Dynamic Media Facilities: From Vision to Technical Progress

By Brad Gilmer, Willem Vermost, and Félix Poulin

The transition to software-defined media infrastructure is about more than replacing dedicated hardware with software. It presents an opportunity to rethink how media facilities are built, provisioned, operated, and paid for. The industry is starting to transition from IP to IT, and the changes this brings will be profound.

That is the vision behind the Dynamic Media Facility (DMF): an open, interoperable environment in which media functions can be dynamically deployed and orchestrated on shared compute resources. Instead of building infrastructure around fixed endpoints and peak requirements, a DMF can allow resources to be provisioned, reconfigured, and scaled according to the needs of a particular production.

This flexibility can apply across different deployment models. A DMF may operate on-premises, in the cloud, or as a hybrid of the two. The goal is not to prescribe a particular platform or technology, but to give users greater freedom to select the applications, infrastructure, vendors, and deployment models that make sense for their workflows.

Turning that vision into a practical, multivendor ecosystem requires solving a number of technical and business challenges. The Joint Task Force on Dynamic Media Facilities (JT-DMF), a joint initiative of the Advanced Media Workflow Association (AMWA) and the European Broadcasting Union (EBU), is doing that work.

A year after the JT-DMF was officially launched, that work is producing tangible results today.

An Iterative Approach to DMF

One of the most important principles behind JT-DMF is iteration. Dynamic media facilities touch many parts of a media production environment: compute, networking, media exchange, timing, orchestration, monitoring, security, resource management, and more. There are also important business questions around total cost of ownership, licensing, and how dynamically provisioned resources are purchased and used.

To address these challenges, JT-DMF is deliberately taking small, time-bounded steps. Working groups identify a specific problem that needs industry alignment and work on it during a defined sprint. At the end of that period, the group is expected to deliver something. It may not be a finished specification. Instead, it may be a set of requirements, an agreed approach to a problem, reference code, a common best practice, or another piece of work that moves the industry closer to an implementable solution.

We call these deliverables Increments, which make work visible as it develops. Rather than waiting until every issue has been resolved, the industry can see the direction of the work, review it, comment on it, and contribute to it. As the work matures, an Increment may eventually be superseded by a specification or another implementable output.

This iterative model is already delivering results. Six JT-DMF Increments have now been published, addressing flow connection, timing, compute resource management, and business requirements.

Flow Connection

One of the first questions JT-DMF addressed was fundamental: How should media flows be connected between software-based media functions?

The Flow Connection Working Group began by examining API requirements for controlling media exchange layer (MXL) flows within a host. That work resulted in the first JT-DMF Increment, **IN-001: API Requirements—Control of MXL v1.0**.

Importantly, this work did not stop at requirements. IN-001 provided the basis for AMWA BCP-007-03, which specifies how NMOS can be used to manage connections between senders and receivers implementing MXL. This provides an early example of the JT-DMF iterative process working as intended: identify a focused interoperability problem, agree on requirements, publish the work, and use it as the foundation for an implementable industry specification.

The Flow Connection Working Group has since expanded its focus beyond making a single connection to examining the broader lifecycle of a media flow. This second phase is represented by **IN-004: Flow Connection Phase 2 Requirements and Gaps**, currently a work in progress. The group is reviewing additional requirements for flow orchestration and identifying areas that need further standardization or industry agreement.

1545-0279/26©2026SMPTE

An important feedback loop is also developing between specifications and implementations. As vendors build MXL implementations and demonstrate them in real-world environments, new questions and integration issues emerge. Those experiences can then inform the priorities for subsequent JT-DMF sprints. That feedback between implementation and specification is essential if DMF is to address real operational requirements rather than theoretical ones.

Compute Resource Management

Dynamic deployment also raises a basic resource-management question: If a media function can be deployed onto shared infrastructure, how does the system know what resources that function requires?

The Compute Resource Management Working Group is addressing this through **IN-003: Compute Resource Management Manifest and Examples**.

The objective is to provide a vendor-agnostic way of declaring the compute requirements of a media function. A compute resource manifest can describe characteristics such as what goes into and comes out of the media function, the media formats involved, required CPU resources, memory bandwidth, and the reference platform used when determining those requirements.

The work uses a YAML-based manifest, with a schema that can be used to validate the resource description. Examples are also being developed to illustrate how media functions can be profiled and deployed.

An important aspect of this work is what it does *not* attempt to specify. JT-DMF is not prescribing a particular deployment technology. It does not require Kubernetes, containers, or other specific infrastructure choices. Nor does the resource manifest prescribe exactly how vendors must benchmark their media functions.

Instead, the objective is to create a common way to communicate information that an orchestration or resource-management system can use when planning how workloads will run.

Future work for this group is expected to address areas including observability and additional types of media functions and resources.

Timing and Synchronization

Timing is another fundamental challenge. Maintaining precise synchronization across media signals and devices remains a sticking point in IP-based production, and dynamic, software-based environments introduce additional layers of timing complexity. Media may move through multiple workloads running on shared compute infrastructure. Signals entering the environment may also arrive with very different timing characteristics.

The End-to-End Synchronization Working Group is examining the principles required to preserve those timing relationships throughout a DMF production workflow. Its first deliverable, **IN-002: Time and Identity in the Dynamic Media Facility**, addresses timing principles within a DMF media workload: what assumptions can be made inside the compute environment and how timing relationships are maintained.

The group's second Increment, **IN-005: Timing Principles of External Signal Ingress for DMF Media Workloads**, turns to a different part of the problem: what happens when signals enter the DMF environment. Those incoming signals might originate from system cameras, remote cameras, microphones, or other sources. They may have different frequency, phase, synchronization, and timing characteristics.

The group has adopted a useful principle to describe the approach: **conform on ingress**. The idea is to treat incoming signals appropriately as they enter the DMF so downstream media functions can rely on known timing assumptions. Establishing those assumptions at ingress can simplify what happens throughout the rest of the workflow.

There is considerably more work ahead. Areas being considered include timing planes for alignment, timing between media workloads, an ingress registry to record properties and conformance decisions for incoming flows, egress timing principles, upstream signal timing requirements, multi-cluster timing, and status and metrics logging. As with other JT-DMF activities, these topics will be prioritized and addressed incrementally.

Making DMF Work Commercially

Technical interoperability alone will not make dynamic media facilities successful.

The transition from dedicated equipment toward dynamically provisioned media functions also changes how facilities are planned, purchased, licensed, and operated. For that reason, JT-DMF includes a Business Working Group alongside its technical groups.

The fundamental question for this group is straightforward: What business issues need to be addressed to make DMF successful?

Its current work is represented by **IN-006: DMF Business User Stories**. One of the first areas being explored is total cost of ownership. Organizations considering DMF need ways to understand and express the value of moving from existing infrastructure toward a more dynamic model.

That calculation is not necessarily straightforward because the perspective can change depending on the user. A facility owner may evaluate cost and utilization differently from a production or show owner responsible for the budget of a particular program.

Licensing is also emerging as an important area for fur-

ther work. Traditional licensing models may become difficult to apply when software-based media functions can be turned on and off, scaled up and down, and combined differently for each production. Multiple existing licensing approaches can add complexity both for users purchasing capabilities and vendors providing them.

Developing a shared understanding of these business requirements is therefore an important complement to the technical work.

Six Increments and Growing

The current JT-DMF work has produced or initiated six Increments:

- IN-001 – API Requirements – Control of MXL v1.0
- IN-002 – Time and Identity in the Dynamic Media Facility
- IN-003 – Compute Resource Management Manifest and Examples *(work in progress)*
- IN-004 – Flow Connection Phase 2 Requirements and Gaps *(work in progress)*
- IN-005 – Timing Principles of External Signal Ingress for DMF Media Workloads *(work in progress)*
- IN-006 – DMF Business User Stories *(work in progress)*

These documents are not specifications. They are snapshots of convergence: places where a working group has been able to document requirements, principles, or approaches on which participants have reached agreement.

Making these Increments public is intentional. It allows the wider industry to see where JT-DMF is heading and provide feedback while the work is still developing.

What Comes Next

JT-DMF has many areas it has not yet tackled fully, and that is by design.

Security is one obvious example. A dynamic environment built on IT infrastructure must address security throughout the architecture, whether the deployment is on-premises, cloud-based, or hybrid. Work is beginning in this area in coordination with EBU experts.

A common glossary is another requirement. As the industry introduces concepts around media functions, flows, domains, orchestration, resource management, and other elements of DMF, agreeing on terminology becomes increasingly important.

Other future areas include minimum requirements for packaging and onboarding media functions, deployment practices, resiliency models, DMA networking, and MXL-related best practices.

Building DMF Together

The JT-DMF was created because no single broadcaster, vendor, standards organization, or technology provider can define a truly open DMF ecosystem on its own. The effort combines the EBU's community of media organizations with AMWA's community of users, vendors, and systems integrators to address interoperability and business challenges together.

Industry support is growing as well. Nearly 60 organizations have already signed on as DMF Supporters, expressing support for an open and interoperable DMF ecosystem, industry collaboration, open interfaces, multi-vendor interoperability, and freedom of choice across technology platforms and deployment models.

The objective is not to design tomorrow's entire media facility in one step. Instead, we are identifying practical problems, bringing the right people together, reaching agreement where we can, publishing that work, implementing it, and using what we discover to determine the next priority.

One increment at a time, the Dynamic Media Facility is moving from an industry vision toward something the industry can build together.

To find out more about JT-DMF, explore the latest increments and educational resources, or discover how you can contribute to an AMWA working group or on GitHub, visit www.amwa.tv/jtdmf.

About the Authors



Brad Gilmer is president of Gilmer & Associates, Inc., and executive director of both the Video Services Forum (VSF) and the Advanced Media Workflow Association (AMWA), each earning Emmy Awards during his tenure. A SMPTE Fellow, he was the first recipient of the SMPTE Workflow Systems Medal.



Willem Vermost heads Media Production Facilities at VRT. He has over 20 years' broadcasting experience and specializes in IP-based studios, live production, and dynamic media facilities.



Félix Poulin directs CBC/Radio-Canada's Media-over IP Architecture & Lab, formerly led Live IP initiatives at EBU, and co-chairs AMWA's NMOS Steering Committee.

A contribution received from the authors for the 2026 Progress Report.

Copyright © 2026 by SMPTE.

DOI: 10.5594/JMI.2026/TRRJ3929
Date of publication: 15 November 2026