The specifics vary between manufacturers, but nevertheless, these systems were built with post-production in mind. The tape name and timecode can often be derived from the source file's metadata as well.

Whichever source you use for the tape name metadata, it is important to make sure that information is available to the conform editor in the original media they are given so that they can match to those files quickly and accurately.

In cases where there is no reliable tape name, such as GoPro or specialty cameras not designed with offline/online post-production workflows in mind, it is important that the DIT or editorial team produce a unique reference that will translate to the conform. This can be done by manually producing unique file names for each clip during DIT ingest and dailies processing. This is particularly important in instances where camera timecode is regenerated at the start of each clip (e.g. each clip starts at 00:00:00:00). It will be difficult to conform a sequence if a timeline clip matches dozens of source clips, all without unique tape names and all starting with the same timecode.

It is also critical to maintain a common time base between footage and editing timelines. For example, mixing 24.000 fps footage with 23.976 fps footage on a 23.976 fps timeline can produce timecode errors that increase the complexity of the conform, and sound that drifts progressively against picture. The majority of American independent features are shot and edited at 23.976 fps, but that is a convention rather than a rule: see Frame Rate: 23.976 or 24.000 for the decision and the sync arithmetic behind it.

EDL

CMX-style EDLs, or Edit Decision Lists, are the oldest and simplest edit list, dating back to tape-to-tape linear online editing systems. They are incredibly universal and all digital intermediate systems can accept them. They are, however, limited in a number of ways. In practice EDLs carry only static timewarps, where speed does not change over time, though the limit is tooling rather than the format: SMPTE 258M does define a syntax for variable-rate moves, but essentially nothing in a modern post chain writes or reads it, so treat variable speed ramps as unavailable in an EDL and hand the DI the ramp some other way. EDLs cannot describe resize/reposition effects, flips, or flops. Support for multi-layered EDLs varies from platform to platform, but in general they are not recommended. Typically, each layer in the offline edit must be exported as an individual EDL file.

The advantages of EDLs are that they are universally accepted and easily human readable. This makes procedural corrections and adjustments to troublesome conforms rather easy in most cases. EDLs can also contain ASC CDL (Color Decision List) data in convenient EDL comments, both `*ASC_SOP` and `*ASC_SAT`, so unlike FCPXML an EDL carries the full ten values.

Reel names and the 8-character limit

A strict CMX 3600 EDL allows only **8 characters** for the reel (tape) name, and this is the single most common source of tape-name mismatch in a conform. When a source name is longer, the EDL writer truncates it, and if several clips share the first eight characters, they collapse into identical reel names that no longer identify anything. Most EDL writers add the full name back as a `* FROM CLIP NAME:` comment, so the information is usually recoverable, but the machine-readable field is not the one you want to be relying on.

In practice almost every modern EDL parser accepts far longer reel names, and most of the time long names round-trip without complaint. The risk is not a lack of support. It is that the one tool in the chain that *doesn't* is often the one you discover last. Agree the reel-naming scheme with the DI up front, keep source names unique within their first eight characters where you reasonably can, and test a round trip rather than assuming.

A strict CMX 3600 EDL is also capped at **999 events**, which a feature reel can exceed on a cut-heavy sequence. If it does, split the list.

XML

In the context of digital intermediate finishing, "XML" typically refers to the Final Cut Pro 7 XML Interchange Format, adhering to the interchange specifications defined by Apple. These are the standard XMLs produced by both FCP 7 and Adobe Premiere. XMLs support a variety of image effects, including, but not limited to, resize and reposition. They also support speed ramps, although these may be interpolated differently by the DI system and require finessing by the conform editor. XMLs support multi-layer timelines.

XMLs are human readable text, but are not always easy to edit quickly.

AAF

An open interchange specification published by the **AMWA** (Advanced Media Workflow Association, formerly the AAF Association), primarily utilized by Avid for picture and audio turnovers, and also used by other editing systems for audio turnover to Pro Tools. It is frequently mistaken for a SMPTE standard. **MXF** is the SMPTE one (ST 377-1), and its structural metadata derives from a subset of AAF's class model. AAFs contain much of the same metadata as an XML. In practice, AAFs out of current Avid Media Composer carry baked keyframe data for speed ramps, which is particularly helpful in producing accurate time ramps in the conform process without interpolation discrepancies between keyframes. Avid writes these as private parameters that no published specification documents, so the behaviour is something you confirm by testing your own round trip rather than by reading a spec: Resolve reconstructs them per-frame on import. AAFs support multi-layer timelines. AAFs are stored as binary data and are not human readable or editable.[22]

FCPXML

A newer flavor of XML used by Final Cut Pro. Compatible with Blackmagic DaVinci Resolve, but not with all other conform systems. It improves on FCP 7 XML in most respects and supports multi-layer timelines.

Do not rely on it to carry colour, though. FCPXML's ASC CDL support is weaker than it sounds: it carries **slope, offset and power but never saturation**: nine of the CDL's ten numbers, and since FCPXML 1.5 the values are not a real element at all, only an XML *comment* inside `filter-video`. Apple's own documentation for the version where it *was* a proper element noted that it "is not processed during import." Nothing has changed through FCPXML 1.14 (Final Cut Pro 12). If colour decisions have to travel with the cut, send an **EDL, ALE, CC or CCC** alongside it. Those are the formats CDL round-trips in practice, and they carry saturation.

OTIO

OpenTimelineIO (OTIO) is an open-source interchange format and API for editorial timelines, originated at Pixar and now an Academy Software Foundation (ASWF) project. Unlike the formats above, each tied to a particular NLE lineage, OTIO is designed as a **vendor-neutral interchange and a programmable data model**: a JSON-based timeline (`.otio`) plus a set of adapters that read and write EDL, FCP7 XML, AAF, FCPXML, and more, so it can act as a hub that translates between them.

Adoption is strongest on the **pipeline and engineering** side rather than the artist-facing one. It is widely used inside studio and VFX pipelines (it ships with or is supported by Nuke Studio/Hiero, RV, Autodesk Flow/ShotGrid, and many in-house tools) and has broad library support, but native **save-as-OTIO from mainstream NLEs is still limited**. DaVinci Resolve has added OTIO import/export, while Avid Media Composer and Premiere still lean on their traditional formats. In practice, for an independent conform in 2026 you will still most often hand the DI an EDL, XML, or AAF. OTIO's real value today is as the **programmatic glue** for translating, comparing, and validating those lists, and its role is growing.

Choosing a format: what each can carry

Every one of these formats carries the same non-negotiable core: the **clip list with a source tape name and timecode**, which is what a conform cannot do without. Beyond that shared core they diverge:

Capability	EDL	FCP7 XML	AAF	FCPXML	OTIO
Human-readable	✓ text	✓ verbose	✗ binary	✓ verbose	✓ JSON
Tape name + timecode	✓	✓	✓	✓	✓
Multi-layer timelines	✗ one/file	✓	✓	✓	✓
Resize / reposition	✗	✓	✓	✓	✓
Variable speed ramps	✗ in practice[3]	✓ interpolated	✓ baked keyframes[1]	✓	✓
ASC CDL color	✓ SOP + SAT, in comments	✗	~	~ SOP only, as a comment[2]	✓
Origin / lineage	CMX tape online	FCP7 · Premiere	Avid · Pro Tools	Final Cut Pro X · Resolve	ASWF (Pixar), cross-NLE
DI acceptance	✓ universal	✓ common	✓ common	~ Resolve+	~ via adapters

[1] Baked speed-ramp keyframes in AAFs from recent Avid Media Composer versions.

[2] Slope, offset and power only: no saturation, and carried as an XML comment rather than a real element since FCPXML 1.5. Unchanged through FCPXML 1.14 (Final Cut Pro 12).

[3] SMPTE 258M defines a variable-rate syntax, but tooling support is effectively absent.

In short, EDL is the lowest common denominator. It supports one layer and no effects, but is universal and editable. XML and FCPXML add layers and effects. AAF adds baked speed ramps (at the cost of being binary). And OTIO spans all of them as a translating hub rather than a native turnover format.

!!! tip “List-management tooling” Wrangling, comparing, and repairing edit lists by hand is error-prone. Video Village’s DI tools include **EDL and list-management utilities** for exactly this: inspecting, cleaning, and reconciling lists before they reach the conform.

Conform Checks

Critical to the conform process is a visual comparison between the offline cut and the final conform. This is how the conform editor visually verifies that the conform is accurate. During the “conform check” or “confidence check” process, they can identify any incongruities between the conform and the intended edit, including temporal discrepancies and missing opticals or transitions that may not have translated from the edit list correctly.

The offline reference is a same-as-source QuickTime export from the offline editing application, typically Avid DNx LB or Apple ProRes 422 Proxy / LT. (DNxHD 36 remains in circulation on HD-raster shows and in older studio specs, but its compression IDs were reclassified as the legacy “HD Profile” in SMPTE ST 2019-1:2016 and cannot carry a UHD raster.) Occasionally H.264s are used. However, they are discouraged because MPEG motion artifacts can produce visual incongruities that would be flagged as potential temporal discrepancies, making it difficult to use as an accurate reference.

Offline QuickTime Reference

The provided offline QuickTime reference is essential in the initial conform process. As the edit changes, and few edits are truly locked, production editorial will provide additional check QuickTimes to production editorial to the DI to verify picture changes.

At some point during the DI, either following a significant milestone or cut change, it is advisable for the conform editor or colorist to render a DNxHD or ProRes QuickTime of the edit, with or without temp color grading, optionally including source clip name and timecode burn-in, and send it to the production editorial team. They will then perform their own visual inspection and verify that the DI conform matches their offline edit as expected.

It is also useful to export an EDL of the DI timeline for the production editorial team to verify.

Opticals List

Many basic effects are easily reproduced during the conform process. These effects are usually resize, reposition, stabilization, flips, flops, simple timewarps, dissolves, effect transitions, and locked-off split-screen composites. These effects are commonly referred to as “opticals” because they required the use of optical printing in the days of negative assembly conforms. Depending on the complexity and volume of shots, they may be delegated to a dedicated VFX artist or vendor.

It is important to identify these shots for the conform editor and provide notes on what the desired effect is and whether you want them to precisely match your offline reference or make improvements upon it.

These lists can be produced and communicated a number of different ways, including timeline markers, locators, or other comments in an edit list, or even a simple spreadsheet listing event timecode and descriptions.

Media Consolidation and Delivery

Major studio features utilize LTO tape archives containing their camera original media. These are usually managed and stored in their digital intermediate vendor's vault during production and throughout the DI. This facilitates easy clip pulls for dailies transcoding, visual effects, and ultimately the onlining and conform of the final edit for color grading. This is common as most production companies do not have the infrastructure to store and pull shots as quickly or efficiently as large facilities.

However, independent productions typically assume responsibility for the storage and archival of their production assets: usually on redundant external hard drives. Even small productions shooting with multiple cameras and recording Camera RAW often span multiple archival drives by the end of production. When it comes time to turn media over to a digital intermediate facility or independent colorist, it is important to consolidate the original media necessary for the conform to as few drives as possible. Trimming media with handles, or simply isolating only the necessary source clips, can massively reduce the amount of data I/O time required to ingest the media at the facility. These are tasks that can sometimes be completed by assistant editors at production editorial. Delivering disorganized media spanning across multiple drives will only increase the amount of time a facility or independent operator needs to conform the project and ultimately the cost associated with it.

Delivering VFX to DI

It is common for editorial to follow the initial DI conform turnover with subsequent change notes and visual effects updates. Those subsequent list turnovers may or may not include a new offline reference QuickTime if there are no timing changes and the visual effects shots have a usable temporal sync reference. If timing can still be visually confirmed using the original offline reference, a new one is not necessary unless requested.

Edit lists delivered containing visual effects shots must exactly match the complete VFX shot name, including version as delivered from the VFX vendor, in order to facilitate an accurate conform. It is important that the timecode of the VFX clips in the offline NLE matches the timecode of the files presented to the DI, using header timecode, or more commonly the frame number, as agreed upon by the DI. It is a common problem that QuickTimes from the VFX vendor may not reflect the accurate timecode, or the NLE interprets it incorrectly. This can be troublesome for a DI and should be tested early on before many shots are delivered.

Common Issues and Solutions

It is not uncommon for details to be lost in translation during a conform. Specific effects produced in the offline NLE often cannot translate across XMLs or AAFs and rarely translate across EDLs at all. Additionally, tape names or clip names may be close matches to the original media, but may require a little tailoring by the conform editor to make it work with their individual workflow. It is recommended to provide EDLs in addition to XMLs or AAFs so that there is a redundant supply of metadata in case something doesn't work correctly with the XMLs or AAFs. This reduces the amount of downtime and troubleshooting between the DI and production editorial.

Supplying the conform editor with a copy of the NLE project containing the turnover timeline can be of help if the conform editor has access to and is familiar with the same NLE production editorial uses. This project does not necessarily need to include media files and is helpful as a reference in cases where metadata in the EDLs, XMLs, or AAFs is not linking back to original media correctly.

An Overview of Digital Intermediates

The digital intermediate (DI) is the stage where picture is conformed, graded, mastered, and turned into deliverables. It brings camera originals, VFX, titles, and graphics into one managed pipeline without photochemical generation loss.

Where the name comes from, and whether it still fits

“Intermediate” comes from the photochemical elements between the camera negative and release prints: the interpositive and internegative. A digital intermediate replaced that middle with a scan, digital color timing, and a film-out. In short: **film in, digital in the middle, film out.**

Most shows now originate, finish, and exhibit digitally, but the name remains useful. “Color grading” describes only part of the work. The DI is the convergence point for every acquisition format, resolution, color space, bit depth, VFX shot, title, master, and deliverable.

Responsibilities of the Colorist

The colorist develops the final look with the filmmakers and maintains visual continuity across changes in camera, lighting, location, and VFX. Judge the work on the intended reference display: a calibrated theatrical projector for cinema or a calibrated monitor for home distribution.

Involve the colorist during camera, lighting, wardrobe, and look tests. Early tests expose technical limits, inform camera and workflow choices, and establish a look that production, editorial, VFX, and finishing can reproduce.

Camera RAW

The DI often processes camera originals natively. RAW formats preserve sensor data and defer white balance, exposure interpretation, and other color decisions until post instead of baking them into acquisition.

Common examples of such cameras recording in a RAW format include:

Camera family	RAW format
ARRI ALEXA 35 (also Mini LF, ALEXA 265)	ARRIRAW
Sony VENICE 2 (also BURANO)	X-OCN (16-bit)
RED V-RAPTOR / V-RAPTOR [X] (also KOMODO-X)	R3D (REDCODE RAW)
Blackmagic URSA Cine family	Blackmagic RAW (BRAW)
Canon EOS C400 / C500 Mark II	Cinema RAW Light

The ARRI ALEXA 35 records ARRIRAW with **LogC4 / ARRI Wide Gamut 4** (ARRI REVEAL color science). LogC3 / AWG3 remains correct for the ALEXA Mini LF and earlier bodies. ARRI added **ARRICORE** with the ALEXA 35 Xtreme in 2025. It is an **RGB rather than Bayer** codec at roughly half ARRIRAW’s data rate that still leaves sensitivity, white balance and tint unbaked, so it keeps the parts of RAW that matter downstream while costing a great deal less storage. Sony’s VENICE 2 and BURANO record 16-bit **X-OCN**. RED, acquired by

Nikon in 2024, records **R3D (REDCODE RAW)** across the V-RAPTOR line. Blackmagic dropped CinemaDNG for **Blackmagic RAW (BRAW)** in 2018, now recorded across the **URSA Cine** family. Canon's current Cinema EOS bodies record **Cinema RAW Light**.[23]

RAW is common on features and VFX-heavy shows, but not required. Many episodic and documentary productions use compressed I-frame camera originals. Choose the recording format around image requirements, storage, turnaround, and the tested post pipeline.

RAW flexibility also creates room for error. Record the source gamut, transfer function, decode settings, and any creative transform. Never make VFX pulls through a display-referred tone map. The lost dynamic range cannot be recovered later. Test the decode and round trip before production, and involve the colorist or DI facility when the pipeline is uncertain.

See Scene-Referred and Display-Referred Imagery and Look-Up Tables for the concepts behind this workflow.

Working with Visual Effects Mattes

Provide VFX mattes when the colorist may need to isolate CG elements, foregrounds, or backgrounds. Reusing production mattes avoids new rotoscoping, keying, and tracking in the DI.

Deliver mattes at the composite's format and resolution in normalized linear encoding.[24] Pack up to three mattes into RGB, or embed additional channels in the OpenEXR composite to reduce conform work.

RGB composite

Figure 2: RGB composite.

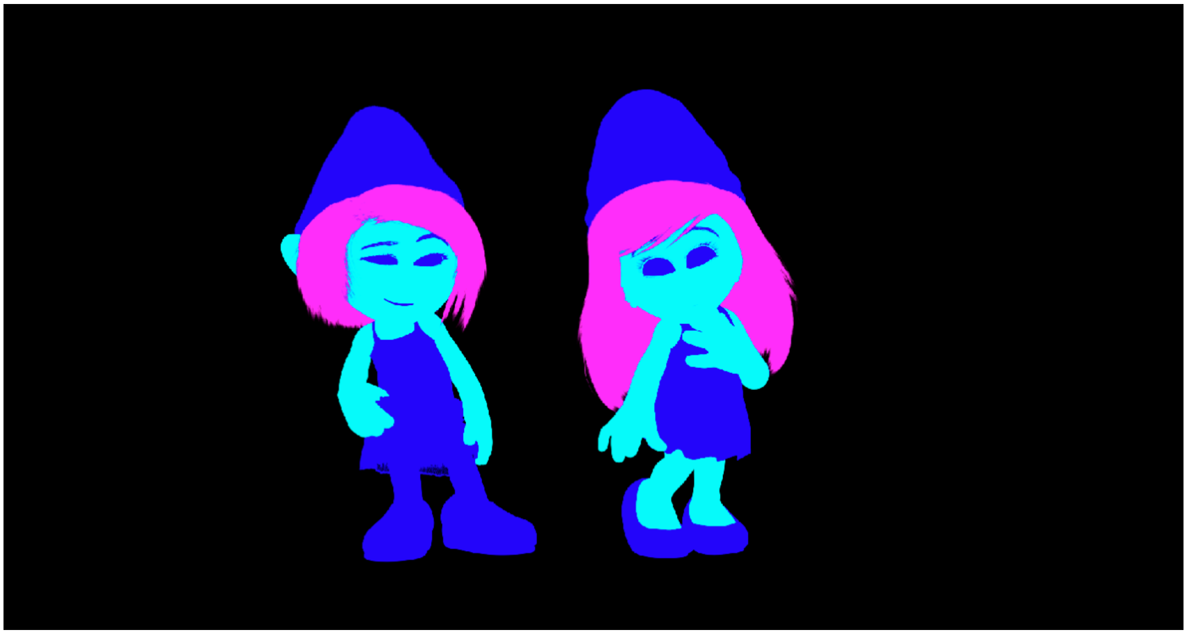

Three-channel DI matte

Figure 3: Three-channel DI matte.

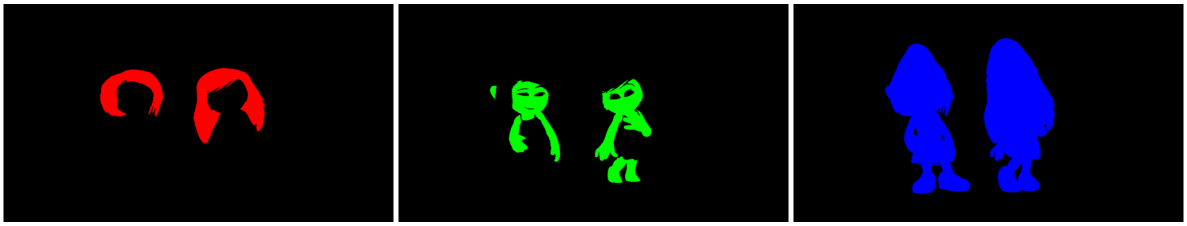

Three-channel DI matte with separated RGB channels

Figure 4: Three-channel DI matte (separated RGB channels).

Color Management

Color Management and OpenColorIO

ACES defines the transforms. **OpenColorIO** (OCIO) distributes and applies them consistently across applications. A show can specify ACES correctly and still produce mismatched renders if one artist uses a different config.

What OCIO is

OpenColorIO is an open-source color management library, originally developed at Sony Pictures Imageworks and now an Academy Software Foundation project. It is embedded in Nuke, Maya, Houdini, Katana, Blender, RV, and most other tools in a VFX pipeline.

The unit of configuration is an **OCIO config**: a file (plus, historically, associated LUTs) that enumerates the color spaces a show uses, the transforms between them, and the named roles that applications look up:

- **Color spaces**: `ACEScg`, `ACEScct`, `ARRI LogC4 / AWG4`, `sRGB - Display`, and so on.
- **Roles**: abstract names that decouple applications from specific spaces. `scene_linear`, `compositing_log`, `color_picking`, `data`. An application asks for `scene_linear`. The config decides what that means on this show.
- **Displays and views**: the viewing transforms available in an artist's viewer.

Roles allow a working space to change without editing every artist's scripts. Redefine `scene_linear` once in the config instead of changing it in each Nuke script. ### How ACES is actually delivered

Facilities consume ACES as an OCIO config, not by implementing the CTL reference transforms directly.

The Academy publishes reference transforms in CTL. Artists receive an OCIO implementation, usually a smaller `cg-config` for CGI/compositing or a broader `studio-config`. See the OCIO ACES config project.

"The show is ACES" therefore requires every application to use the same OCIO config and version.

What to specify on a show

Alongside the three format specifications, a color-managed show needs these written down:

Item	Example	Why
Config identity and version	`studio-config-v4.0.0_aces-v2.0_ocio-v2.5.ocio`	Two artists on different config versions can produce different renders from identical scripts, and ACES 2.0's output transforms differ substantively from 1.x, so a 1.3 config and a 2.0 config are not two versions of the same render
Working space	`ACEScg`	See the AP0/AP1 warning
Grading space	`ACEScct`	ACEScc and ACEScct are not interchangeable
Plate color space	`ARRI LogC4 / AWG4`	Must match what the plate pull actually produced
Delivery color space	`ACES2065-1`, uncompressed EXR	The DI's requirement, not the vendor's preference
Viewing transform	Output Transform, P3-D65 or Rec.709	Artists must review under the same view the DI uses

Distribute the config with the plates in `support_files/`, alongside the Show LUT and CDLs. Do not ask vendors to source it independently.

Versions

OCIO 2.x replaces baked LUT approximations with native ACES transforms and matching GPU and CPU processing paths. Use the annual VFX Reference Platform to align ACES, OCIO, and OpenEXR versions across vendors. CY2026 specifies ACES 2.0, OCIO 2.5.x, and OpenEXR 3.4.x. ### Checking that it works

Run these checks in a confidence package before shot work:

1. **Round-trip a plate.** Convert to the working space and back without operations, then difference against the original. Allow only floating-point precision error.
2. **Round-trip an extreme.** Repeat with a plate containing genuine highlight clipping and deep shadow. Faults that are invisible on a mid-tone chart show up here as clamping.
3. **Compare viewers.** The same frame, under the same view transform, in Nuke and in the DI system. If they differ, resolve it before anyone grades anything.

Use a difference check for the first two tests.

The Academy Color Encoding System (ACES)

ACES standardizes a scene-referred interchange encoding, working spaces, and input and output transforms. Multiple vendors can implement the same pipeline without negotiating a custom color workflow.

What ACES actually is

ACES is not a look, and it is not a LUT. It is four things:

1. **A scene-referred interchange encoding**: ACES2065-1, standardized as SMPTE ST 2065-1.
2. **Working spaces** derived from it for compositing and grading.
3. **Defined transforms** at the input and output of the pipeline.
4. **A container** for the image data, standardized as SMPTE ST 2065-4.

The ACES pipeline as a left-to-right image path: camera → (Input Transform) → ACES2065-1 interchange → ACEScct grade → graded ACES2065-1 master → (Output Transform) → each deliverable. VFX vendors are a tangent off the interchange: plates go out and renders come back in ACEScg, round-tripping through ACES2065-1, not through the grade.

Everything converts to ACES2065-1. VFX round-trips through the interchange in ACEScg. Grading converts internally to ACEScct and back to ACES2065-1. That graded ACES2065-1 is the Graded Archival Master, and the Output Transform derives each deliverable from it. The ungraded interchange is the non-graded assembly master.[25]

Color primary sets

ACES defines two sets of primaries, and confusing them is the most common source of ACES-related pipeline errors.[26]

Set	Role	Chromaticities (x, y)
AP0	The ST 2065-1 encoding primaries. Encompasses the entire CIE 1931 visible spectrum. The primaries are non-physical.	R (0.73470, 0.26530) · G (0.00000, 1.00000) · B (0.00010, −0.07700)
AP1	Working-space primaries, closer to achievable display primaries. Basis of ACEScg, ACEScc, and ACEScct.	R (0.713, 0.293) · G (0.165, 0.830) · B (0.128, 0.044)

AP0 uses non-physical primaries so the interchange encoding can represent the visible spectrum without clipping. That makes it suitable for archival interchange but unsuitable for image processing. AP1 exists for rendering, compositing, and grading.

!!! warning “The most common ACES mistake” Rendering or compositing in **AP0** because it is “the ACES space.” Lighting and shading math performed against non-physical primaries produces hue shifts that look like a lighting bug and get chased as one. Composite in **ACEScg** (AP1). Reserve ACES2065-1 (AP0) for interchange, delivery between vendors, and archival. #### Why ACES Is “D60” in a D65 World

All ACES encodings share a white point of **x = 0.32168, y = 0.33767**.

The ACES white point is D60-like but not CIE D60. It is an encoding neutral, not a required display or creative white. Grade to the calibrated display white. The Output Transform handles adaptation. See Academy TB-2018-001 §§3–4 for the derivation.

What it means in practice

- **You do not grade to D60.** Set your grading display to its calibrated white point (P3-D65 or DCI for theatrical, D65 for home video) and grade normally.
- **The Output Transform handles the adaptation.** Converting ACES to a D65 output applies a chromatic adaptation. That is expected behavior, not an error.
- **When neutrals drift, look here first.** If grays read warm or cool between the grading display and a deliverable, a mismatched or doubled white-point adaptation in the transform chain is the usual cause, more often than a monitor calibration fault.
- **Do not "correct" the ACES white point.** Productions occasionally try to force ACES to D65 by inserting an adaptation before the Output Transform. This double-adapts and produces exactly the neutral drift it was meant to fix.
- **Use D60-sim only when D60 rendering is intentional.** It preserves ACES white on a D65 display without adding an extra adaptation before the Output Transform.

Encodings

Encoding	Primaries	Transfer	Use
ACES2065-1	AP0	Linear	Interchange between vendors and archival. The "full fidelity" exchange format.
ACEScg	AP1	Linear	CGI rendering and compositing.
ACEScc	AP1	Logarithmic	Color grading, where controls expect a log relationship to scene exposure.
ACEScct	AP1	Logarithmic with a toe	As ACEScc, with a film-like toe near black that makes lift/offset behave the way colorists expect.
ACESproxy	AP1	Log, integer	On-set look management and transmission over HD-SDI. **Not** for storage, production imagery, or final grading.

ACES2065-1 values are 16-bit floating point, with 1 sign bit, 5 exponent bits, and 10 mantissa bits. They encode relative exposure values linearly.[27]

!!! note "ACEScc vs ACEScct" Use **ACEScct** unless you have a specific reason not to. ACEScc is purely logarithmic to black, which makes a lift adjustment behave very differently from the log grading controls most colorists have muscle memory for. ACEScct adds the toe and is the usual choice for grading. Pick one and specify it in the workflow document. Mixing them across vendors produces subtle mismatches in the shadows that are painful to diagnose late. #### Transforms

ACES 1.0 deliberately renamed the transforms in user-facing terms. TB-2014-012 records both the old and new names, and explicitly deprecates "RRT" in end-user documentation:

Current name	Formerly	Does
Input Transform	Input Device Transform (IDT)	Converts camera-native data to ACES2065-1
Look Transform	Look Modification Transform (LMT)	Applies a global, show-wide look upstream of the Output Transform
Output Transform	"RRT plus ODT"	Converts ACES data to display code values

The Look Transform is a show-wide creative transformation analogous to a Show LUT. See Academy TB-2014-010 and Look Transforms.

Grade under the Output Transform. Do not bake it into the render. Keep the Graded Archival Master scene-referred and apply the display transform last.

Container

ACES image data is carried in the ACES Image Container, SMPTE ST 2065-4. It is OpenEXR-compatible but **feature-restricted**: a constrained profile of EXR rather than arbitrary EXR. It stores 16-bit half-float pixels and mandates a fixed set of header attributes. ST 2065-4:2023 lists ten required attributes, one of which is `compression`.[28]

!!! warning "ST 2065-4 containers are uncompressed" SMPTE ST 2065-4 applies only to **ACES2065-1 (AP0, linear)** interchange and archival files. It requires uncompressed 16-bit half-float EXR. ACEScg plates, renders, and working files are ordinary OpenEXR files and may use normal ZIP, PIZ, or DWA compression.

```
| File | Typical encoding | ST 2065-4 container? | Compression |
| --- | --- | --- | --- |
| Archival / source master | ACES2065-1 (AP0) | Yes | **Uncompressed** (mandated) |
| Full-fidelity vendor interchange | ACES2065-1 (AP0) | Yes | **Uncompressed** (mandated) |
| VFX plate pull | ACEScg or camera-log EXR | No | Compress freely (ZIP/PIZ/DWA) |
| VFX render back to the DI | ACEScg (AP1) EXR | No | Compress freely |
| Comp working files | ACEScg (AP1) EXR | No | Compress freely |

A compressed ACES2065-1 EXR may open normally but is not a conformant ST 2065-4 container.
```

Clip-level metadata travels in an **ACES Metadata File** (AMF specification). The earlier ACESclip form is described in Academy TB-2014-009. LUTs travel in the **Academy-ASC Common LUT Format** (CLF specification, S-2014-006).

The **Reference Gamut Compression** (RGC specification) added in ACES 1.3 handles out-of-gamut camera values. Use it when saturated practical lights or lasers produce artifacts outside AP1.

ACES 2.0

!!! note "Pin your ACES version" ACES 2.0 is now the version named by the VFX Reference Platform for CY2025 and CY2026, and it ships in DaVinci Resolve 21 and Baselight v7, but plenty of in-flight shows and older OCIO configs are still pinned to 1.x. Confirm which version each vendor is actually running before committing a show, and verify specifics against the ACES documentation.

ACES 2.0 replaces the separate RRT and ODT with one Output Transform built on the **Hellwig 2022 JMh** color appearance model. It adds a unified tone scale, chroma compression, volumetric gamut compression, and improved invertibility.[29]

Version history: 1.0 (2014), 1.1 (2018), 1.2 (2020), 1.3 (2021), 2.0 (2024–25). Many existing projects and configs remain pinned to 1.3.

!!! warning "Do not switch versions mid-show" ACES 2.0 output does not match 1.x. Finish on the version used to start the show. Whether a full color appearance model belongs inside a display rendering transform is an open argument in the field, with ACES 2.0 on one side and approaches like FilmLight's TCAM on the other. It is not a settled question.

Should an independent production use ACES?

Yes when:

- You are working with **multiple camera systems and multiple VFX vendors**. ACES gives everyone a specified pipeline to implement, so plates and renders interchange without a bespoke negotiation per vendor.
- You **lack the color-science infrastructure or resources to produce and vet a custom workflow**. A standardized, documented pipeline you can adopt is safer than a homemade one you cannot fully test.

Maybe not when:

- You have a **single camera acquisition format**. Much of what ACES buys you is interchange across disparate sources. With one source, a well-run camera-native workflow may be simpler.
- You have a **strong creative reason to deviate from the ACES Output Transform**. If the look you want fights the standard rendering, you are working against the system rather than with it.
- You have a **fixed, well-understood set of deliverables and the preparation to accommodate them**. If nothing about the job needs neutral interchange or archival, the setup cost may not pay for itself.

Adopting ACES badly is worse than a well-run camera-native pipeline. It costs setup, testing, and vendor discipline.

Choose the ACES version in pre-production, record it in the format specifications, and round-trip a confidence package before shot work.

Pitfalls

- **Compositing in AP0.** Use ACEScg instead.

- **Two Output Transforms.** Applying a display transform in the comp and again in the DI creates the classic double-tone-map. Renders delivered to the DI should be scene-referred, with the Output Transform applied by the colorist.
- **An Input Transform that does not match the camera settings.** Input Transforms are specific to camera, and often to color science version and recording mode. “ALEXA” is not a specification. “ALEXA 35, LogC4, AWG4” is.
- **Assuming ACES removes the need for a look.** ACES gives you a neutral rendering, not a look. Productions that skip look development and grade against the bare Output Transform tend to discover the film has no visual point of view somewhere around the DI.
- **Treating “we’re ACES” as a workflow document.** It specifies neither resolution, nor bit depth, nor handle length, nor naming. Those still need to be written down.

Application-Native Color Management: Resolve and FilmLight

Blackmagic **DaVinci Resolve** and FilmLight **Baselight** include scene-referred color-management frameworks. Each assigns source color spaces, maps clips into one working space, and applies a display transform at output. Unlike ACES, these systems live inside a vendor product.

These frameworks automatically apply a defined chain of input, working-space, and output transforms. A colorist can instead build that chain manually with **color space transforms (CSTs)**, **LUTs**, **DCTLs**, and a chosen display rendering transform. Manual construction offers more control but requires the colorist to maintain consistency.

DaVinci Wide Gamut is documented, and FilmLight publishes portable color-space files. Neither is an open standard like ACES (ST 2065-1) or has ACES's ecosystem-wide support through OpenColorIO. They are **published and partially portable, but single-vendor**.

Each system works in its own wide-gamut internal space. The three you will meet are ACES's **AP1** (the ACEScg working gamut), **DaVinci Wide Gamut**, and FilmLight's **E-Gamut** (what T-CAM renders from). All are deliberately wide, encompassing most camera and display gamuts, with virtual primaries that fall outside the spectral locus:

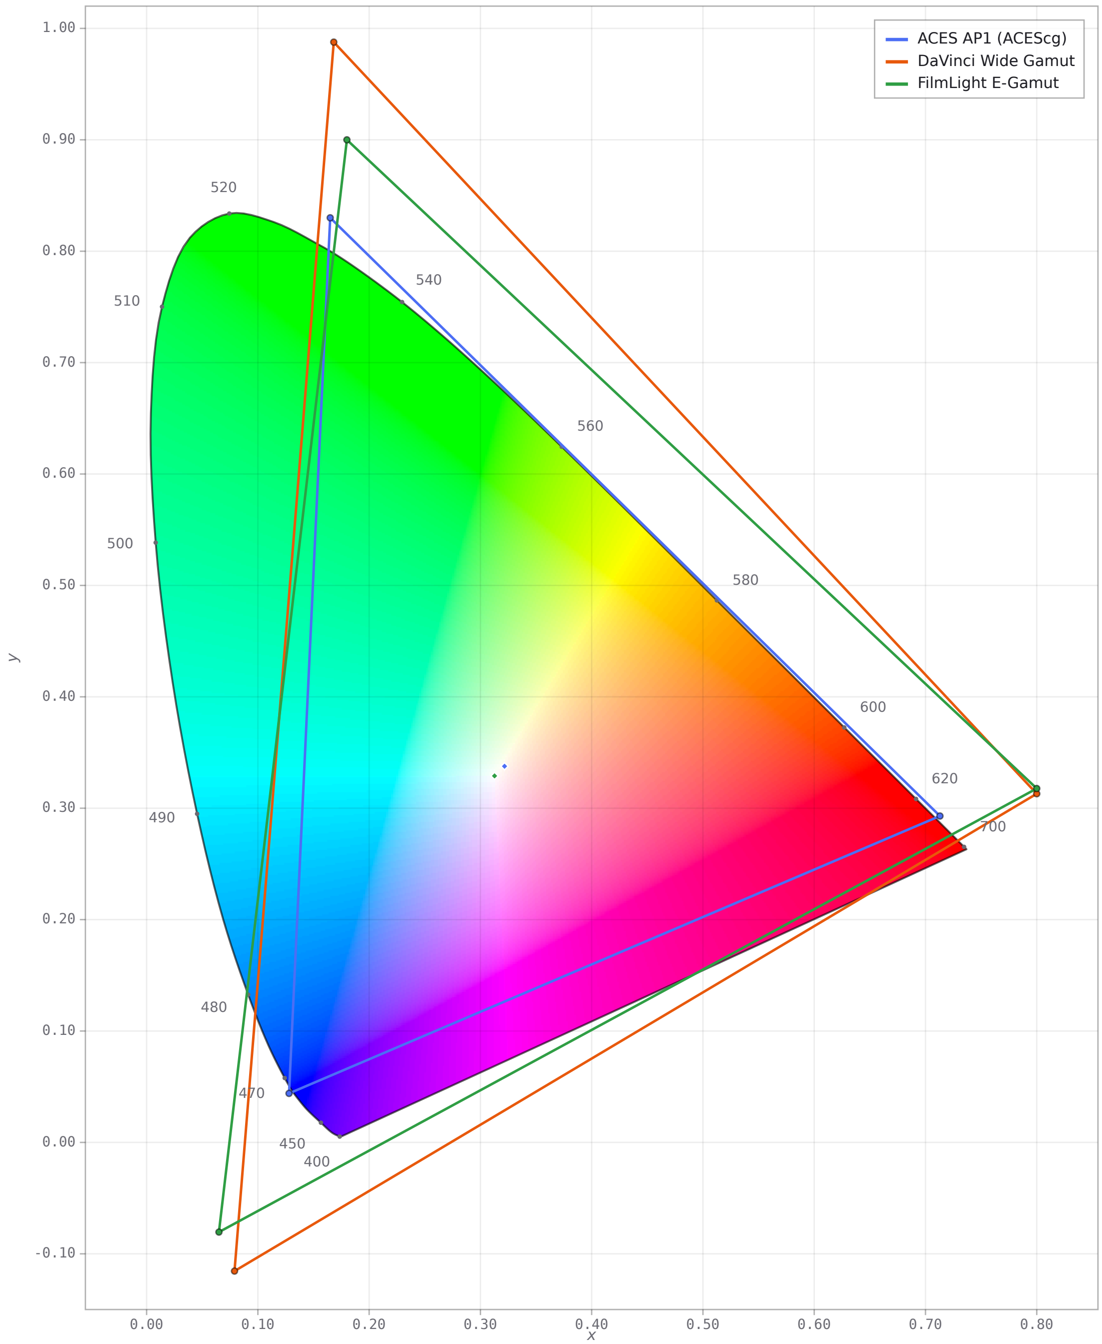

ACES AP1, DaVinci Wide Gamut, and FilmLight E-Gamut on the CIE 1931 diagram

The three managed working color spaces on the CIE 1931 diagram: ACES AP1 (ACEScg), DaVinci Wide Gamut, and FilmLight E-Gamut, the gamut T-CAM renders from. DaVinci Wide Gamut is the widest. All three reach beyond the spectral locus, but not to the same degree: DWG and E-Gamut do so dramatically, with negative blue-y and greens close to y = 1, while AP1 does so only slightly. Its blue primary sits at y = +0.044 and its green at y = 0.830. Interactive: drag to pan, scroll to zoom. Static version.

DaVinci Resolve Color Management (RCM)

Objective: normalize multiple cameras into one scene-referred working space and retarget deliverables inside Resolve. Assign each clip's camera profile. Resolve applies the corresponding log and gamut transform.[30]

RCM inside one application: each source clip is assigned an input color space, mapped into the DaVinci Wide Gamut / Intermediate working space, then rendered through one output display transform (the Resolve DRT, applied last) to each deliverable: Rec.709, P3-D65, and Rec.2100.

RCM and ACES are both automatic, scene-referred systems. RCM is a configurable, single-vendor framework with selectable working spaces and output tone-mapping methods. Resolve can also run ACES directly, so grading in Resolve does not preclude an ACES pipeline.[31]

Blackmagic recommends **DaVinci Wide Gamut (DWG)** with the **DaVinci Intermediate** log curve.[32] Its primaries, white point, matrices, and log equations are published and can be reproduced in OCIO. Resolve's tone and gamut mapping remain internal. **The encoding is specified. The rendering is not.**

FilmLight: Baselight, Daylight, and the Truelight Color Spaces

Objective: manage a grade in **E-Gamut / T-Log** and render through the **T-CAM v2** display transform. Baselight and Daylight implement this as the Truelight Color Spaces system. These are FilmLight product names, not standards. See FilmLight Truelight.

The same shape as RCM, inside Baselight / Daylight: each source clip is assigned a Truelight input, mapped into the E-Gamut / T-Log working space, then rendered to each display through the T-CAM v2 display transform (applied last).

T-CAM uses a lighter appearance model plus viewing-condition correction. ACES 2.0 uses a Hellwig-derived color appearance model. The difference affects display rendering and should be evaluated with the colorist.[33]

FilmLight publishes `.flspace` and `.fltransform` files, a Truelight OCIO config, and an Autodesk Flame color policy. Baselight can also run a full ACES pipeline. These files improve portability but remain vendor-controlled rather than independently implemented standards.[34]

Application-native color management vs. ACES: easier and harder

Setup effort depends on whether the pipeline crosses applications and vendors:

	Application-native (RCM / TCS)	ACES
Assigning input color spaces (camera normalization)	The same metadata-driven menu	The same metadata-driven menu: **virtually identical labor**
Single-app, single-facility finish	On by a project setting	Also a project setting. Runs in one app just as well
Multi-vendor VFX interchange	**More work**: no path to VFX vendors. The colorist hand-authors LUTs or transforms and workflow guidance for each vendor	**Less work**: VFX tools already ship OCIO/ACES configs. Vendors are already set up
Neutral archival master	No: a working space, not an interchange encoding	Yes: ACES2065-1

- Camera normalization requires similar labor in either system.
- ACES and OpenColorIO already ship in standard VFX applications through the VFX Reference Platform. Extending RCM or TCS across vendors requires the colorist to supply matching LUTs, transforms, and documentation.

Use an application-native system for a single-facility finish. Prefer ACES when the pipeline crosses vendors or requires a neutral archival master.

Why choose application-native color management over ACES?

Choose an application-native system when:

- **The finish stays in one facility or application.** For a documentary, commercial, or Resolve-only grade, ACES's interchange advantage may go unused.[35]
- **Grading-tool behavior and working-space preference.** Colorists often choose a timeline space because the grading controls *behave* the way they like in it: you can set RCM's timeline space to DaVinci Wide Gamut or any log format for exactly this reason.[36]
- **You prefer its display rendering.** The colorist may favor T-CAM or Resolve's DRT over the ACES Output Transform.[37]
- **Only part of the pipeline needs ACES.** Resolve and Baselight can use native working spaces while retaining ACES where interchange requires it.

!!! note "Using a custom display transform" Either system can manage inputs and the working space while a custom LUT, DCTL, or DRT handles output. In ACES, this removes the guarantee that standard transforms alone can reconstruct the delivered look. Archive the custom transform with the master in a durable, documented form. ### Documentary vs. single-camera narrative

Camera count and vendor count are separate considerations:

- **Documentary:** many source formats but usually one finishing facility. Application-native color management is often the lower-friction choice.
- **VFX-heavy narrative:** fewer source formats but multiple vendors, applications, and deliverables. ACES usually provides better interchange and archival support.

Many cameras do not inherently favor ACES. Many vendors, applications, deliverables, or an archival requirement do.

When to decide

Choose the **working space and display transform** in pre-production. Record them in the format specifications and prove the pipeline with a confidence package before shot work begins. See Choosing a Color Management Approach.

HDR Mastering

HDR is not a look or an export setting. It encodes brighter highlights and ties code values to measurable display luminance.

Three things change relative to an SDR master:

- **The transfer function.** This is the curve that translates the numbers stored in a file into light on a screen. SDR and HDR use different curves, which is why an HDR file played as though it were SDR does not simply look dim. It looks wrong.
- **The container is wider.** HDR encodes into a colour space large enough to hold more saturated colour than Rec.709 can describe. That is headroom, not an instruction to fill it.
- **Metadata travels with the picture**, describing the display the grade was made on, so a television with different capabilities can adapt the image sensibly instead of guessing.

All three must be correct for predictable playback outside the grading room.

The standard: ITU-R BT.2100

BT.2100 is the recommendation that defines HDR television. Its parameters:[38]

Parameter	**Value**
Primaries	R (0.708, 0.292) · G (0.170, 0.797) · B (0.131, 0.046): the BT.2020 primaries
White point	D65
Bit depth	10 or 12 bit
Transfer functions	**PQ** (Perceptual Quantization) or **HLG** (Hybrid Log-Gamma)
Reference display peak	≥ 1,000 cd/m²: for *small-area highlights*, not full-screen white

BT.2100 defines a **container gamut**, not a requirement to grade to the BT.2020 primaries, and no display on the market reaches those primaries anyway. See BT.2020 is a container below.

PQ and HLG are not interchangeable

PQ (SMPTE ST 2084) is absolute: each code value maps to a luminance up to 10,000 cd/m². Use it for cinematic mastering and streaming.

HLG is relative and designed for live production and broadcast, with partial SDR compatibility.

For scripted streaming content, use **PQ** unless the delivery specification requires HLG.

Metadata

An HDR master is incomplete without metadata describing the display it was made on. This is a genuine change from SDR practice, where the mastering display was implied by the standard.

Standard	Carries	Kind
SMPTE ST 2086	Mastering display color volume: primaries, white point, min/max luminance	Static
MaxCLL / MaxFALL	Maximum content light level. Maximum frame-average light level	Static
SMPTE ST 2094	Dynamic metadata for color volume transform, per-scene or per-frame	Dynamic

ST 2086 describes the mastering display. MaxCLL and MaxFALL describe the content. Record both accurately. Incorrect values cause downstream tone-mapping errors that may not appear on the reference monitor.

Distribution formats

Format	Metadata	Notes
HDR10	Static (ST 2086 + MaxCLL/MaxFALL)	The baseline. PQ, 10-bit, Rec.2020 container. Royalty-free.
HDR10+	Dynamic (ST 2094-40)	Scene-by-scene metadata without a proprietary grading step.
Dolby Vision	Dynamic (ST 2094-10)	Adds an authored trim pass. Effectively mandatory for premium streaming delivery.

Netflix does not accept HDR10-only masters for originals. Plan its HDR delivery as Dolby Vision, with L1 analysis metadata covering every shot.

Dolby Vision changes the schedule. Grade the HDR master, then have the colorist author **trims** for lower-capability displays and SDR. This is a creative pass, not an export setting.[39]

The metadata comes in levels, and the three you will hear named are **L1**, **L2** and **L8**.

L1 is generated analysis metadata. One L1Min/L1Mid/L1Max set describes each shot. It is not editable. Shot events must cover the timeline without gaps or overlaps. Dissolves, transitions, and long-play analysis may use per-frame events.

Build accurate shot boundaries from an EDL where possible. Boundaries that miss the picture cut can produce luminance shifts or flashing and commonly fail Dolby Vision QC.

!!! tip “The shot that swings from dark to bright” A single continuous shot that moves from a dark interior to a bright exterior is the classic problem case: one L1 set has to describe both ends, so the analysis lands between them and neither maps well. You cannot fix it by keyframing L1. Those values are generated, not editable.

```
 **Cut at the transition midpoint and add a dissolve across the cut.** Each half receives
separate
analysis while the dissolve creates smooth per-frame metadata.
```

L2 and L8 are per-shot, per-target creative trims. L2 belongs to CM v2.9. L8 belongs to CM v4.0 and adds more controls. Confirm the required content-mapping version before booking the trim pass. Changing versions requires another creative pass.

!!! warning “Budget the trim pass” Budget the Dolby Vision trim and SDR review as creative work, not metadata export.

One master or two?

Choose before finishing whether SDR will be **graded separately** or **derived** from HDR trim metadata.

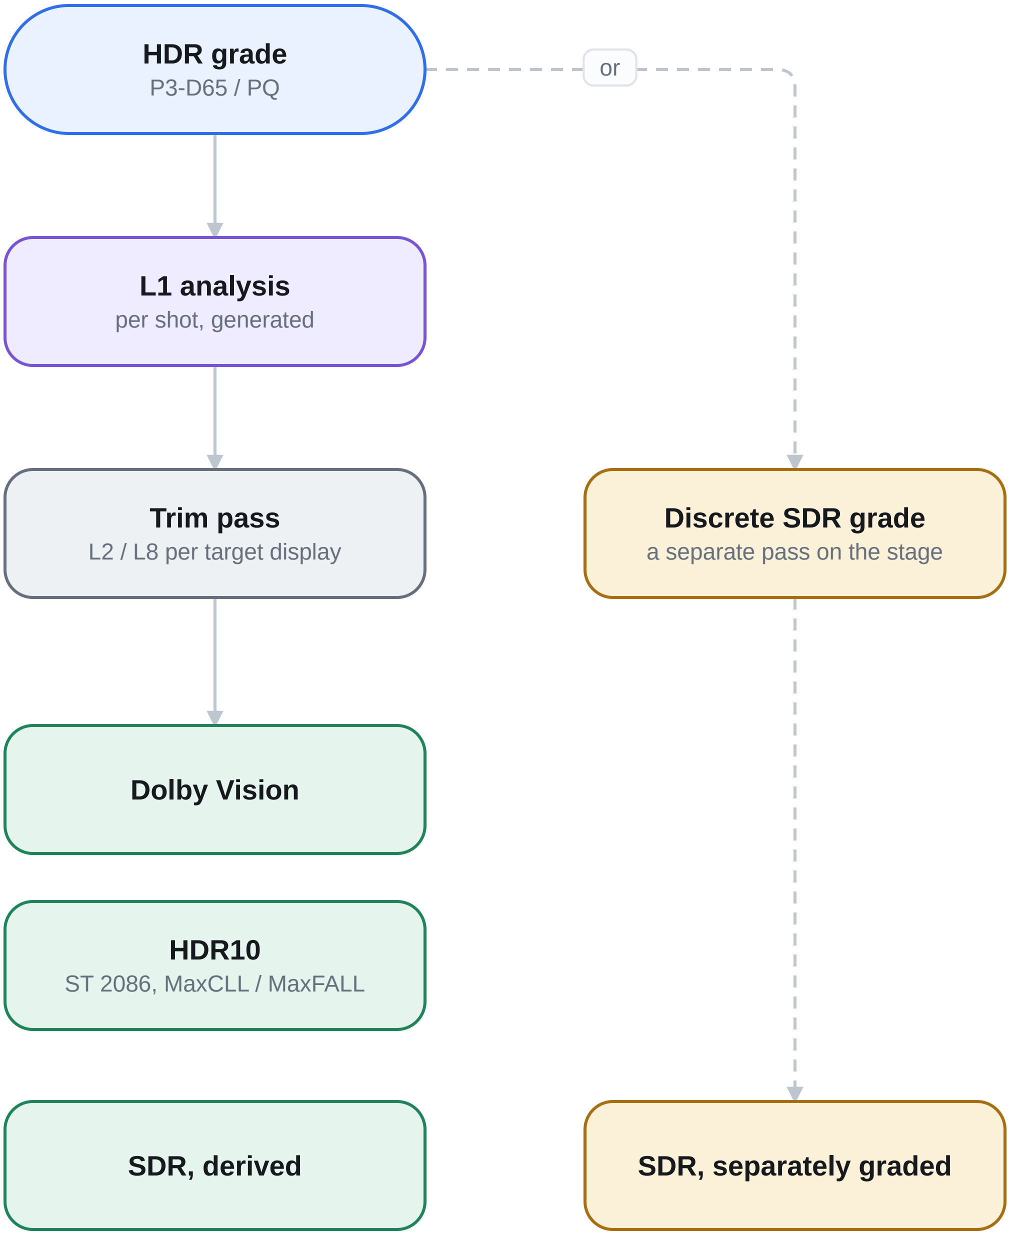

One master or two. The derived route runs through the trim metadata. A discrete SDR grade is a separate pass on the stage.

Requirements vary:

- **Netflix** takes a single Dolby Vision master plus its sidecar XML and derives the SDR streams from the metadata. There is no separately graded SDR in the chain.
- **Sky** requires the opposite: a separate Rec.709 grade where an HDR deliverable is commissioned, and states that it will not convert HDR to SDR itself.
- **Disney** permits either, but requires that derived-versus-discrete be settled with mastering early in finishing.

- **Dolby** requires the colourist to check the SDR Rec.709 100-nit mapping across the whole timeline whichever route you take, because that mapping is what SDR displays fall back to.[40]

Confirm whether “SDR master” means a derived file or a separate grade. The difference may be several days of stage time.

L2 and L8 trims are image-global. They provide no windows, keyers, or spatial secondaries. If the look depends on regional corrections, budget a separate SDR grade.

Netflix requires L1 but not L2 or L8. Complete a trim pass anyway to verify SDR rendering.

Mastering in practice

Grade in P3-D65. That is what almost every HDR master is actually made in, at a stated peak luminance: commonly 1,000 or 4,000 cd/m². You grade inside the volume your mastering display can genuinely measure and reproduce, not inside the container the file is encoded in.

The container is a separate question, and the delivery specification names it. Rec.2100/BT.2020 is the common case, but **Netflix asks for a P3-D65 container**, so read the schedule rather than assuming the wider one.

BT.2020 is a container

BT.2020 uses monochromatic primaries on the spectral locus.[41] Real displays emit broader bands. Only laser systems approach full coverage. Treat BT.2020 as encoding headroom, not a grading target. Grade within the measured volume of the mastering display, usually P3-D65.

A typical deliverable set for a streaming feature:

- **HDR master**: PQ, P3-D65 inside Rec.2100, 1,000 nit, with ST 2086 + MaxCLL/MaxFALL
- **Dolby Vision metadata**: trims authored against that master
- **SDR master**: Rec.709 / BT.1886, derived from the HDR grade with a colorist pass

Theatrical

Theatrical HDR is a separate presentation, not the home master, and it arrives two ways: brighter projection, and self-emissive LED walls.

Projected HDR includes Dolby Vision Cinema and HDR by Barco. Both exceed the 48 cd/m² DCI reference and require their own delivery specifications.

- **Dolby Vision Cinema** uses dual-laser projection and dynamic per-scene metadata.
- **HDR by Barco** uses Lightsteering to redirect light toward highlights.[42]

Direct-view LED displays replace projection with a self-emissive wall. DCI-certified HDR targets about **300 nits** peak across the display. Examples include Samsung Onyx, Cinity LED, and LG Miraclass.[43]

Either way, a theatrical HDR pass is its own grade, and its own line item.

Consequences for visual effects

Plan HDR requirements with VFX before shot work:

- **Bit depth.** The argument only strengthens for HDR: 10-bit log is not enough. 12-bit log is the floor for HDR finishing, 16-bit float preferable. See Camera Log.
- **Highlights are visible now.** Specular detail, practical lights, and skies that clipped invisibly in a Rec.709 grade are plainly visible at 1,000 nits. Comps that “worked” in SDR reveal blown highlight rolloff, clipped CG speculars, and mismatched black points.
- **Review at the target.** Use an HDR reference display rather than assuming Rec.709 review will translate.
- **Use exposure checks.** Raise and lower exposure to compare artificial highlight clipping with the plate.

Mastering and Delivery

Distribution Deliverables for Independent Film

A distributor's *delivery schedule* lists the picture masters, audio stems, caption files, artwork, and legal documents required before release and payment. Treat it as its own production schedule.

Exact requirements usually arrive with the distribution agreement. They vary by buyer, change over time, and may be negotiable. Building a complete package before receiving the schedule risks paying for the same work twice.

!!! tip "Finish archival-ready. Produce specifics on demand" Get the **foundational, expensive-to-recreate** elements right during finishing: a clean high-quality archival master, textless elements, and a complete audio stem set. **Defer the format-specific deliverables until a buyer hands you real specs.** Done this way, whatever a distributor later asks for can be produced quickly and cheaply from materials you already have, without an expensive archival restore or a re-mix.

What to lock now, and what to defer

Lock now: while the stage is open	Defer: until the buyer sends a spec
Master mezzanine render: what every other deliverable is made from	ProRes flavour, raster and frame-rate build
Archival master, graded *and* ungraded	Loudness-normalised near-field mix
Textless elements for every shot carrying text	Caption and subtitle formats, and language versions
Mix session, discrete stems, fully-filled M&E	Head and tail builds, slates, distributor bumpers
Conform project and LTO archive	Platform-specific encodes

Items on the left require the stage, vendors, or source elements to recreate. Items on the right vary by buyer and can be generated from those retained assets.

Master format: ProRes, not IMF

Independent distributors overwhelmingly ask for a self-contained ProRes QuickTime (`.mov`), not an IMF package. ProRes is the picture master across VOD, TVOD, and home-video distribution. The distributor makes its own platform-specific encodes downstream. IMF is a studio-and-streamer interchange format. Produce it only when a studio, major streamer, broadcaster, or large catalog aggregator specifically asks for it.

ProRes 422 HQ is the common default. Some distributors require 4444 or 4444 XQ. Retain a high-quality archival master so any requested flavor can be generated without re-finishing. See Common Working Formats.

The picture master baseline

A few things hold across almost any HD picture master and are safe to assume before you have a spec:

- **Raster and rate.** 1920×1080, progressive, at the film's native frame rate, usually **23.98 fps**. Deliver at the original aspect ratio, letterboxed to *fill* the 1920×1080 raster (mattes baked in, not a cropped 1920×858 file).
- **No up-conversion.** A 4K master is meaningful only for a film actually *finished* in 4K. Do not up-convert HD to 4K or SD to HD.
- **Color. Resolution and dynamic range are separate decisions: do not assume 4K means HDR.** A 4K SDR master is ordinary and is still **Rec.709**. Plenty of indie 4K deliverables are SDR. An **HDR** master is the one that moves to a **BT.2020 container with the PQ / ST 2084 EOTF**, at 10- to 16-bit, usually with an SDR derivation. Where HDR is required, expect **HDR10, Dolby Vision, or both**. The buyer decides, so read the schedule rather than assuming. See HDR Mastering.
- **Heads and tails.** Modern file deliverables want the file *ready to play*: about one second of black at head and tail, no bars/tone/slate/2-pop, but some schedules still want a full broadcast head (bars, tone, slate, program at `01:00:00:00`). This is a classic per-buyer variable: **build it when you know the spec**, not before.

Audio: what to mix, and to what target

Budget for both the **stem ladder** and separate **theatrical and near-field mixes**. Reconstructing either after the session closes is expensive.

The stem ladder: build it during the mix

Produce **linear PCM at 48 kHz/24-bit**, a 5.1 mix, a true stereo or Lt/Rt fold-down, discrete dialogue/music/effects stems, and a **fully-filled M&E** containing every effect in the composite. A partial M&E can block foreign dubbing and international licensing.

Print these elements while the mix session is open. They are inexpensive to create during the mix and expensive to reconstruct later.

Theatrical vs. near-field: two mixes, two targets

A cinema mix and a home/streaming mix are separate deliverables:

- A **theatrical mix** is monitored on a dubbing stage at **reference level** (each screen channel at 85 dB SPL from −20 dBFS pink noise, Dolby fader at 7) through the cinema **X-curve** (SMPTE ST 202). It is mixed *by ear* with wide dynamic range and is **not loudness-normalized to a number.** This is what feeds a DCP.
- A **near-field mix** (home / home-theater / streaming) is monitored on near-field speakers at lower level and **delivered to a specific integrated-loudness target** measured per ITU-R BS.1770 (LKFS/LUFS), with narrower dynamic range so it holds up on a TV, laptop, or phone.

Do not deliver the theatrical printmaster to a streamer. Derive a near-field pass with the required dynamic range and loudness target.

Loudness targets by destination

Destination	Integrated loudness	Max true peak
Theatrical / DCP	Not normalized. Mixed at calibrated reference (85 dB SPL, X-curve)	Headroom. No fixed cap
Netflix (near-field mix)	**−27 LKFS** (±2 LU), **dialog-gated** (BS.1770-1)	**−2 dBTP**
Apple (iTunes / Apple TV feature master)	**−24 LKFS**: stated bare, with no tolerance or measurement standard given	**−6 dB** (unit not specified in the source)
Amazon (Prime Video Direct)	**No figure published**	**No figure published**
US broadcast (ATSC A/85 / CALM Act)[44]	**−24 LKFS** (±2 dB). Dialogue-anchored for long-form, but **full-program mix for commercials**	**−2 dBTP** (±0.5)
Streaming (ATSC A/85:2026 Annex L)	Pick **one** value between **−23 and −27 LKFS** and apply it consistently to everything on the service	−2 dBTP
EBU R128 (European broadcast)	**−23 LUFS** (±1), **program-gated**	−1 dBTP

!!! warning “These numbers are not directly comparable” Netflix and ATSC A/85 long-form are measured **dialogue-gated**. EBU R128 is **program-gated**. Match both the number and the gating method. Under A/85, use dialogue gating for long-form and full-program measurement for commercials. For long-form, use BS.1770-1 or disable relative-level gating on a later meter.

Netflix accepts either **−27 LKFS dialogue-gated** (BS.1770-1) or **−24 LKFS full-program** (BS.1770-4), ±3, and flags a mix only when both fail. Do not measure −27 LKFS program-gated. That produces an unnecessarily quiet mix.

ATSC A/85:2026 permits streaming services to choose a consistent target between −23 and −27 LKFS. The apparent spread partly reflects different measurement methods.

Without a buyer specification, ask the mixer for roughly **−24 LKFS full-program with a −2 dBTP ceiling**, and log the dialogue-gated value at the same time. This preserves a practical route to Netflix, US broadcast, and EBU R128 deliveries.

Independent VOD schedules remain inconsistent. Many name no LUFS target. **Confirm both the target and gating method in the buyer’s schedule.**

Keep the printmaster, stems, and mix session archived and accessible. A changed loudness target may require a new stage pass and QC cycle.

What to ask your mixer to bid

At bid time, specify every plausible destination:

- **If theatrical is at all possible, bid a theatrical mix *and* a near-field mix.** Mixing theatrical first and deriving the near-field is the cheaper order. Retrofitting theatrical from a home mix is not.

- **Bid the full stem and M&E ladder**: 5.1 and stereo, discrete DME stems, and a fully-filled M&E, produced during the mix.
- **Bid the near-field loudness master to a real spec** (e.g. −27 LKFS for streaming), not "a stereo bounce."
- Even a streaming-only film needs that near-field mastering pass done to the platform's loudness target. Budget it instead of discovering it at QC.

Textless, forced narrative, and the "three masters" problem

Distributors may require three picture versions:

- the **texted** master (titles, credits, and any on-picture text present),
- a **non-subtitled** master (all creative text present, but no burned-in subtitles), and
- a **fully textless** master: every shot that has text over picture supplied clean, so the distributor can localize it. (Rolling end-credit crawls are usually exempt. End-credit *cards over picture* are not.)

Keep **forced narrative** subtitles for foreign dialogue or signs separate so each territory can replace them. **Never burn subtitles into the only master.** Rebuilding a clean version later becomes a new package and may limit international sales.

Captions and subtitles

Caption and localization files are their own deliverable stack:

- **Closed captions**, usually English SDH, in the buyer's required format such as `.scc`. Legal captioning obligations do not prescribe one file format. Reading speed, line length, duration, and line-count limits also vary. Follow the delivery schedule rather than a generic caption template.
- **Subtitles** in a format that varies by house, usually **`.srt`**, **`.stl`**, or **`.itt`**, with foreign-language subs (Latin-American Spanish and French are the most-requested) expanding the sale.
- A **CCSL** (combined continuity and spotting list) and/or **dialogue list** is a standard legal and localization deliverable. See the DCP chapter for how caption and continuity files also feed accessibility tracks (OCAP/CCAP) on the theatrical side.

AI captioning

Use automatic captioning to create a rough transcript or initial timing pass, not a delivery master.

!!! warning "AI captions are a draft, not a deliverable" Raw automatic output commonly fails platform QC on:

```
 - **SDH non-speech information**: captions must identify speakers and describe music, sound
   effects, and off-screen audio. Auto-transcription captures none of it.
- **Reading speed and formatting**: hard limits on words-per-minute / characters-per-second,
line
  length (≈ 32 characters), two lines, and minimum/maximum durations that raw output
ignores.
- **Timing and sync**: in/out must track the edit precisely. Auto-timing drifts, especially
over
  overlapping dialogue.
- **Accuracy on the hard parts**: names, technical terms, accents, and homophones are
exactly
  where it errs, and exactly what a viewer notices.

A distributor may repair a failed file and charge back the cost, sometimes with a penalty.
```

Budget a professional captioning vendor or, at minimum, a human conform and QC pass against the delivery specification.

Quality control

Almost every distributor requires the masters to **pass QC at an approved third-party facility** before delivery is accepted, and the licensor usually pays for it. This is not a formality:

- Facilities are frequently **pre-approved**: some distributors will not accept delivery from a lab not on their list.
- QC failures are typically **recouped against you**: a distributor may repair a failing file and charge back the cost, sometimes with a penalty, before paying revenue.
- The bar can be a **specific standard**: an international-broadcast QC rather than a domestic-release pass, or a clause that HD masters are not approved until a Blu-ray QC pass.

Include a QC cycle and possible re-export and re-QC in the delivery timeline.

The legal and metadata package

A full schedule can contain **thirty-plus legal and metadata items**, including:

- **E&O insurance** (errors & omissions): effectively universal. Limits commonly cluster around **$1M per occurrence / $3M aggregate** with a deductible near $25K and a multi-year term, and can run higher at the top end. Confirm the exact limits and exclusions your distributor requires.
- **Chain of title** and a current (≤ 30-day) **copyright & title report**, proof of copyright registration, and often a **UCC search** confirming no secured creditor interests in the film.
- A **music cue sheet** and executed **music licenses** (and composer agreements).
- A **CCSL / dialogue list**, **metadata** sheet, **billing block**, **key art** (usually layered, with fonts), festival laurels, cast/crew list, and photo IDs & clearances.
- **Lab access letter**, **certificate of origin**, **talent agreements**, and, for some buyers, an **ISAN** number.

Timelines are long and front-loaded: aggregators commonly want **all approved elements 90–120 days before the street date**, and the delivery clock often does not start until *every* item is received and QC'd. Set a release date only against a realistic deliverables schedule.

For a large schedule, budget a **post-production supervisor or delivery coordinator**.

Trailer and marketing deliverables

Match the trailer to the feature's aspect ratio, frame rate, finish, and colour pipeline. If another facility handles marketing, supply the same LUTs, transforms, and graded elements. Confirm whether the distributor wants a rating card supplied or omitted. Also expect a vertical 9:16 cutdown, platform-specific key art, and an H.264 review screener.

A trailer is not one deliverable

Plan for repeated trailer versions: announcement, dated, in-theaters, TVOD/SVOD, broadcast, social, and an **evergreen** version with no date.

Preserve the source package:

- **Deliver the trailer textless**, the same way you deliver the feature textless. Every one of those reissues is a new end card over the same picture, and re-cutting from a texted master is avoidable work.
- **Make sure the graphics package survives the campaign.** Layered project files, fonts and key art need to be available *and licensed for continued use*. Artwork commissioned for a theatrical campaign is not automatically cleared for a home-video card two years later, and discovering that when you need a new title treatment is expensive.

Trailer shots are VFX work nobody scheduled

Trailers often contain earlier shot versions and a deliberately stronger grade. Track these as marketing versions rather than treating them as feature-final shots.

Budget trailer-specific VFX versions for:

- different **frame ranges** than the feature version needs
- **censored** or softened variants for a green-band cut
- alternate or **replaced dialogue**
- **spoiler-free** versions that conceal a reveal the feature depends on

Charge this work to the marketing budget and schedule it as overage work.

Trailer requests interrupt feature work. Decide in advance which artists and shots can absorb the schedule impact.

DCP: The Digital Cinema Package

A DCP is the theatrical deliverable muxed from the DCDM. Festival delivery requires decisions about format, resolution, audio, encryption, accessibility, transport, and QC. See Generalized Production Workflow for the mastering path.

What a DCP is

A DCP stores essence in MXF track files, assembles it with a **CPL** (Composition Playlist), inventories it with a **PKL** (Packing List), and locates it with an **ASSETMAP**. IMF later reused this architecture for mastering and distribution.[45] DCP essence targets cinema playback:

Component	Encoding
Picture	**JPEG 2000**, intra-frame, 12-bit, in **DCI X′Y′Z′**: gamma 2.6, 48 cd/m² reference white (see Color Spaces). X′Y′Z′ is itself gamut-unbounded. The practical P3 limit is the *projector's*, not the container's
Audio	24-bit linear **PCM**, up to 16 channels at 48 or 96 kHz
Subtitles / captions	**SMPTE Timed Text** (ST 428-7) or Interop XML, with PNG or font-rendered glyphs

What is actually in the folder

A DCP is a folder. Almost all of its size is a small number of very large **MXF track files**: the picture, the sound, and any subtitle or caption essence. The rest is three small XML files, together a few tens of kilobytes describing a few hundred gigabytes.

Those three are what people find confusing, because **all three describe the same assets, for three different purposes**:

File	The question it answers	What it records for each asset
ASSETMAP	*Where is it?*	The asset's UUID and its path on the delivered medium
PKL (Packing List)	*What is in this delivery, and did it arrive intact?*	UUID, file size, hash, and asset type
CPL (Composition Playlist)	*What plays, in what order?*	Which reel and role (picture, sound, or subtitle), plus the entry point and duration used from that file

Two distinctions matter:

- **The CPL is the version.** The same track files with a different CPL is a different cut. That is how a version file (VF) reuses an original version's picture and sound without copying a byte of them, and it is the same idea IMF later built its supplemental packages on.

- **The PKL is an inventory, not a playlist.** It can perfectly well list files that no CPL references, which is exactly why it exists as a separate document.

The ASSETMAP is neither. It is a lookup table that turns a UUID into a file path, so that a server handed a folder can find the assets the other two documents refer to.

Two package flavors exist:

- **Interop (IOP)** is the older format. It supports fewer features but plays on legacy servers.
- **SMPTE DCP** is the modern standardized format. High frame rates, immersive audio, and properly signalled closed captions require it.[46]

Default to SMPTE DCP using the Bv2.1 profile (SMPTE RDD 52:2020). Major festivals still accept Interop, but Venice prefers SMPTE and Berlinale recommends it. Use Interop only for a venue known to require legacy compatibility. Confirm the festival specification before authoring.

Resolution: 2K, 4K, and the bit-rate cap

DCPs come in two resolution classes, **2K** and **4K**, each with three containers. The rasters are fixed, and Flat is **not** the full width of the container:

Container	2K	4K
Flat (1.85:1)	1998 × 1080	3996 × 2160
Scope (2.39:1)	2048 × 858	4096 × 1716
Full (1.90:1)	2048 × 1080	4096 × 2160

DCI caps JPEG 2000 image essence at **250 Mbit/s** for both 2K and 4K.[47] A 4K DCP therefore spends the same data budget on four times the pixels.

A 4K DCP plays on a 2K projector. Do not author a separate 2K package for 2K screens. A 2K DCP likewise plays on 4K systems.

JPEG 2000 is **resolution-scalable**: a 4K codestream contains a complete 2K resolution level. A 2K server decodes that level without scaling or transcoding. This is **single inventory**: one package for both resolutions.

If the show finished in 4K, deliver 4K. The fixed bit-rate cap means the gain is smaller than the pixel count suggests and is most visible in detailed, high-contrast material. Deliver 2K when smaller transfers matter more. Either package remains broadly compatible.

Frame rates and the general-release envelope

For broad compatibility, follow **SMPTE DCP Bv2.1**: 2D 4K at 24 fps. SMPTE ST 429-2 permits 4K at 24, 25, or 30 fps, and DCI's 2024 HDR addendum permits 2D 4K at 48 or 60 fps with a 450 Mbit/s ceiling, but those modes fall outside the general-release profile. Stereoscopic DCP remains 2K.

Bv2.1 is an interoperability profile, not a physical limit. Dolby Cinema, IMAX, and other premium formats use private specifications that may exceed it. **Get those specifications from the vendor.**

Audio: channels and how it was mixed

DCP audio is discrete, uncompressed multichannel PCM, but *how many* channels, and *how it was mixed*, matter more than the container.

Prefer 5.1 over 2.0. Its center channel anchors dialogue to the screen for every seat. A 2.0 mix creates a phantom center that pulls toward the nearer speaker away from the room's centerline. Many festivals expect 5.1.

2.0 is not Dolby Stereo. LoRo is discrete left/right with no center or surround information. Lt/Rt is matrix-encoded and decodes to L, C, R, and surround. Identify a two-channel delivery explicitly. Treating one as the other misroutes dialogue and surround.

Mix for the theater, not the nearfield. Use a calibrated dubbing stage at reference level with the cinema X-curve (SMPTE ST 202 / ISO 2969).[48] A nearfield mix usually plays too loud and bright in a theater. If theatrical exhibition is possible, budget a theatrical mix rather than a nearfield bounce.

Playback level and the loudness spiral

Cinema reference is Dolby fader **7.0**, calibrated to **85 dB SPL per screen channel**. Some theaters play below reference because some films are mixed excessively loud. Filmmakers then mix louder in anticipation, giving theaters more reason to turn down. The cycle sacrifices headroom and dynamic range, while films mixed properly at reference are punished by playing too quietly.

Mix at reference, or only slightly below it when that is a deliberate creative choice. Do not pre-compensate for an assumed low playback level by crushing the mix or pushing everything louder.

For private and festival screenings:

- QC the DCP in the actual auditorium whenever possible. Start playback at **7.0** and work down only if necessary. Agree on the final level with the projectionist.
- Account for the audience. A full auditorium absorbs more sound than an empty one, so a mix that feels slightly hot during an empty-room QC may land correctly with a full house.
- If you cannot attend, send the projectionist a note explaining that the film was mixed on a calibrated theatrical stage at reference. Ask them to begin at **7.0** and preserve the intended level unless the room presents a specific problem.

Encrypted or unencrypted

An encrypted DCP uses **AES-128**. Playback requires a **KDM** (Key Delivery Message) encrypted to one server certificate and valid for a specified time window. An unencrypted DCP plays on any compliant server.

Choose encryption before authoring. Changing it later requires new MXF track files, CPL, PKL, hashes, and keys. The JPEG 2000 codestream need not be recompressed, but the package must still be rebuilt.

For festivals, **default to unencrypted** unless a right-holder requires encryption. Venice recommends unencrypted DCPs, Sundance strongly prefers them, and Berlinale accepts either.[49] Encryption requires a certificate and time-limited KDM for every playback server. A wrong certificate, expired window, or last-minute room change prevents playback.

If encryption is required, confirm every server and screening window with the festival and arrange 24-hour key-management support.[50]

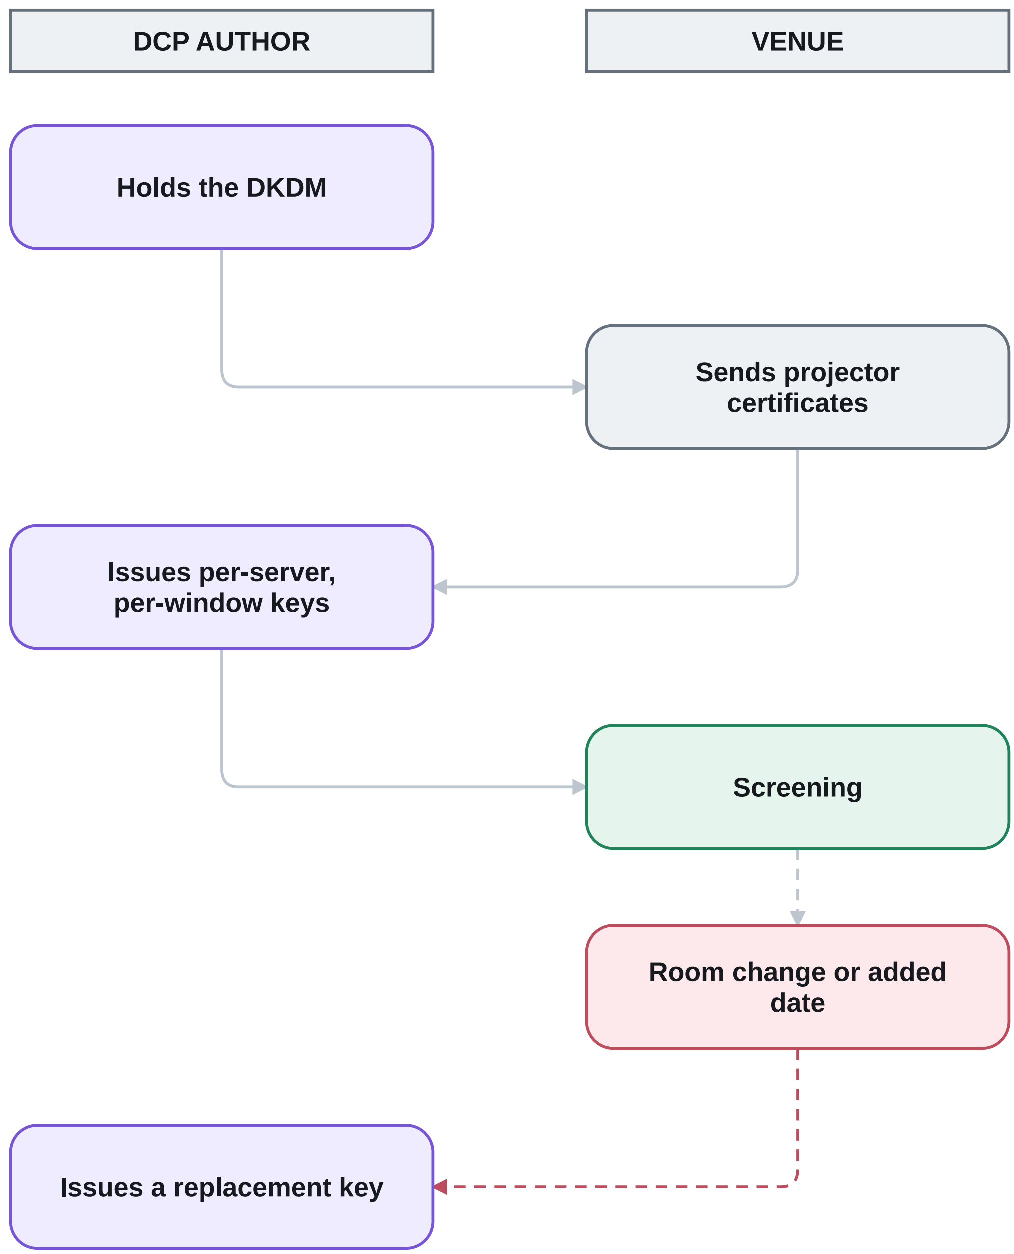

The KDM exchange. Keys are cut per server and per window, so a late room change means going back to the author for another one.

Accessibility: captions and described audio

Accessibility requirements vary sharply by festival:[51]

- **Sundance:** English-language projects deliver both CCAP and OCAP DCPs. Non-English projects deliver standard subtitles plus English SDH.
- **SXSW:** closed captions for every project. English subtitles for non-English content.
- **Hot Docs:** closed captions for English-language features.
- **IFFR:** SDH and audio description strongly encouraged.
- **Berlinale, Cannes, Venice, and Locarno:** no published mandate.

Captions, for Deaf and hard-of-hearing viewers, come in two forms, and the DCP naming convention makes you declare which:[52]

- **OCAP (open captions)** are always visible and require no venue hardware. They are the most reliable accessible screening format.
- **CCAP (closed captions)** feed personal devices such as CaptiView readers or caption glasses. They work only when the venue has functioning caption hardware.

SDH differs from ordinary translation subtitles: beyond the dialogue, it carries speaker IDs and non-dialogue sound cues (music, effects).

Described audio (VI-N) narrates key visual action for blind and low-vision viewers. A separate **Hearing Impaired (HI)** track carries a dialogue-boosted mix. Map HI to channel 7 and VI-N to channel 8.[53] Channel 15 carries sign-language video. Channels 9, 10, and 16 remain blank. Validate the channel map: a misplaced track can pass structural QC but never reach the assistive device.

Provide the tracks named in the festival specification. When venue capability is unknown, an **OCAP version** is the safest accessible deliverable because it requires no caption hardware.

Delivering to a festival

For **physical delivery**, use the medium and file system named by the festival. CRU DX115 drives remain common. USB 3.0 SSDs are increasingly accepted.[54] For maximum server compatibility, format Linux ext2/ext3 with **128-byte inodes** under ISDCF Doc 3. Some festivals also accept exFAT or NTFS.[55]

Electronic delivery is rising: Venice takes files by **Aspera**, Berlinale through its own **Digital Cinema Portal**, and services such as CineSend and Massive move DCPs over managed transfers.[56] A feature DCP may be hundreds of gigabytes, so budget enough upload time and bandwidth to meet the deadline.

Name the package under the **ISDCF Digital Cinema Naming Convention (DCNC)**.[57] Write that name into **ContentTitleText** and, where used, **AnnotationText** inside the CPL and PKL. Renaming the folder alone does not change the title shown by a cinema server.

QC on a real cinema server

Before delivery, **play the DCP from a cinema server through a cinema projector.** Desktop playback does not prove that the package will ingest or screen correctly.

Run Video Village's DCP Inspector before the theater QC. It performs deep validation of the CPL, PKL, ASSETMAP, hashes, reel structure, frame rate, audio layout, and timed-text references. Validation catches structural faults but does not replace cinema-server playback.

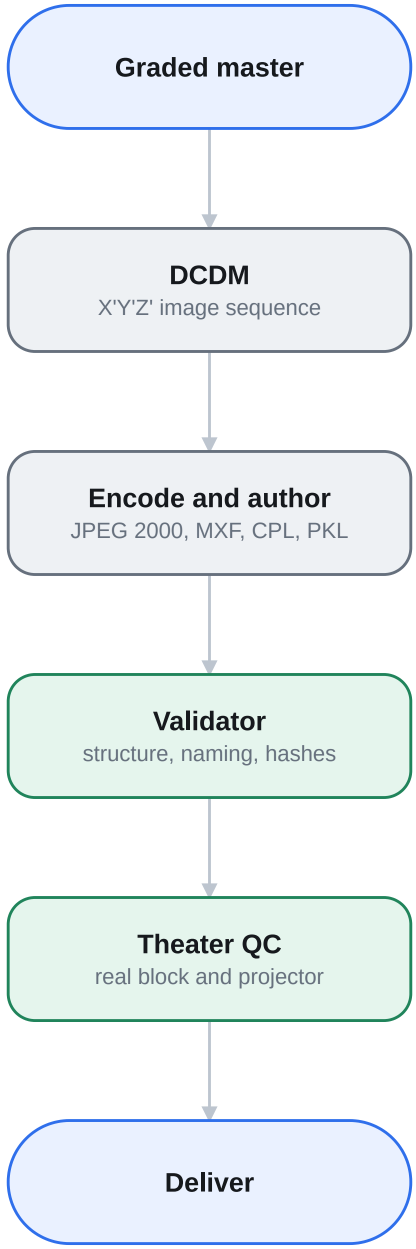

Authoring through to delivery. A validator and a theater QC pass catch different classes of fault, and either can send the package back to authoring.

Common DCP authoring failures

Check these before delivery:

!!! tip “Validate before returning to the authoring vendor” DCP Inspector can repair many common CPL errors directly. Use it to validate and fix the package before sending it back to the authoring vendor or platform.

- **Naming:** follow the ISDCF convention. Match the aspect-ratio token to the raster so the server selects the correct masking preset.
- **Content kind:** match the CPL’s `ContentKind` (`feature`, `trailer`, `test`, `psa`, and so on) to the composition name.
- **Leaders and markers:** omit bars, tone, slates, countdowns, 2-pops, and tail pops. If used, keep `FFOC` and `LFOC` paired, ordered, and inside the correct reel.
- **Audio:** use DCI channel order (L, R, C, LFE, Ls, Rs, HI, VI-N), confirm every track contains the intended mix, and check true peak for inter-sample clipping.
- **Timed text:** start tracks at `EntryPoint=0` and preserve the track structure across every reel. Interop OV and VF packages each need their own subtitle files and UUIDs.
- **Integrity and packaging:** regenerate PKL hashes after any re-wrap, include every CPL-referenced asset, remove stray files, and apply encryption consistently.
- **Bit rate:** check peak instantaneous rate, not only the average. Detailed frames can exceed the cap and stutter even when the package average is compliant.

Validation catches structural errors. Theater QC confirms presentation.

A theater QC pass on a DCI media block (Dolby, GDC, Sony, Barco Alchemy, Qube) and projector is the only place to catch:

- **Ingest and validation** errors, and Interop/SMPTE incompatibilities with a particular server generation.
- **Audio channel mapping.** Verify that center, surrounds, LFE, HI, and VI-N land where they should, not rotated or swapped.
- **Subtitle and caption rendering.** Check timing, position, glyphs, and, for CCAP, that the captions reach the seat device in sync.
- **Framing and masking**: the right aspect, no cropping, correct Flat/Scope container.
- **Frame-rate support**, dropped frames, and audio-video sync across the whole runtime.
- **KDM validity.** For encrypted packages, confirm that the key opens on *that* server within the window.

Play the film start to finish. Listen to the full mix and verify captions on the actual assistive device.

Delivery checklist

- **Decide encryption early**, and default to unencrypted for festivals unless a rights-holder demands otherwise.
- **Plan accessibility early.** Author OCAP/SDH and described-audio assets at the correct frame rate before festival deadlines.
- **Get each festival’s exact spec first.** “A DCP” is not a specification: Interop vs SMPTE, media and delivery method, naming, and which accessibility tracks all need to be settled before mastering.
- **QC on real hardware.** A desktop player proves the package parses. Only a cinema server proves it screens.

IMF: The Interoperable Master Format

IMF is a component-based mastering and delivery format used for streaming and international distribution. Unlike a flat master, it stores reusable picture, audio, and subtitle elements, then assembles each version from a playlist.

The problem IMF solves

A **flat** deliverable is one continuous render per version. Domestic, airline, language, music-rights, SDR, and HDR variants therefore duplicate nearly identical picture.

IMF stores the essence once as track files. Each version is a playlist referencing the required segments. A changed version can therefore contain only its replacement segments and a new playlist. **IMF is a versioning format.**

Structure

The unit of delivery is an **IMP**: Interoperable Master Package. ST 2067-2 §7.2 defines it as one IMF Packing List plus all the assets it references. A *Delivery* is the ASSETMAP-mapped file set carrying one or more IMPs. “IMF package” is loose usage that will be understood, but IMP is the term the specifications and the QC reports use.

An IMP is a set of files with defined roles:

Component	Role
Track Files	The essence itself as separate picture, audio, and subtitle MXF files
CPL (Composition Playlist)	Describes one *version*: the track files, order, and timing
PKL (Packing List)	Inventory of the files in this delivery, with hashes for integrity
ASSETMAP	Maps asset identifiers to file paths on the delivered medium
OPL (Output Profile List)	Optional. Describes how to derive a specific output from a CPL
Sidecar assets	Optional non-essence files carried with the package (ST 2067-9): scripts, QC reports, artwork

A **Supplemental Package** contains replacement segments and a new CPL while referencing unchanged track files from the original package.

A CPL resembles an EDL: it references external essence by identifier and timecode instead of baking one continuous render.

Validating a package

Photon is Netflix’s open-source IMF validator. It checks the ASSETMAP, PKL, CPL, App #2E constraints, and MXF header metadata against the declared essence descriptors.

Photon does **not** decode pixels or verify PKL hashes. It may also warn rather than fail on unsupported namespaces. Use it for structural validation, then verify hashes and watch the picture separately.

DaVinci Resolve Studio exposes Photon validation directly. Run the full local validator before using Netflix's reduced delivery-interface check.[58]

Application #2E

IMF is a framework with **applications** that constrain it to a usable profile. App #2E is the one that matters for streaming delivery. Per ST 2067-21, it uses image essence coded as a **JPEG 2000** codestream and audio as **linear PCM**, with track files conforming to SMPTE ST 379-1 and ST 422, and picture essence described by a CDCI descriptor (for Y′C′BC′R) or an RGBA descriptor (for R′G′B′).

JPEG 2000 is intra-frame and frame-addressable, so a version can change at any frame without temporal dependencies. **ST 2067-21:2023 added HTJ2K** (high-throughput JPEG 2000), which encodes and decodes far faster than the original block coder at the same quality.

Other applications include **App #4** for archival cinema mastering, **App #5** for ACES essence, and **RDD 59-1** for ProRes essence under RDD 45.

Netflix requires IMF App #2E conforming to SMPTE ST 2067-21:2016, :2020, or :2023. Immersive-audio packages must also conform to ST 2067-201:2019 (IAB) Level 0.[59]

Production requirements

IMF authoring is usually handled by a mastering facility or specialist vendor. The production must supply material from which that vendor can build the package:

- **Supply textless elements.** Every title, caption, and dub card that may vary by territory needs a clean background with handles.
- **Plan versions during post.** Organize known alternate cuts and elements during the DI. Adding them later means reopening the master.
- **Enforce naming and versioning.** Confirm every supplemental CPL references the intended segment versions.
- **Get the complete specification before the DI.** "IMF" alone is insufficient. Confirm the application, audio configuration, HDR variant, and whether supplemental packages are expected.

Archival

Archive the masters, projects, dependencies, and documentation needed to remaster the show. Confirm the studio or distributor specification, then include the following elements in addition to the distribution master.

Ungraded Archival Master

Render the complete conformed picture with final VFX, opticals, titles, and textless elements at the highest applicable resolution. Use camera-native log or scene-linear encoding, typically 16-bit DPX or 16-bit float OpenEXR. This is the digital equivalent of an assembled negative.

Graded Archival Master

Match the ungraded archival master's format and resolution, but include the final grade in the DI **working color space**. In a managed pipeline, archive the scene-referred grading space, not any camera-native encoding. Keep grading operations before the display transform so future masters can be generated from the same graded source.

Show Look Up Tables (LUTs)

Archive every LUT and transform used by the project: the display transform, creative looks, camera normalizations, technical conversions, and DCTLs. A future session must be able to rebuild the original pipeline.

Two dependency traps are worth checking before an archive is considered complete:

- **Encrypted DCTLs (`.dctle`)** can carry an **expiry date** or a machine/license lock. They may simply stop working after a certain date or on a different system, silently breaking a re-grade years later. Treat an encrypted DCTL as an archival liability and obtain an unencrypted or license-free equivalent wherever possible.
- **OpenFX (OFX) plugin dependencies.** A grade that leans on third-party OFX plugins will not reproduce without those plugins and the correct versions and licenses. It is the same failure mode as a missing LUT.

LUTs and their dependencies can be gathered and bundled for archive with Video Village's DRP Inspector, which reads a DaVinci Resolve project and **reports any unexpected or at-risk dependencies** (external LUTs, DCTLs, OFX plugins) so they can be collected before the project is stored.

Reference EDLs

Flattened EDLs of the finalized DI timeline, matching the master and archival renders, are generated so that they may be easily split (or notched) at edit points in a later session, possibly on a different platform than the initial DI.

Archival Documentation

An archive without instructions is only a collection of files. Write the documentation for a qualified person who did not work on the production and may restore it years later on different hardware and software.

The document should identify the authoritative version, explain how to rebuild it, and warn against plausible but incorrect uses of the material. Include:

- **Document control:** Project, archive version, date, author, technical contacts, scope, superseded documents, and the package that represents final creative intent.
- **Source material:** Asset families, originating department or vendor, file formats, rasters, frame rates, image state and color encoding, timecode source, naming rules, handles, mattes, stereo structure, and known gaps.
- **Project environment:** Application, version and build, operating system where relevant, project or database organization, reel or longplay structure, bin layout, track logic, and versioning method.
- **Conform:** FFOA and LFOA, frame ranges, source-to-timeline relationship, editorial replacements, layer rules, and any alternate cuts or branches.
- **Framing and scaling:** Working raster, pixel aspect ratio, resize method, framing charts, output rasters, extraction rules, and deliverable-specific reframes that cannot be recreated by a simple crop.
- **Color pipeline:** Input encodings, grading space, color-management configuration, LUT and transform order, display targets, white point, graphics treatment, output transforms, and the exact dependency files.
- **Text and graphics:** Fonts, source applications, alpha or matte handling, color encoding, texted and textless ranges, and instructions for recompositing elements.
- **Exceptions:** Shot or timecode ranges that use alternate sources, eye swaps, speed treatment, unusual transforms, special titles, or other choices that are not evident from the media.
- **Master lineage:** Every graded, ungraded, theatrical, video, HDR, SDR, texted, textless, language, and accessibility master, plus which versions are historical references and which should drive future work.
- **Relinking and restore:** Archive root, folder map, relink rule, expected offline media, external dependencies, reconstruction order, validation procedure, checksums, storage locations, and the most recent restore-test result.

Use exact filenames, paths, versions, frame ranges, and hashes in the production document. Put long inventories and exception lists in appendices or machine-readable sidecars. Do not bury the only manifest, checksum list, or EDL inside a PDF.

Keep the documentation both **inside the archive package and separately with the archive catalog**. Do not include passwords, license keys, or active credentials. Record the credential owner and access procedure instead.

Start with the DI Archival Documentation template.

Storage: LTO, cloud, and the trouble with hard drives

Use **LTO**, **cold cloud storage**, or hard drives as nearline copies. Compare full lifecycle cost, not media price alone.

The calculator below estimates the cost of **archiving a given amount of data once, holding it, then fully restoring it**: drag the sliders for your data size and retention period. It is a planning aid, not a quote: cloud rates and egress tiers change often, and the LTO figures separate the tape stock (hard costs) from a per-TB service rate.[60]

> The interactive archival-cost calculator is available in the web edition.

```
<table class="ac-t">
  <thead><tr><th>Option</th><th>Hard costs</th><th>Archive</th><th>Storage</th>
<th>Restore</th><th>Total</th></tr></thead>
  <tbody id="ac-body"></tbody>
</table>
```

Estimates at mid-2026 rates for a single full restore to the open internet. Cloud storage is each provider's cold/archive tier. Egress tiers and provider rates vary: verify before budgeting. Backblaze and Wasabi assume the restore falls within their included free-egress ratio. LTO assumes you keep the tapes (no vault fee). See the footnote for the service-rate assumptions.

Read the total, not the headline rate. Cold-cloud tiers are cheap to *store* (~$1/TB/month) but bill heavily on retrieval and egress to get data back. Flat-rate clouds such as Backblaze and Wasabi offer free egress but cost 5–7× more to store. LTO inverts both. You pay a service to write and read the tapes, but the stock is ~$6/TB and holding is free, so its total stays flat and, at any realistic multi-year retention, it is the cheapest option. Glacier Deep Archive only undercuts it for holds under ~two years (drag the years slider to see the crossover). Cloud's real value is as an *offsite* copy with managed redundancy and no hardware to migrate. It is also a permanent bill exposed to price hikes and account risk.

Hard drives are not an archive. The usual claim that "an unpowered drive only holds data for a few years" conflates two different things and misses the real risk. Magnetic remanence on a shelved platter is good for a long time. What actually fails is everything around it.

The real failure modes are mundane and much more likely:

- **The drive nobody is monitoring.** An archive is a *process*, not an object. A drive on a shelf is not being checksummed, so corruption is not detected, and there is no schedule that would surface a problem while a second copy still exists.
- **Mechanical trouble on spin-up.** Drives that sat idle for years can fail to start: stiction, dried lubricant, seized bearings, degraded electronics. The data may be intact and unreachable without a recovery house, which is the moment you discover the archive was never tested.
- **Silent bit rot.** Without checksums, corruption propagates into your only copy and into every backup made from it, undetected until someone tries to use the file.
- **Single point of failure.** One shock-sensitive object, one filesystem, one enclosure, one interface standard that will eventually stop being current.

Hard drives are fine as a nearline copy. Never use one as the only long-term copy.

On a bonded or financed film, proper archival is contractual, not optional. A completion guarantor guarantees delivery *to spec*, including preservation elements. Negative and digital-media insurance assume professional storage, so archival masters generally must live on **LTO with checksums, redundant copies, and offsite/vaulted storage**. A lone hard drive satisfies neither guarantor, distributor QC, nor insurer.

LTO is the archival standard, but it requires migration. Tape is offline and air-gapped, rated ~15–30 years, and self-describing via **LTFS**. The catch is generational, and it just got sharper. Historically a drive read two generations back and wrote one, which is where the familiar "migrate every 7–10 years" advice comes from. **LTO-10 (shipping since June 2025) dropped backward compatibility altogether: an LTO-10 drive cannot read *or* write LTO-9 media.**[61]

The practical rule is therefore not a calendar interval. **Migrate on the life of your drive, not the life of the tape, and never retire the last drive that can read your oldest cartridges.** If you hold an LTO-9 archive, buying an LTO-10 drive does not extend it. It strands it unless you keep a working LTO-9 drive or migrate the data first. The tapes will outlast the hardware that reads them. That is the failure mode, and it is silent until the day you need a restore.

!!! tip "The smart default: 3-2-1, not one of anything" Keep **three copies, on two different media, with one offsite.** For an independent film that usually means **two LTO copies, or one LTO plus a cold-cloud offsite copy**. Verify each with checksums (an **MHL**, media hash list). Retain the graded and ungraded archival masters, textless elements, the audio stems and M&E, and the conform project. Whatever you choose, do not let a single hard drive be the only thing standing between you and a future remaster.

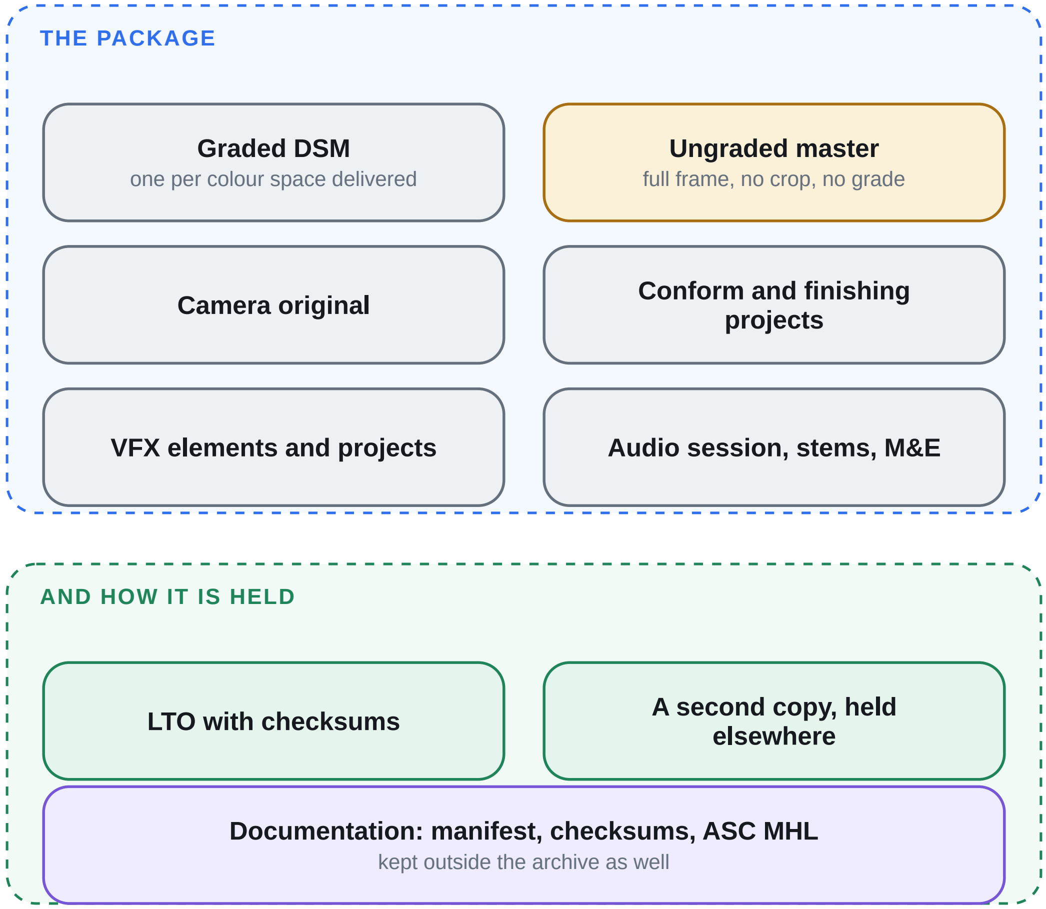

What belongs in a preservation package. The documentation travels separately as well as on the tape: a manifest that exists only on the LTO it describes is no help when deciding whether to restore.

Production Tools

Production Checklists

Use these as gates, not as substitutes for a tested show specification.

Pre-Production Pipeline

- ☐ Name an owner for the image pipeline.
- ☐ Name an owner for every handoff and define the escalation path.
- ☐ Confirm every camera, sensor mode, recording format, raster, and frame rate.
- ☐ Choose and document the color-management approach.
- ☐ Define input transforms, working spaces, viewing transforms, and output transforms.
- ☐ Establish the Show LUT or viewing pipeline.
- ☐ Review principal and VFX lens choices with the DP and VFX supervisor.
- ☐ Shoot distortion grids and lens charts for every production lens.
- ☐ Test Cooke /i, ARRI LDS, ZEISS eXtended Data, or other required camera and lens metadata.
- ☐ Establish sequence, shot, file, folder, frame-number, and version conventions.
- ☐ Define the tracking system and source of truth for assets, versions, notes, and approvals.
- ☐ Define transfer, access, security, checksum, backup, and retention methods.
- ☐ Set editorial proxy format and metadata requirements.
- ☐ Set VFX plate, working, review, and final-delivery formats.
- ☐ Set DI turnover and conform requirements.
- ☐ Obtain likely distributor, streamer, broadcaster, and festival specifications.
- ☐ Define HDR, SDR, theatrical, and archival masters.
- ☐ Issue one versioned show specification and vendor-onboarding package.
- ☐ Complete camera-to-dailies-to-VFX-to-DI round-trip tests.
- ☐ Review the tests on calibrated target displays.

VFX Vendor Onboarding

- ☐ Send the vendor contact list and escalation path.
- ☐ Send the image-pipeline and delivery specifications.
- ☐ Send the current OCIO configuration, LUTs, DCTLs, or other transforms.
- ☐ Confirm application versions and the VFX Reference Platform.
- ☐ Confirm transfer method, security policy, and directory structure.
- ☐ Confirm shot names, file names, version numbers, and frame numbering.
- ☐ Confirm plate handles and required comp handles.
- ☐ Confirm review media, burn-ins, slate, and audio requirements.
- ☐ Send a representative test shot.
- ☐ Approve the round trip before assigning production shots.

VFX Plate Pull

- ☐ Remove timewarps and pull every source frame needed to rebuild them.
- ☐ Remove editorial color, display transforms, resizes, and repositions unless specified.
- ☐ Include the agreed handles.

- [ ] Preserve source dynamic range and gamut.
- [ ] Match the specified raster, pixel aspect ratio, and data window.
- [ ] Preserve or provide source clip name, timecode, tape name, and camera metadata.
- [ ] Include framing reference, editorial reference, and pull manifest.
- [ ] Validate first and last frame numbers.
- [ ] Check the plate through the approved viewing transform.

VFX Final Delivery

- [ ] Use the approved shot name and version.
- [ ] Deliver the exact frame range and handles.
- [ ] Match the specified format, compression, channels, and bit depth.
- [ ] Match the plate raster, pixel aspect ratio, display window, and data window.
- [ ] Render in the specified scene-referred output encoding.
- [ ] Remove slates from the production frame range unless instructed otherwise.
- [ ] Include required mattes and label every channel.
- [ ] Check for missing, duplicated, frozen, or corrupt frames.
- [ ] Compare against the plate and previous approved version.
- [ ] Gamma-, exposure-, and difference-check the complete frame.
- [ ] Confirm the review movie and full-quality render show the same version.
- [ ] Include a delivery manifest and checksum.

DI Editorial Turnover

- [ ] Lock or clearly label the turnover cut.
- [ ] Flatten unnecessary tracks and remove obscured clips.
- [ ] Decompose nested sequences and multicam clips.
- [ ] Join through-edits where possible.
- [ ] Keep titles, VFX, opticals, and camera originals on agreed tracks.
- [ ] Preserve source names, tape names, timecode, and frame rate.
- [ ] Export the requested EDL, XML, AAF, or FCPXML.
- [ ] Export a reference movie with timecode, source name, and sequence name.
- [ ] Provide leaders, two-pops, tail-pops, and sync references.
- [ ] Identify every optical, speed change, resize, reposition, freeze, and split screen.
- [ ] Identify the exact approved version of every VFX shot.
- [ ] Consolidate and checksum required source media.
- [ ] Compare the first conform against the offline reference.

Mastering and Delivery

- [ ] Obtain the current delivery schedule from the recipient.
- [ ] Confirm raster, aspect ratio, frame rate, color encoding, codec, and wrapper.
- [ ] Confirm HDR and SDR mastering requirements.
- [ ] Confirm audio configuration, stems, M&E, sample rate, and loudness.
- [ ] Confirm captions, subtitles, accessibility tracks, and language versions.
- [ ] Confirm texted and textless picture requirements.

- [ ] Confirm DCP, IMF, ProRes, or other package requirements.
- [ ] Confirm naming, metadata, artwork, trailer, and legal deliverables.
- [ ] Run automated file validation.
- [ ] Watch and listen to the complete final deliverable.
- [ ] Validate captions and accessibility tracks on the target system.
- [ ] Retain delivery manifests, checksums, and acceptance reports.

Archival

- [ ] Ungraded archival master
- [ ] Graded scene-referred archival master
- [ ] Textless elements
- [ ] Audio masters, stems, and M&E
- [ ] Conform, grade, VFX, title, and graphics projects
- [ ] LUTs, OCIO configs, DCTLs, fonts, and plugin dependency list
- [ ] Reference movie and flattened EDL
- [ ] Archival documentation with master lineage, relink instructions, exception list, and restore test
- [ ] File manifest and checksums
- [ ] Three copies on two media, with one copy offsite
- [ ] Documented restore test and media-migration owner

Technical Specifications

Complete these fields before work begins. Replace every bracketed item. Do not leave choices implied.

Image-Pipeline Specification

Field	Production value
Show	`[title / code]`
Pipeline owner	`[name / email / phone]`
Color-management system	`[ACES / OCIO / RCM / TCS / manual]`
Configuration and version	`[config filename / repository tag / hash]`
Camera inputs	`[camera, gamut, transfer function, decode settings]`
Scene-referred interchange	`[ACES2065-1 / other]`
VFX working space	`[ACEScg / other]`
Grading space	`[ACEScct / DWG-Intermediate / other]`
Viewing transform	`[name and version]`
Creative look	`[LUT / LMT / DCTL / BLG / CDL and version]`
SDR output	`[Rec.709 / gamma / luminance]`
HDR output	`[P3D65 or Rec.2020 / PQ or HLG / luminance]`
Theatrical output	`[P3DCI / DCI X'Y'Z' / 48 cd/m²]`
Reference displays	`[model / calibration target / date]`
Negative-value policy	`[preserve / clamp only at named output]`
Out-of-gamut policy	`[transform / gamut compression / approved method]`

Attach every configuration, LUT, transform, and test image referenced above.

VFX Plate Specification

Field	Production value
Plate format	`[OpenEXR / DPX / ProRes]`
Compression	`[none / ZIP1 / other]`
Bit depth and channels	`[16-bit float RGB / RGBA / other]`
Color encoding	`[ACEScg / camera log + gamut / other]`
Raster	`[width × height]`
Pixel aspect ratio	`[1.0 / other]`
Display/data windows	`[must match / approved overscan rule]`
Plate handles	`[frames before / after]`
Frame-number start	`[1001 / source timecode / other]`
Naming pattern	`[pattern]`
Source metadata	`[clip name / tape name / timecode / lens data]`
Framing reference	`[filename / burn-in requirements]`
Checksum	`[xxHash64 / MD5 / other]`

VFX Final-Render Specification

Field	Production value
Final format	`[OpenEXR / DPX / ProRes]`
Compression	`[none / ZIP1 / other]`
Bit depth and channels	`[value]`
Color encoding	`[value]`
Comp handles	`[frames before / after]`
Alpha	`[none / straight / premultiplied]`
Required mattes	`[objects and channel names]`
Data-window rule	`[match display window]`
Version pattern	`[v### / vendor-coded v####]`
Slate	`[required / not required / location]`
Delivery path	`[server / folder pattern]`
Manifest and checksum	`[required format]`

Review-Media Specification

Field	Production value
Codec and wrapper	`[H.264 .mp4 / ProRes .mov / other]`
Raster and frame rate	`[value]`
Color encoding	`[Rec.709 gamma 2.4 / other]`
Viewing transform	`[name and version]`
Burn-ins	`[shot / version / frame / timecode / date]`
Slate fields	`[artist / vendor / notes / duration]`
Audio	`[reference mix / silent]`
Review platform	`[Frame.io / cineSync / other]`
Approval authority	`[role / name]`

DI Turnover Specification

Field	Production value
Locked-cut identifier	`[date / version / runtime]`
Timeline frame rate	`[23.976 / 24.000 / other]`
Reel or longplay structure	`[value]`
First frame of action	`[timecode]`
Interchange file	`[EDL / AAF / XML / FCPXML]`
Reference movie	`[codec / raster / burn-ins]`
Camera-original delivery	`[source / consolidated / storage location]`
VFX-delivery location	`[path / manifest]`
Title and graphics delivery	`[format / alpha / color encoding]`
Optical-effects list	`[filename]`
Color metadata	`[CDL / LUT / stills / ALE / other]`
Audio reference	`[guide mix / stems / two-pop]`
Conform approval	`[responsible editor / process]`

Reusable Production Templates

Copy these into the production’s preferred system. Replace every bracketed field and delete anything that does not apply.

VFX Vendor Onboarding

```
PROJECT
Title / code:
Production company:
VFX producer:
VFX supervisor:
VFX editor:
Primary coordinator:
Emergency contact:

SCOPE
Awarded sequences / shots:
Current count:
First turnover:
First WIP:
Final delivery:
Known trailer / marketing work:

PIPELINE
Color-management system and version:
Configuration package:
VFX working space:
Viewing transform:
Show look:
Plate specification:
Final-render specification:
Review-media specification:
Reference application versions:

NAMING AND VERSIONING
Shot-name pattern:
File-name pattern:
Frame-number rule:
Version-number rule:
Plate handles:
Comp handles:

TRANSFER AND SECURITY
Transfer service:
Delivery path:
Encryption requirement:
Checksum and manifest:
Access expiration:
Incident-reporting contact:

REVIEW
Review platform:
Submission cutoff:
Review schedule:
Required burn-ins:
Who may approve:
Where final notes are recorded:
```

```
ONBOARDING TEST
Test shot:
Plate sent:
Vendor render returned:
Round trip approved by:
Approval date:
```

Image-Pipeline Handoff

```
Subject: [PROJECT] image-pipeline package [VERSION]

This package supersedes: [PREVIOUS VERSION / NONE]
Effective date: [DATE]

Working configuration:
- Color-management system:
- Config filename and hash:
- VFX working space:
- Grading space:
- Viewing transform:
- Creative look:

Package contents:
- [FILE]
- [FILE]

Install location or instructions:
[INSTRUCTIONS]

Validation:
1. Open [TEST IMAGE].
2. Assign [INPUT].
3. View through [DISPLAY / VIEW].
4. Compare against [REFERENCE].
5. Return [OUTPUT] without the display transform.

Do not begin production shots until the returned test is approved.

Questions and failures:
[PIPELINE OWNER / CONTACT]
```

VFX Delivery Notice

```
Subject: [PROJECT] VFX delivery [DELIVERY ID]: [SHOT COUNT] shots

Delivery:
Date and time:
Vendor:
Transfer location:
Manifest:
Checksum type:

Shots:
[SHOT] [VERSION] [FRAME RANGE] [STATUS]
[SHOT] [VERSION] [FRAME RANGE] [STATUS]

Technical specification:
Format / compression:
Raster / pixel aspect:
Color encoding:
Alpha / mattes:
Viewing transform used for review:

Notes and exceptions:
[SHOT / VERSION / NOTE]

Supersedes:
[PRIOR DELIVERY OR VERSION]

Delivered by:
QC completed by:
```

Review Note

```
Project:
Sequence / shot:
Version reviewed:
Frame or timecode:
Reviewer:
Date:

Status: [APPROVED / APPROVED PENDING TECH FIX / REVISE / HOLD]

Note:
[ONE ACTION PER NOTE]

Reference:
[ANNOTATION / STILL / PREVIOUS VERSION / PLATE]

Applies to:
[THIS SHOT / SHOT RANGE / ENTIRE SEQUENCE]

Next delivery requested:
[DATE / REVIEW]
```

Shot Change Order

```
Project:
Shot:
Current awarded scope:

Requested change:

Reason:
[EDITORIAL CHANGE / CREATIVE REVISION / NEW SOURCE / TECHNICAL FAILURE / MARKETING]

Inputs supplied:
- [FILE / VERSION]

Schedule impact:
Cost impact:
Downstream impact:

Authorized by:
Authorization date:
Vendor acknowledgment:
```

DI Turnover Notice

```
Subject: [PROJECT] DI turnover [CUT VERSION]

Locked-cut identifier:
Runtime:
Frame rate:
Reel / longplay structure:
First frame of action:

Delivered:
- Edit interchange:
- Reference movie:
- Camera originals:
- VFX:
- Titles / graphics:
- Optical-effects list:
- Color metadata:
- Audio reference:
- Checksums / manifest:

Known exceptions:
[ITEM / OWNER / EXPECTED DATE]

Conform review:
Date:
Production editorial representative:
Approval method:
```

Delivery Acceptance Record

```
Project:
Recipient:
Delivery schedule version:
Delivery ID:
Date:

Package contents:
[ITEM / VERSION / CHECKSUM]

Automated validation:
[TOOL / VERSION / RESULT]

Human QC:
Picture watched by:
Audio checked by:
Captions / subtitles checked by:
Accessibility checked by:

Recipient status:
[ACCEPTED / CONDITIONALLY ACCEPTED / REJECTED]

Outstanding issue:
Owner:
Due date:
Final acceptance date:
```

DI Archival Documentation

```
=== DOCUMENT CONTROL ===
Project:
Production / rights holder:
Archive version:
Prepared by / date:
Technical contacts:
Purpose, scope, and supersedes:
Authoritative package or build:
=== ARCHIVE SCOPE ===
Cuts and versions included:
Deliverable families included:
Elements excluded:
Source of final creative intent:
Historical reference versions:
Versions approved for future mastering:
=== SOURCE MATERIAL (repeat for each asset family) ===
Asset family:
Originating department / vendor:
File format and compression:
Raster and pixel aspect ratio:
Frame rate:
Image state / color encoding:
Timecode source:
Clip / reel naming rule:
Frame-number rule:
Handles:
Alpha / mattes:
Stereo / multiview structure:
Known gaps or exceptions:
=== PROJECT ENVIRONMENT ===
Application:
Version and build:
Operating system:
Required hardware:
Project / database name:
Project organization:
Reel / longplay structure:
Bin / folder structure:
Track / layer logic:
Versioning method:
=== CONFORM AND EDITORIAL ===
Timeline raster:
Timeline frame rate:
Start timecode:
FFOA / LFOA by reel or program:
Source-to-timeline relationship:
Update and replacement method:
Alternate cuts / branches:
Reference movie:
Flattened EDL:
=== FRAMING AND SCALING ===
Working raster:
```

```
Framing chart:
Resize filter:
Input scaling:
Output scaling:
=== OUTPUT (repeat for each intended output) ===
Output name:
Output raster:
Aspect ratio:
Scaling / crop / offset:
Active image area:
Extraction restrictions:
=== COLOR PIPELINE ===
Color-management system and version:
Configuration file and hash:
Source encodings:
Grading / working space:
Input transforms:
Creative look:
Display transforms:
Output transforms:
White point:
Reference display targets:
Graphics treatment:
Required LUT / DCTL / shader package:
Transform order:
=== TEXT, TITLES, AND GRAPHICS ===
Source application and version:
Fonts:
Source files:
Image state / color encoding:
Alpha / matte handling:
Texted / textless ranges:
Rebuild or recomposite instructions:
=== EXCEPTION LIST (repeat for every non-obvious exception) ===
Reel / shot / timecode / frame range:
Affected version:
Exception:
Required source or treatment:
Do not derive from:
=== MASTER INVENTORY (repeat for every master or project archive) ===
Identifier:
Type:
Cut / version:
Texted / textless / language:
2D / 3D:
Frame rate:
Raster / aspect ratio:
Color encoding / display target:
Audio configuration:
Source project:
Archive path:
Manifest / checksum:
Status: [AUTHORITATIVE / DERIVATIVE / HISTORICAL REFERENCE / SUPERSEDED]
Permitted future use:
=== RELINK AND RECONSTRUCTION ===
```

```
Archive root:
Folder map:
Relink base path:
Expected offline media:
External assets:
Reconstruction order:
Render instructions:
Validation reference:
Known application-version issues:
=== DEPENDENCIES (repeat for each dependency) ===
File / plugin / font / application:
Version:
Archive path:
License requirement:
Replacement or fallback:
Owner and access procedure:
=== KNOWN LIMITATIONS ===
Missing material:
Unsupported derivations:
Expired or machine-locked dependencies:
Decisions that cannot be inferred from the project:
=== STORAGE AND INTEGRITY ===
Package manifest:
Checksum format:
Primary / second / offsite copies:
Archive catalog record:
Media generation:
Migration owner:
Next migration review:
=== RESTORE TEST ===
Test date / restored by:
Hardware / software:
Files verified:
Project opened:
Media relinked:
Test render compared:
Result:
Open issues:
Next test:
```

Reference

Glossary

Terms as used in this handbook.

A

AAF: Advanced Authoring Format. An open interchange specification published by the **AMWA** (Advanced Media Workflow Association, formerly the AAF Association), primarily used by Avid for picture and audio turnovers. Binary, supports multi-layer timelines and keyframed speed ramps. Not to be confused with **MXF**, which *is* a SMPTE standard (ST 377-1) and whose structural metadata derives from AAF's data model. See Edit List Generation.

ACES: Academy Color Encoding System. A standardized colour-management framework: one scene-referred interchange encoding (ACES2065-1), working spaces derived from it, defined transforms at each end of the pipeline, and a container. See ACES.

AMF: ACES Metadata File. An XML sidecar that records which ACES transforms a shot or show uses: Input Transform, look transforms, Output Transform, so a pipeline's colour configuration travels with the material instead of living in someone's notes.

AOV: Arbitrary Output Variable. A supplementary render pass (normal maps, z-depth, and so on) stored as an additional channel in an OpenEXR file. See OpenEXR.

ASC CDL: American Society of Cinematographers Color Decision List. A deliberately restrictive color correction interchange format limited to slope, offset, power, and saturation. The ten values transfer reliably **only when sender and receiver agree on the working space**: the CDL defines no colour space of its own, so the same numbers applied to differently encoded images give different results. Agreeing the space is the whole job. The numbers are the easy part. See Plate Pre-Grading.

ASC MHL: Media Hash List. An ASC-published format for recording checksums of copied media so that a transfer can be verified rather than assumed. The practical defence against silent corruption during offloads and archival.

B

BRAW: Blackmagic RAW. Blackmagic's camera format, which replaced CinemaDNG in 2018 and is recorded across the URSA Cine family. Partially demosaiced in-camera, which is why it behaves more like a light RAW than a full sensor dump.

C

Camera RAW: Unprocessed sensor data preserved for creative interpretation later. Inherently scene-referred. Inefficient as a VFX working format because it must be debayered on every access. See Camera RAW.

CBB: Could Be Better. A visual effects shot deemed passable as-is but slated for improvement if schedule permits. See Digital Intermediate.

Colorist: The artist responsible for the final look of the picture, working with the filmmakers to craft it and maintaining continuity across the show. See Responsibilities of the Colorist.

Comp: Composite. A finished visual effects shot, assembled from a plate plus rendered or painted elements. Also the verb for making one.

Concatenation: The collapsing of multiple successive image transforms into a single net operation, avoiding the cumulative softening of repeated resampling. Broken by interposing color or shader nodes. See Concatenated Image Transforms.

Conform: Reconstructing the show's timeline at full quality from camera original and VFX elements, using metadata from the offline edit. See Editorial Turnover for DI.

D

Dailies: Footage processed and delivered for review shortly after it is shot, usually with a viewing LUT applied, for the director, cinematographer and editorial. See Digital Dailies.

DCDM: Digital Cinema Distribution Master. An unencrypted render in DCI-X'Y'Z', almost always carried as **16-bit TIFF**, though note the distinction between container and content: ST 428-1 defines the *code values* as **12-bit**, which the 16-bit TIFF simply pads. The exception is the packed-image form of ST 428-24, which stores those 12-bit values in losslessly compressed JPEG 2000 instead. A DCDM is the source a DCP is made from, but not *only* that. It is also delivered as a mastering element in its own right, and archived. See Theatrical Distribution Masters.

DCP: Digital Cinema Package. The compressed, usually encrypted package delivered to theaters. See DCP.

Data window / display window: In OpenEXR, the region containing image data versus the visible frame area. These may differ, and in practice DI systems handle the difference inconsistently, Resolve, for instance, simply truncates pixels outside the display window. Worth testing on your own setup rather than assuming the overscan survives. See OpenEXR.

Debayer: Demosaic. Reconstructing full-colour pixels from a sensor's Bayer-patterned photosites, where each site records only one of red, green or blue. Debayering happens when Camera RAW is decoded, which is why RAW files carry no fixed resolution or colour until they are processed.

Digital intermediate (DI): The stage where a show is conformed, VFX are integrated, colour grading happens, and every deliverable is mastered. Named for the photochemical intermediate elements it originally replaced. See Where the name comes from.

Display-referred: Image encoding defined in relation to display characteristics, with no numeric relationship to real-world exposure. The state of any image intended for exhibition. ISO 22028-1 calls this same state **output-referred**. The two terms are interchangeable and you will meet both. See Display-Referred Imagery.

DIT: Digital Imaging Technician. Responsible for on-set ingest, backup, and often on-set look creation. See Digital Dailies.

DNx: Avid's codec family, rebranded in 2025 from the DNxHD/DNxHR names to five quality levels (444, HQX, HQ, SQ, LB). Resolution-independent, and all levels now support 8- to 16-bit encoding and an optional alpha channel. See Avid DNx.

Dolby Vision trim: An authored pass in which a colourist sets, per shot or per scene, how the HDR grade should be reduced to lower-capability displays. The trim is carried as dynamic metadata rather than baked into a separate master. See HDR Mastering.

DSM: Digital Source Master. The graded master a show's deliverables are derived from. DCI deliberately **does not specify** its format, colour space or encoding. It is defined by what the production and its distributors agree on, and in practice one show may carry several (Rec.709, Rec.2020/HDR10, P3-DCI and Dolby Vision versions are all commonly requested separately). Sometimes skipped in favour of producing a DCDM directly. See DSM.

E

EDL: Edit Decision List. The oldest and simplest interchange format. Universally supported, human readable, but unable to describe variable speed ramps or geometric effects. See EDL.

EOTF / OETF / OOTF: Electro-Optical, Opto-Electronic, and Opto-Optical Transfer Function. The OETF maps scene light to code values (a camera's job). The EOTF maps code values to displayed light (a display's job). The OOTF is the net scene-to-display rendering, including the intended artistic rendering. Conflating the OETF and EOTF of a standard is a common and consequential error.

EXR: OpenEXR. The standard image format for scene-linear and VFX work, storing half-float (IEEE 754 binary16) or full-float pixels, arbitrary channels, and a data window that may differ from the display window. See OpenEXR.

F

FFOA: First Frame of Action. The reference point marking where the program picture begins, used to align head leaders and the 2-pop. Also written **FFOP** (First Frame of Picture). FFOA is the form used in most delivery specifications. See Leaders and Tail-Pop.

Frame padding: Deliberately recording image area outside the intended frame lines to preserve latitude for stabilization and reframing. See Frame Padding.

Framing chart: A photographed chart or generated pixel-accurate reticle communicating intended framing unambiguously. See Framing Charts.

H

Handles: Extra frames beyond the edit at the head and tail of a plate or comp, giving editorial room to roll a cut. See Frame Numbering and Handles.

HLG: Hybrid Log-Gamma. A *relative* HDR transfer function designed for broadcast, where the display maps the signal to its own peak. Contrast with PQ, which is absolute. Both are defined in ITU-R BT.2100.

I

IDT / ODT / RRT: Input Device Transform, Output Device Transform, Reference Rendering Transform: the original ACES 1.0 names. Current ACES documentation says **Input Transform** and **Output Transform**. Since ACES 1.1, the RRT and ODT are concatenated into a single Output Transform rather than applied in sequence. The old names persist in software menus.

IMF application: A constrained profile of the Interoperable Master Format specifying which essence encodings a package may use. App #2E (ST 2067-21) is the common JPEG 2000 profile. App #4 is for archival cinema masters, App #5 carries ACES, and RDD 45 carries ProRes. See IMF.

L

LFOA: Last Frame of Action. The reference point marking where the program picture ends, used to align tail leaders and the tail pop. Also written **LFOP** (Last Frame of Picture). LFOA is the form used in most delivery specifications. See Leaders and Tail-Pop.

LMT: Look Modification Transform. In ACES, a look applied in the scene-referred domain, before the Output Transform, as distinct from a display-referred grade. The right place for a show LUT-style creative look inside an ACES pipeline.

Longplay: The film assembled as a single continuous sequence rather than divided into reels. See Working in Reels.

LUT: Look Up Table. A pre-computed color transformation, either technical (standards-based) or creative (subjective, s-shaped, scene-to-display). A **1D LUT** maps each channel independently and so cannot remap a gamut or mix channels. A **3D LUT** samples the RGB cube and can. The Academy/ASC **Common LUT Format (CLF)** is the interchange format that carries either, along with ASC CDL operations, in a single documented XML container. See Look Up Tables.

M

Matte: A channel isolating an element of a composite, delivered to the DI so the colorist can grade composite regions selectively without re-rotoscoping. See Working with Visual Effects Mattes.

MaxCLL / MaxFALL: Maximum Content Light Level and Maximum Frame-Average Light Level. Static HDR metadata describing the *content*, specifically the brightest pixel and brightest frame average. This contrasts with ST 2086, which describes the mastering display. See HDR Mastering.

MXF: Material Exchange Format. A SMPTE-standardized container (ST 377-1) that wraps essence with structural metadata. Used for IMF and DCP track files and for Avid's OP-Atom media. Its metadata model derives from AAF's.

N

Nit: Informal name for the candela per square metre (cd/m²), the unit of luminance. SDR mastering references 100 nits. HDR grading displays are commonly 1,000 or 4,000.

O

OCIO: OpenColorIO. An open-source colour-management library, now an Academy Software Foundation project, that lets applications share one colour configuration. The practical mechanism by which a facility applies a consistent pipeline across different software. See Color Management and OpenColorIO.

Opticals: Effects reproducible during conform: resizes, repositions, stabilization, flips, flops, simple timewarps, dissolves, locked-off split screens. Named for the optical printing required in the negative-assembly era. See Opticals List.

OTIO: OpenTimelineIO. An open-source, vendor-neutral timeline interchange format and API, originated at Pixar and now an Academy Software Foundation project. Used mostly as programmatic glue for translating and validating edit lists rather than as a native turnover format. See OTIO.

OTT: Over-the-top. Streaming distribution outside traditional broadcast and cable.

Output Transform: In ACES, the transform that converts graded ACES data to a specific display encoding. Replaced the separate RRT + ODT pair as a user-facing concept in ACES 1.0, and became a genuinely concatenated transform in ACES 1.1.

P

Plate: Camera original footage delivered to visual effects as the basis for a shot. See Visual Effects Plate Pulls.

PQ: Perceptual Quantization, standardized as SMPTE ST 2084. An *absolute* HDR transfer function: a code value maps to a specific luminance in cd/m², so the creative intent is pinned to measurable light. Contrast with HLG. See HDR Mastering.

Previz: Previsualization. CGI animatics used to plan complex shots and sequences before production. See Pre-Production.

ProRes: Apple's professional codec family, published as SMPTE RDD 36, with MXF carriage in RDD 44 and an IMF application in RDD 45. The 4444 variants sustain multiple generations without meaningful degradation. See Apple ProRes.

ProRes RAW: Apple's camera-acquisition RAW format (2018). Despite the shared name it is not a member of the ProRes intermediate family and is not applicable to transcoded VFX plate pulls.

Proxy: A lower-resolution or lower-bandwidth copy of media, used so that editorial and review can work without the storage and throughput the full-resolution material demands. Distinct from a *deliverable*: a proxy is never what ships.

R

Reel: A ~20-minute division of a feature, inherited from the physical limits of print stock and retained because it parallelizes editorial and grading labor. See Working in Reels.

RGC: Reference Gamut Compression. An ACES transform that maps out-of-gamut values, which a wide encoding like AP0 readily produces from real cameras: back into a workable range, so that downstream operations behave predictably instead of producing artefacts around saturated highlights.

Roundtrip: Sending material out of a system and bringing it back, then verifying that nothing changed but the intended work. The standard proof that a pipeline is correctly configured. See zero net change.

S

Scene-linear: Linear encoding of relative light values in floating point, with middle gray at 0.18 and each stop doubling. Values may exceed 1.0. See Scene-Linear.

Scene-referred: Image encoding defined in relation to the scene's relative exposure values, retaining the camera's native dynamic range. See Scene-Referred Imagery.

Show LUT: The single scene-to-display transformation chosen for a project. Drives on-set monitoring and sits at the end of the grading pipeline. The digital equivalent of print stock. See The Show LUT.

Slap comp: A rough temporary composite made in editorial to convey a story beat before the real shot exists. See Creative Editorial.

Slate frame: An informational frame prepended to a delivered comp carrying shot, version, vendor, artist, and framing details. See Frame Numbering and Handles.

ST 2086: SMPTE's static HDR metadata standard describing the **mastering display**, including its primaries, white point, and minimum and maximum luminance. Frequently confused with MaxCLL/MaxFALL, which describe the content instead. See HDR Mastering.

T

Tape name: The unique source identifier that, with timecode, links a timeline clip back to camera original media. Without it, conform is guesswork. See Edit List Generation.

TD: Technical Director. At a VFX studio, the person responsible for defining and vetting the imaging pipeline. See Visual Effects Quality Control.

Turnover: The delivery of edit lists, media, and references from editorial to the DI or to VFX. See Editorial Turnover for DI.

Two-pop: A single-frame audio and picture sync reference two seconds before first frame of picture. See Leaders and Tail-Pop.

U

USD: Universal Scene Description. Pixar's open-source framework for describing, composing and interchanging 3D scenes, now an Academy Software Foundation project and the emerging standard for moving scene data between applications.

W

WIP: Work in progress. An unfinished VFX shot delivered so grading can begin in parallel. See Digital Intermediate.

Z

Zero net change: The principle that non-VFX regions of a plate must come back through the compositing pipeline unaltered, with no unintended colour space or encoding shift. Where the plate and the comp share a format and encoding this is literally bit-identical and should be verified by differencing. Where the pipeline crosses formats, such as a DPX plate returning as EXR, the requirement is that the difference be zero within the precision of the encoding, not that the bits match. See Visual Effects Production.

Supplemental Reading & Resources

The documents, standards, books, and tools cited or referenced across this handbook, gathered in one place. Where a work is freely available, the link goes to the PDF or hosting page. Where it is sold, the link goes to the publisher or vendor page that offers it.

!!! note “Sourcing” Claims in the chapters are checked against the primary sources below and cited inline.

Standards

SMPTE

As of June 2026, all published SMPTE Standards, Recommended Practices, Engineering Guidelines, and Registered Disclosure Documents are **freely downloadable** from the SMPTE standards library. The ACES set also has a curated landing page. Each citation below links to its document page:

- **ST 2065-1**: Academy Color Encoding Specification (ACES2065-1), the AP0 interchange encoding
- **ST 2065-4**: ACES image container file layout (the constrained OpenEXR container)
- **ST 2065-5**: mapping ACES image sequences into the MXF Generic Container
- **ACESclip**: clip-level ACES metadata is Academy **TB-2014-009**, not a SMPTE standard. In current practice it is superseded by AMF
- **ST 2067-21**: IMF Application #2E (the streaming/versioning application)
- **ST 2098-2**: Immersive Audio Bitstream (IAB) specification
- **ST 2067-201**: IMF Immersive Audio Bitstream Level 0 Plug-in (carriage of ST 2098-2 essence in IMF)
- **ST 2084**: PQ (Perceptual Quantization) EOTF
- **ST 2086**: Mastering-display color-volume metadata (static)
- **ST 2094**: Dynamic metadata for color-volume transform (Dolby Vision, HDR10+): -1, -10, -40
- **ST 428-1**: D-Cinema Distribution Master (DCDM) X′Y′Z′ image encoding
- **ST 428-24**: DCDM Packed Image (pDCDM): lossless JPEG 2000, 12-bit-packed DCDM exchange
- **ST 379-1** and **ST 422**: MXF generic container and JPEG 2000 mapping
- **ST 2019-1**: VC-3 (Avid DNxHD) bitstream
- **RP 431-2**: D-Cinema projection reference (P3 / 48 cd/m²)
- **SMPTE 195-1993** and **PH22.106** (1957, 1971): anamorphic (“CinemaScope”) aperture and aspect-ratio history

ITU-R (freely downloadable at itu.int)

- **BT.709**: HDTV primaries, white point, and camera OETF
- **BT.1886**: reference EOTF for flat-panel SDR displays
- **BT.2020**: UHDTV wide-gamut primaries
- **BT.2100**: HDR television (PQ and HLG)
- **Report BT.2390**: high dynamic range television (the OETF/EOTF/OOTF discussion)

Academy / ACES

- **ACES documentation hub**: encodings (ACES2065-1, ACEScg, ACEScc, ACEScct), the Input / Output / Look Transforms, AMF, CLF, Reference Gamut Compression, and the white point
- **aces-core**: the reference transforms (CTL). Note the repository was renamed from `ampas/aces-dev`, and the `documents/` tree was removed at v1.3
- **Historical ACES 1.x technical documents** (superseded, with current versions at docs.acescentral.com): Technical Bulletins **TB-2014-004**, **TB-2014-009**, **TB-2014-010**, **TB-2014-012**, **TB-2018-001**, and Specifications **S-2013-001**, **S-2014-003**, **S-2014-004**
- **aces (ASWF)**: the current Academy Software Foundation ACES project
- **ACES 2.0**: the unified Output Transform (replacing the RRT + ODT split), chroma compression, gamut compression, and white limiting

Books

- **Cinematic Color: From Your Monitor to the Big Screen**: Jeremy Selan, VES Technology Committee white paper (2012). Still the best short primer, but predates ACES 1.0 (free PDF)
- **Color and Mastering for Digital Cinema**: Glenn Kennel (Routledge/Focal Press, 2006). The DCI X′Y′Z′ encoding and the 52.37 normalization rationale. Also a free borrow at the Internet Archive
- **Colour: Sense & Measurement**: Richard Kirk (FilmLight, revised edition, March 2026). The T-CAM rationale and the “simplest model” argument. FilmLight also hosts a free online/PDF edition

Vendor documentation

- **Blackmagic Design: DaVinci Resolve manuals & training**: the *DaVinci Resolve Reference Manual* (Ch. on Data Levels, Color Management, and ACES), *The Colorist Guide to DaVinci Resolve*, and the *DaVinci Resolve 17: Wide Gamut Intermediate* white paper (DWG/Intermediate primaries and matrices)
- **FilmLight: Truelight / TCS / T-CAM**: the Truelight Color Spaces, T-CAM display transform, downloadable color-space files, and OCIO config
- **Dolby: Dolby Vision (professional)**: the *Dolby Vision Color Grading Best Practices Guide* and *Dolby Vision Professional Tools User Manual* (provided to licensed facilities). Trim-pass and metadata-level workflow

Tools & platforms

- **OpenColorIO**: the color-management engine that delivers ACES to VFX applications
- **OpenEXR**: the scene-referred image format underlying the ACES container
- **VFX Reference Platform**: the annual specification pinning ACES, OCIO, OpenEXR, and the rest of the stack
- **Filmbox** (Video Village): a Kodak Vision3 film-print-emulation display transform
- **Lattice** (Video Village): LUT authoring and conversion
- **ColourSpace** (Light Illusion): display calibration and profiling (successor to LightSpace CMS)
- **Codex Backbone**: on-set media management, discontinued and no longer in Codex’s current lineup
- **Apple ProRes white paper**: the ProRes family and data rates
- **DCI: Digital Cinema Initiatives**: the digital-cinema specifications and compliance test plans

- **Technicolor Pulse**: remote review and dailies, now defunct with its service and domain retired
- **Academy/ASC Common LUT Format specification** (S-2014-006): the freely available normative statement of the ASC CDL transfer functions (`ASC_CDL` ProcessNode). The primary document, Pines and Reisner's *ASC CDL Transfer Functions and Interchange Syntax*, Release 1.2, is not publicly hosted. Request it from `asc-cdl@theasc.com`. For a plain-language walkthrough of SOP/Sat versus lift/gamma/gain, see Pomfort's in-depth look at ASC-CDL based color controls

Digital cinema delivery (DCP)

- **ISDCF**: the Inter-Society Digital Cinema Forum: the Digital Cinema Naming Convention and the Open and Closed Captions registry (OCAP / CCAP)
- **Deluxe Digital Cinema:** ***Specifications for Digital Cinema Source and DCP Content Delivery*** (v5.11, 14 March 2022): a practical DCP delivery spec covering Interop vs SMPTE, audio channel layouts, subtitle and caption rules, accessibility tracks, and delivery media. Some chapters cite the earlier v4.08/2019 "DTDC" edition for historical points
- **Video Village DCP Inspector**: deep DCP validation of the CPL, PKL, ASSETMAP, hashes, reels, frame rate, audio layout, and subtitle and caption references. It can repair many common CPL errors before a theater QC or return to the authoring vendor
- **Cinematiq:** ***Making a DCP: Accessibility***: OCAP/CCAP, SDH, and described-audio in practice
- **Festival DCP technical specifications.** Verified July 2026. Linked to each festival's *evergreen* page rather than a year-stamped PDF, because the PDF paths change every edition (Venice's file name is not even predictable year to year):
  - **Berlinale**: the most complete public spec of any major festival and the clearest pro-SMPTE statement
  - **Venice**: names SMPTE **Bv2.1** as the preferred profile. The year in the URL changes each edition
  - **Cannes: Marché du Film**, where Cannes actually publishes a spec. The Official Selection publishes none
  - **Locarno**: compulsory DKDM handover through a partner lab
  - **Sundance**: the strictest accessibility requirements in this group. Note its detailed technical PDF has not been refreshed since 2023
  - **IFFR** and **Hot Docs**
  - **TIFF, SXSW, Tribeca and Cannes' Official Selection publish no DCP technical specification**. They issue requirements privately after invitation. Do not rely on third-party summaries.

Displays, HDR, and theatrical

- **HDR by Barco** and the **Barco LS4K** Lightsteering projector
- **RTINGS**: per-display Rec.2020 / DCI-P3 gamut-coverage measurements
- **GDC: DCI-certified HDR for cinema**: direct-view LED HDR certification
- **Samsung Onyx** and **LG Miraclass**: direct-view LED cinema displays
- Wide-gamut panel announcements: **FlatpanelsHD**, **TechRadar**

Articles & discussion

- **Resolve Color Management vs ACES, Which Should You Choose?** (Frame.io Insider)
- **ACES vs RCM** (Mixing Light)

- **RCM vs ACEScct discussion** (ACESCentral)
- **FilmLight Baselight v7 Truelight Colour Space improvements** (Newsshooter)
- CinemaScope aspect-ratio history: the **American WideScreen Museum**

Change List

Version 2.0: 2026

Version 2.0 converts the 2017 handbook into a maintained wiki and updates its technical guidance through July 2026. It adds current mastering, color-management, distribution, and archival workflows while preserving the handbook's focus on practical decisions for independent productions.

New chapters

- The Academy Color Encoding System (ACES)
- Color Management and OpenColorIO
- Application-Native Color Management
- Choosing a Color Management Approach
- HDR Mastering
- IMF: The Interoperable Master Format
- DCP: The Digital Cinema Package
- Distribution Deliverables for Independent Film
- Where to Start
- Production Checklists
- Technical Specifications
- Reusable Production Templates

New and expanded sections

- Scene-referred and display-referred image states, camera log, and bit-depth requirements
- ACES 2.0, ACEScg/ACEScct, Output Transforms, gamut compression, and ACES containers
- Resolve Color Management, DaVinci Wide Gamut/Intermediate, FilmLight TCS, and T-CAM
- OCIO configuration, version control, vendor distribution, and round-trip testing
- HDR10, Dolby Vision metadata and trims, HDR/SDR mastering strategy, and HDR VFX review
- IMF applications, supplemental packages, Photon validation, and component-based versioning
- SMPTE DCP, Bv2.1, accessibility tracks, KDM delivery, cinema-server QC, and festival requirements
- Independent distribution masters, audio stems and M&E, loudness targets, captioning, legal deliverables, and trailer requirements
- VFX bidding, count sheets, review and approval, versioning, trailer work, and turnover lifecycle
- Confidence packages, frame-rate planning, pDCDM interchange, and modern mastering paths
- Archival masters, LTO, checksums, cloud storage, media migration, and project reconstruction

General updates

- Replaced 2017 camera examples with current ARRI, Sony, RED, and Blackmagic systems.
- Updated Apple ProRes, Avid DNx, OpenEXR, DPX, H.264/H.265, AV1, and JPEG 2000 guidance.
- Updated HD, UHD, 2K, 4K, Flat, Scope, aspect-ratio, and frame-rate guidance.
- Updated Rec.709, BT.2020, DCI X′Y′Z′, camera gamut, transfer-function, and white-point material.
- Added current SMPTE, ITU, Academy, DCI, ISDCF, AMWA, and VFX Reference Platform references.
- Rebuilt figures from the original source artwork and added current workflow diagrams.

- Corrected terminology, obsolete product references, dead links, and technical inaccuracies.
- Reorganized the handbook into shorter, searchable chapters with consolidated cross-references.
- Tightened the prose throughout for faster scanning and more actionable guidance.
- Reorganized navigation around the production lifecycle and role-based entry points.
- Added page metadata, canonical URLs, structured data, `llms.txt`, and a complete Markdown corpus.
- Added automatically rebuilt EPUB and PDF editions for offline use.
- Made handbook figures expandable in a full-screen, keyboard-accessible lightbox.

Earlier releases

Version	Notes
1.0.1	Corrected typos and removed an erroneous blank page.
1.0.0	First public release.
0.9.4	Added image-processing workflows and imaging-science refinements.
0.9.3	Revised graphics.
0.9.2	Revised figures, added the Cineon transfer function, and corrected layout.
0.9.1	Added display and camera color-space figures.
0.9.0	Initial limited release.
0.1.0	Initial draft.

---

1. Poynton, *Digital Video and HD*, 2nd ed., ch. 32–33 covers the 1000/1001 factor and its NTSC origin. SMPTE ST 170M-2004 §11 defines the rates. **SMPTE ST 12-1:2014** §5.1.2 defines NTSC time as `t(NTSC) = 1.001 × t(REAL)`. Section 7.1.2 gives the ~86 frames (3.6 s) per hour figure. Samples per frame at 48 kHz are tabulated in **SMPTE ST 2101:2015**. For digital cinema, **SMPTE ST 429-2:2023** §8.2 fixes the legal edit rates at 24/1, 25/1, 30/1, 48/1, 50/1 and 60/1. **DCI DCSS** v1.3 §3.1.4.2 requires 24.000 Hz. **SMPTE RP 428-4:2010** §5.1.3 states that digital cinema "does not support fractional frame rates such as 23.976 fps … 2002 or 4004 samples per edit unit are not allowed."

   Drop-frame timecode does not exist at 23.976. It is a 30-frame concept. ST 12-1:2014 §7.2 says there is no counting mode "such as drop frame … applicable to 24-frame counting," and Table 3 leaves the drop-frame flag bit unassigned for 24-frame systems. Several widely-read industry references get this wrong and describe 23.976 as a "drop frame" rate.↩

2. Poynton, *Digital Video and HD*, 2nd ed., ch. 32–33 covers the 1000/1001 factor and its NTSC origin. SMPTE ST 170M-2004 §11 defines the rates. **SMPTE ST 12-1:2014** §5.1.2 defines NTSC time as `t(NTSC) = 1.001 × t(REAL)`. Section 7.1.2 gives the ~86 frames (3.6 s) per hour figure. Samples per frame at 48 kHz are tabulated in **SMPTE ST 2101:2015**. For digital cinema, **SMPTE ST 429-2:2023** §8.2 fixes the legal edit rates at 24/1, 25/1, 30/1, 48/1, 50/1 and 60/1. **DCI DCSS** v1.3 §3.1.4.2 requires 24.000 Hz. **SMPTE RP 428-4:2010** §5.1.3 states that digital cinema "does not support fractional frame rates such as 23.976 fps … 2002 or 4004 samples per edit unit are not allowed."

Drop-frame timecode does not exist at 23.976. It is a 30-frame concept. ST 12-1:2014 §7.2 says there is no counting mode “such as drop frame … applicable to 24-frame counting,” and Table 3 leaves the drop-frame flag bit unassigned for 24-frame systems. Several widely-read industry references get this wrong and describe 23.976 as a “drop frame” rate.↩

3. *VES Handbook of Visual Effects*, p. 509: “generally a minimum of 8 frames at the head and tail of each element or plate.” Vendor delivery specifications commonly set 32-frame plate handles against 8-frame comp handles. See Frame Numbering and Handles.↩

4. Joshua Pines and David Reisner, *ASC Color Decision List (ASC CDL) Transfer Functions and Interchange Syntax*, ASC-CDL Release 1.2, ASC Technology Committee, Digital Intermediate Subcommittee, 4 May 2009. It is not publicly hosted. Obtain it from the ASC by e-mailing `asc-cdl@theasc.com`, which auto-responds with terms and instructions (see the ASC MITC Progress Report 2022, *SMPTE Motion Imaging Journal*, September 2022, p. 47). The transfer functions are normatively restated in the freely available Academy/ASC Common LUT Format specification (S-2014-006), `ASC_CDL` ProcessNode.↩

5. Sources include Blackmagic’s *DaVinci Resolve 21 Reference Manual*, *The Colorist Guide to DaVinci Resolve 20*, and DWG/Intermediate white paper, Kirk’s *Colour: Sense & Measurement*, and FilmLight’s TCS documentation. See Supplemental Reading & Resources.↩

6. https://www.dcimovies.com/↩

7. Hence X’Y’Z’ (pronounced X-prime, Y-prime, Z-prime), prime implying a gamma, rather than XYZ.↩

8. SMPTE ST 428-1, Annex B.1, which works the calculation and reports the same three code values together with the 135-code-value headroom figure. The D55 / D61 / D65 rationale is the standard’s own, from footnote 4 to the transfer-function equation. The 48 cd/m^2 reference luminance it preserves is specified in SMPTE ST 431-1, Table 5.2.↩

9. ARRI ALEXA ProRes, RED V-RAPTOR, Blackmagic URSA Cine, a variety of external third-party recorders, and more…↩

10. Avid, *Avid DNx naming scheme and data rates* and *DNxHR Codec Bandwidth Specifications*. Rebrand announced at IBC 2025.↩

11. Avid, *Avid DNx naming scheme and data rates* and *DNxHR Codec Bandwidth Specifications*. Rebrand announced at IBC 2025.↩

12. JPEG, PNG, Targa for example.↩

13. Dimensions and derivations per SMPTE PH22.106-1957, PH22.106-1971, and **SMPTE ST 195-2000** (*Motion-Picture Prints: Projectable Image Area*, reaffirmed 2012. The revision was proposed in the June 1992 *SMPTE Journal* and first published August 1993). Overview at the American WideScreen Museum.↩

14. Not to imply the crew are “objects”.↩

15. Dropbox, Version history overview and data retention policy. Dropbox Business did once offer unlimited version history. That changed in early 2017.↩

16. Some nameless vendors have been guilty of delivering entire feature films without version numbers, leaving the conform editors to rely solely on the delivery date to know whether they are using the most up to date shot.↩

17. Expected count-sheet (or "turnover sheet" / "lineup sheet") content follows common feature practice: see, for example, Evan Schiff, *Feature Turnover Guide: VFX* and Park Road Post's *Guide to VFX Handovers*.↩

18. Foundry, Nuke documentation: the Transform node's controls are translate/rotate/scale/skew/center/invert/filter/clamp/black-outside/motionblur, with no colour-space knob. And "when using filters that employ sharpening, such as Rifman and Lanczos, you may see a haloing effect. Some nodes, such as Transform and Tracker, include a **clamp** control to correct this problem."↩

19. Adobe After Effects is one common example.↩

20. Sports mode, anyone?↩

21. There are dozens of you!↩

22. Under standard human conditions.↩

23. Current-generation bodies and formats as of 2025–26.↩

24. 0.0–1.0 representing a percentage of opacity. Log mattes produce undesirable results or require pre-conversion to linear.↩

25. The ACES documentation, the aces-aswf/aces repository (formerly ampas/aces-dev), and the SMPTE ACES standards: ST 2065-1 through ST 2065-5, normative and, since 2025, freely downloadable from the SMPTE library.↩

26. Academy TB-2014-004 (informative notes on ST 2065-1) and S-2014-004, published as ACEScg.↩

27. Academy TB-2014-004, §4.2.3. Encoding specifications: ACES2065-1 (ST 2065-1) · ACEScg (S-2014-004) · ACEScc (S-2014-003) · ACEScct (S-2016-001) · ACESproxy (S-2013-001). Component naming follows Academy TB-2014-012, except **ACEScct**, which postdates it and was defined in S-2016-001.↩

28. SMPTE ST 2065-4:2023, *ACES Image Container File Layout*, §6.5.3 and Table 5 (required attributes). The 2023 edition revises ST 2065-4:2013.↩

29. The `aces-aswf/aces` repository carries two non-prerelease `v2.0.0` tags: `v2.0.0+2024.09.05` and `v2.0.0+2025.04.04`, so "released in 2025" and "released in 2024" are both defensible depending on whether you mean the reference implementation or the public launch.↩

30. *DaVinci Resolve 21 Reference Manual* (Blackmagic Design, July 2026), Ch. 9, pp. 225–290. *The Colorist Guide to DaVinci Resolve 20* (Blackmagic Design, 2025), Lesson 4.↩

31. *DaVinci Resolve 21 Reference Manual*, Ch. 9, covers the two color-science modes, the Output DRT tone-mapping options (incl. *None*), Resolve's ACES signal flow, working space, and AMF 2.0 support.↩

32. *DaVinci Resolve 21 Reference Manual*, Ch. 9, pp. 241–243. *DaVinci Resolve 17: Wide Gamut Intermediate* white paper (Blackmagic Design, 2021), including the corrected v1.1 green-x coordinate.↩

33. Richard Kirk, *Colour: Sense & Measurement* (FilmLight, 2022): the T-CAM rationale (Ch. 5–6) and "simplest model" argument. FilmLight hosts a free PDF. The E-Gamut/T-Log/T-CAM-v2 product names and the portability specifics below come from FilmLight's own materials rather than the book.↩

34. FilmLight: Colour Management / Truelight and Baselight v7 Truelight Colour Space improvements.↩

35. *Resolve Color Management vs ACES, Which Should You Choose?* (Frame.io Insider, 2024). See also Mixing Light and the ACESCentral RCM-vs-ACEScct discussion. Both RCM and ACES are scene-referred systems solving the same problem.↩

36. *Resolve Color Management vs ACES, Which Should You Choose?* (Frame.io Insider, 2024). See also Mixing Light and the ACESCentral RCM-vs-ACEScct discussion. Both RCM and ACES are scene-referred systems solving the same problem.↩

37. FilmLight: Colour Management / Truelight and Baselight v7 Truelight Colour Space improvements.↩

38. ITU-R BT.2100-3 (February 2025): Table 2 (system colorimetry), Table 3 (reference viewing environment, whose Note 3c carries the small-area-highlight qualifier), and Table 9 (10- and 12-bit integer representation).↩

39. Dolby *Vision Color Grading Best Practices Guide* v4.0 and the *Dolby Vision Professional Tools User Manual* v4.0.0 cover trim workflow and the metadata levels in detail.↩

40. Netflix Partner Help Center, *Dolby Vision HDR Mastering Guidelines*. Sky's clause is in its *Technical Specification for Delivery of Content to Sky*. Disney's in its *Dolby Vision Best Practices Guide*. The SDR-check requirement in Dolby's *Color Grading Best Practices Guide* (v4.x). No standards body publishes a position on derived versus discrete SDR: it is a distributor-by-distributor question.↩

41. ITU-R BT.2020 (adopted by BT.2100 for HDR) defines the wide-gamut primaries as monochromatic single-wavelength points: red 630 nm, green 532 nm, blue 467 nm.↩

42. HDR by Barco and the LS4K projector. Launched at CinemaCon, April 2024. The contrast and gamut figures are Barco's.↩

43. DCI direct-view display specification and certification. The ~300-nit HDR figure and the product details are from manufacturer and integrator sources (GDC Technology, Samsung Onyx, LG Miraclass).↩

44. ATSC A/85:2026-07, *Techniques for Establishing and Maintaining Audio Loudness for Digital Television*, approved 8 July 2026. It is the first comprehensive revision since 2013, adding **Annex L** (streaming) and **Annex M** (a one-page quick reference). §6 carries the −24 LKFS / ±2 dB / −2 dBTP figures. §5.2 the anchor-element rule. §5.1.1 the meter guidance. Note that A/85 is a *Recommended Practice*: everything in it is a "should". What is legally binding is the FCC rule, and **47 CFR §73.8000(a)(4) still incorporates A/85:2013**, not the 2026 edition. ATSC also revised A/85 in 2021 without the FCC updating its reference. The Commission has an open proceeding (MB Docket 25-72) that asks, among other things, about streaming loudness. So the current *document* and the enforceable *rule* are different editions.↩

45. DCP packaging is defined by **ST 429-7** (CPL), **ST 429-8** (PKL), and **ST 429-9** (ASSETMAP). IMF later reused the model. See SMPTE OV 2067-0 and the Library of Congress format description.↩

46. DTDC, *Recommended Guidelines for Digital Cinema Source and DCP Content Delivery* v4.08 (2019). DCP is defined by SMPTE **ST 428**, **ST 429**, and the DCI Digital Cinema System Specification.↩

47. The DCI specification caps 2K and 4K JPEG 2000 at **250 Mbit/s**. DCSS §4.3.1 and ST 429-2 §8.3 define the rasters. DCI's *HDR D-Cinema Addendum* v1.2.1 (2024) raises the ceiling to 450 Mbit/s for 2D 4K at 48/60 fps.↩

48. Reference monitoring level (85 dB SPL per channel) and the cinema **X-curve** are standardized in SMPTE ST 202 / ISO 2969. DCP audio channel layouts (5.1, 7.1, and the HI/VI assignments) are covered in the DTDC delivery guidelines. Dolby Stereo Lt/Rt is the historical matrix-encoded two-channel theatrical format.↩

49. Published festival specifications, verified July 2026: Venice · Berlinale · Marché du Film · Locarno · Sundance. Specifications change yearly: confirm the current edition before mastering. Note that Sundance's detailed technical PDF has not been refreshed since 2023. **TIFF, Tribeca, Cannes Official Selection and SXSW publish no DCP specification at all**. requirements are issued privately after invitation. Anything attributed to those festivals here comes from a specification supplied to an accepted project, not from a public document.↩

50. Published festival specifications, verified July 2026: Venice · Berlinale · Marché du Film · Locarno · Sundance. Specifications change yearly: confirm the current edition before mastering. Note that Sundance's detailed technical PDF has not been refreshed since 2023. **TIFF, Tribeca, Cannes Official Selection and SXSW publish no DCP specification at all**. requirements are issued privately after invitation. Anything attributed to those festivals here comes from a specification supplied to an accepted project, not from a public document.↩

51. Festival accessibility requirements, verified July 2026: Sundance · Hot Docs · IFFR. The SXSW figures come from the specification it sends accepted projects rather than a public document. SXSW publishes only that a DCP is required. TIFF publishes nothing at all and issues requirements privately, so treat any second-hand statement of its caption policy with caution.↩

52. ISDCF, Open and Closed Captions and the Digital Cinema Naming Convention.↩

53. DTDC, *Recommended Guidelines for Digital Cinema Source and DCP Content Delivery* v4.08 (2019). DCP is defined by SMPTE **ST 428**, **ST 429**, and the DCI Digital Cinema System Specification.↩

54. Published festival specifications, verified July 2026: Venice · Berlinale · Marché du Film · Locarno · Sundance. Specifications change yearly: confirm the current edition before mastering. Note that Sundance's detailed technical PDF has not been refreshed since 2023. **TIFF, Tribeca, Cannes Official Selection and SXSW publish no DCP specification at all**. requirements are issued privately after invitation. Anything attributed to those festivals here comes from a specification supplied to an accepted project, not from a public document.↩

55. ISDCF Doc 3 recommends ext2/ext3 with **128-byte inodes**. Modern Linux defaults to 256-byte inodes, which some older servers reject. DCI does not specify a delivery-drive file system.↩

56. Published festival specifications, verified July 2026: Venice · Berlinale · Marché du Film · Locarno · Sundance. Specifications change yearly: confirm the current edition before mastering. Note that Sundance's detailed technical PDF has not been refreshed since 2023. **TIFF, Tribeca, Cannes Official Selection and SXSW publish no DCP specification at all**. requirements are issued privately after invitation. Anything attributed to those festivals here comes from a specification supplied to an accepted project, not from a public document.↩

57. ISDCF, Open and Closed Captions and the Digital Cinema Naming Convention.↩

58. Netflix/photon on GitHub, Apache-2.0.↩

59. Netflix Partner Help Center: *Interoperable Master Format (IMF): Overview* and *Post Production Branded Delivery Specifications*. Netflix's delivery specifications are versioned and change over time. Confirm current requirements before delivery.↩

60. **LTO hard cost** = tape stock only (`ceil(TB ÷ 18)` cartridges at ~$105 each). **Archive** and **restore** are a managed-service rate (~$50/TB each way, per 2025–26 rate cards), typically with a ~20 TB minimum or an hourly rate for smaller jobs. In-house instead means buying your own deck, software, and hardware. You keep the tapes, so there is no vault fee. **Cloud** figures use each provider's cold/archive tier. Egress assumes one full download. Prices are from mid-2026. Verify them before budgeting.↩

61. LTO Program, LTO-10 generation overview: 30 TB native (40 TB variant), 400 MB/s native, and "LTO-10 tape drives do not support backwards write or read compatibility."↩