

The SDVI Rally Platform in a Hybrid Environment

Architecting media supply chains that span
on-premises infrastructure and the cloud — a guide
to storage, processing, and deployment strategy

Geoff Stedman, SDVI CMO
2026

TABLE OF CONTENTS

The Case for Hybrid 3

Rally Platform Architecture 3

Hybrid Storage 4

Hybrid Processing 6

Trade-offs Summary 8

Deployment Options 9

Many media organizations are at an inflection point: the operational and economic benefits of cloud-based media supply chains are well established, but the practical reality of existing on-premises investments — storage infrastructure, processing clusters, and long-established workflows — means that a full, immediate move to the cloud is neither feasible nor desirable for many. This paper examines how the SDVI Rally platform is architected to support hybrid deployments, covering the options available for integrating on-premises storage and processing infrastructure into Rally-managed supply chains, the technical trade-offs of each approach, and how organizations can use Rally as a foundation for a more agile and efficient modern operating model.

The Case for Hybrid

The shift toward cloud-based media supply chains has gathered momentum across the industry, driven by the operational and economic advantages the cloud offers: elastic compute capacity that scales with workload demand, pay-as-you-go cost models that reduce idle infrastructure spend, and access to a growing ecosystem of cloud-native media processing and AI services.

For established media organizations, however, the reality is more nuanced. On-premises storage systems — NAS, SAN, tape libraries, and object storage appliances — often represent significant capital investment and still have years of useful life remaining. Content archives may contain petabytes of material that cannot be migrated overnight without disruption. And in some cases, regulatory, security, or contractual requirements may mandate that certain content remain within a controlled on-premises environment for the foreseeable future.

The critical insight is that hybrid is not a compromise or a transitional inconvenience — it is a legitimate, long-term architectural choice for many organizations. The goal is not to prescribe a direction of travel, but to ensure that on-premises and cloud resources can operate as a coherent, well-managed whole. A well-executed hybrid model allows organizations to begin realizing the benefits of cloud-based workflows immediately, against existing content and infrastructure, while executing a deliberate cloud utilization strategy at a pace that fits their operational and financial constraints.

Rally Platform Architecture

SDVI Rally is a cloud-native media supply chain platform. Its control plane — the SaaS layer responsible for resource management, workflow orchestration, asset inventory management, job scheduling, monitoring, and analytics — runs in the cloud and is managed by SDVI. This is an architectural benefit: regardless of where content resides or where processing occurs, the Rally control plane is always available in the cloud.

This architecture is significant because it means that hybrid deployments do not create a split-control model. There is no on-premises Rally instance to maintain, no synchronization between on-premises and cloud control layers, and no reduction in Rally's orchestration capabilities when on-premises infrastructure is involved. The cloud-hosted control plane has full visibility into, and control over, every component of the supply chain — whether those components are cloud-resident or on-premises. Furthermore, the Rally architecture is designed to leverage cloud architectures such as multi-region and multi-availability zones to ensure high availability, resiliency, and business continuity. For more information on the resiliency of Rally, please see www.sdvi.com/resources/maintaining-continuity-and-resilience-in-media-operations/

The hybrid capability in Rally is therefore about the integration reach of the platform — the mechanisms through which cloud-hosted orchestration can connect to, monitor, and drive on-premises storage and processing resources. These mechanisms are the subject of the sections that follow.

Architectural Principle

The Rally SaaS control plane is always cloud-hosted. Hybrid deployments extend Rally's reach into on-premises environments through well-defined integration mechanisms — they do not split or replicate the control layer. This ensures consistent orchestration, monitoring, and governance across all workflow components, regardless of where infrastructure physically resides.

Hybrid Storage

Storage and processing should be treated as independent architectural dimensions in a hybrid deployment. Each can reside on-premises, in the cloud, or in a combination of both, and the decisions made for storage do not necessarily constrain the decisions made for processing. Rally is designed to respect this separation, providing distinct integration mechanisms for each.

Rally supports two integration models for storage — managed and unmanaged — and each model applies equally to on-premises and cloud storage environments. The choice between them has implications for the level of visibility, reliability, and workflow automation available.

Managed Storage

Managed storage is the preferred integration model for content that needs to be fully visible and actionable within Rally. In this model, files stored on the storage system are tracked as inventory items within Rally: Rally is notified when files are added, modified, or removed, and those changes are reflected automatically in the asset inventory. Media analysis runs automatically against content on managed storage, and the full range of Rally supply chain capabilities — including asset registration, QC, transformation, and delivery — can be applied to that content.

On-premises managed storage: Achieving true managed storage integration with on-premises infrastructure requires more than basic compatibility with cloud object storage protocols. Many on-premises storage systems offer partial support, but in practice lack two capabilities that Rally requires for full inventory management: proper event notification (via SNS/SQS messaging) to inform Rally of file changes, and correct role-based access controls to enforce storage permissions. Where these capabilities are not available natively, a storage gateway can bridge the gap.

Spectra Vail, from Spectra Logic, is a software-defined storage gateway that fully satisfies these requirements. Vail presents on-premises storage to Rally as a fully compatible cloud object storage endpoint — including event notifications and role-based access — making on-premises storage functionally equivalent to cloud object storage from Rally's perspective. SDVI remains open to evaluating other storage gateway solutions that can demonstrably meet the same requirements.

Cloud managed storage: For cloud storage, Rally is cloud-agnostic and supports all major object storage providers — including AWS S3, Google Cloud Storage, and Microsoft Azure Blob Storage. Support for European providers such as Scaleway and StackIT is under consideration. This gives organizations the freedom to choose storage providers based on their own strategic, financial, regulatory, or data sovereignty preferences.

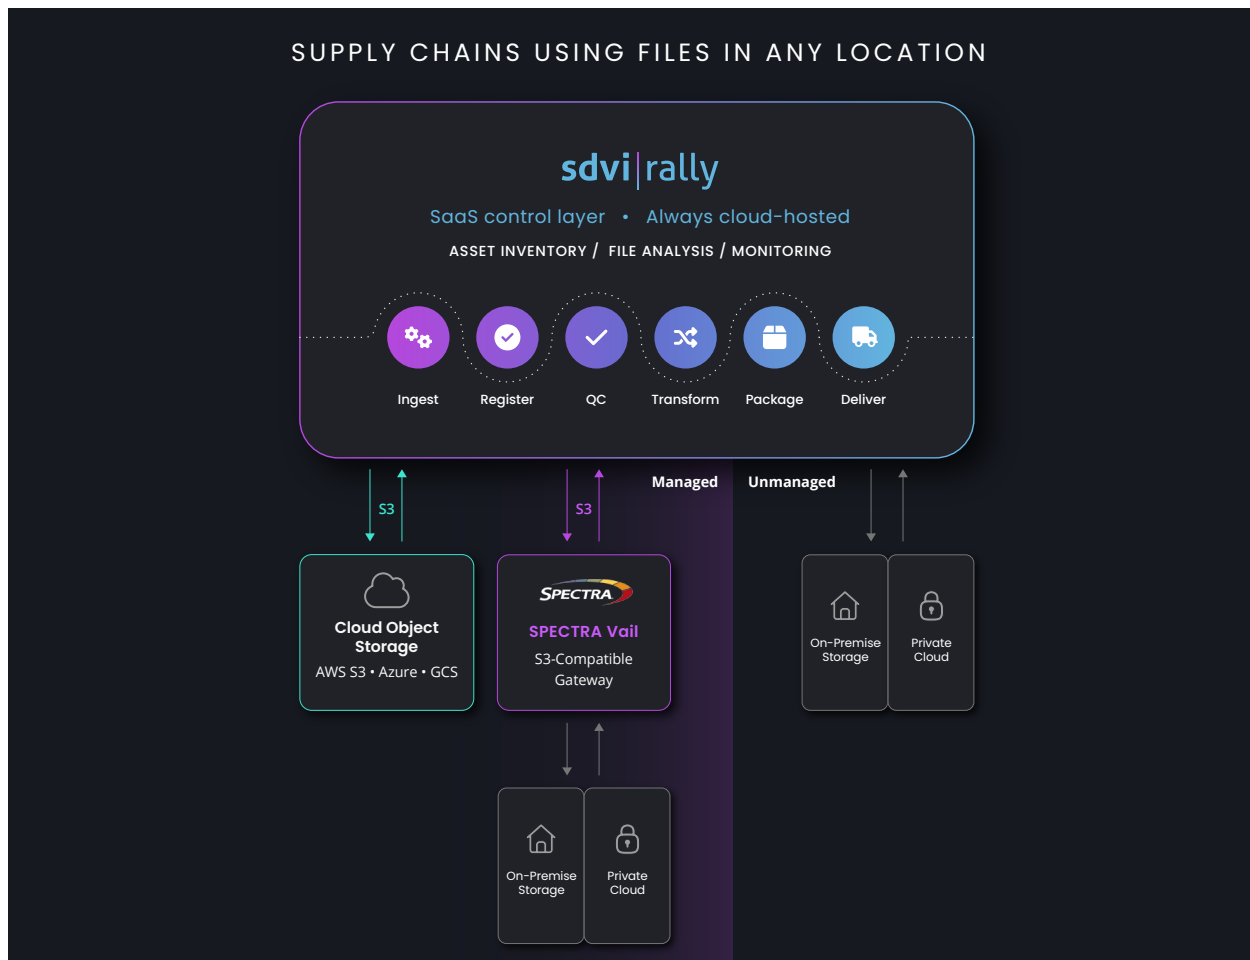
Unmanaged Storage

Rally can also work with storage systems in an unmanaged capacity — meaning content is not tracked within Rally's file inventory. File paths can be stored as metadata attributes on Rally assets, and content can be moved into Rally-managed workflows using SDVI Mover, SFTP, or third-party transfer tools.

On-premises unmanaged storage: On-premises storage systems that do not support the event notification and access control requirements for managed integration can still participate in Rally-managed supply chains in an unmanaged capacity. This provides a practical path to integrating legacy or proprietary storage infrastructure without requiring hardware replacement.

Cloud unmanaged storage: While less common, cloud storage environments can also be integrated in an unmanaged capacity — for example, where access control or organizational policies prevent full managed integration. The same transfer-based mechanisms apply.

Key consideration for unmanaged storage: The key limitation of unmanaged storage — whether on-premises or cloud — is visibility: Rally has no awareness of what happens to files after they are placed on unmanaged storage. Files may be moved, renamed, modified, or deleted without Rally's knowledge. If a file is required for reprocessing at a later stage and is no longer available at its recorded path, the supply chain will require manual intervention to resolve. Organizations using unmanaged storage should implement appropriate content governance practices to mitigate this risk.



Hybrid Processing

Processing — the execution of compute tasks such as transcoding, quality control, audio normalization, and AI-based media analysis — is the second dimension of a hybrid deployment. As with storage, on-premises and cloud processing can coexist within a single Rally-managed supply chain, and the integration model chosen has direct implications for the level of automation, scalability, and operational insight Rally can deliver.

On-Premises Processing via User Defined Providers

For organizations with existing on-premises media processing infrastructure — transcoding clusters, hardware-based QC systems, or other fixed installations — Rally's User Defined Provider framework provides a structured mechanism to integrate these systems into Rally-managed supply chains, rather than requiring their premature retirement.

User Defined Providers are implemented by clients directly within the Rally UI, and support file handling, progress reporting, and job status surfacing in Rally's monitoring views.

There are structural characteristics of on-premises processing integration that organizations should be aware of. Because Rally does not control the underlying infrastructure, it cannot dynamically scale compute resources in response to workload demand, cannot reprioritize tasks across the cluster, and cannot provide the same level of cost and performance metrics as it can for cloud-managed processing. Operational insight is limited to what the on-premises system exposes through its API. Organizations should evaluate whether this level of visibility meets their requirements, balanced against the value of continuing to use existing infrastructure through its natural depreciation cycle.

Note on software licensing

Organizations that hold existing licenses for processing software — such as transcoding or QC applications — can in many cases deploy those licenses within cloud-hosted Rally supply chains, without requiring a hybrid deployment at all. This can be an effective way to extend the value of existing software investments while operating fully within Rally's cloud-native model.

Cloud Processing

Cloud processing is the native operating model for Rally. The platform is designed to spin up cloud compute resources on demand, execute processing tasks, and release those resources immediately upon completion — providing elasticity without idle infrastructure cost. Rally supports two cloud processing models:

SDVI-managed application services: For application services accessed through Rally Application Services Market, which is a SDVI marketplace for 3rd party processing applications, SDVI manages all underlying infrastructure — provisioning, scaling, monitoring, cost attribution, and optimization. This model provides full operational metrics including task duration, queue depth, and cost per job.

SaaS application integrations: Rally can trigger processing tasks against third-party SaaS applications via API — for example, cloud-based QC platforms or transcoding services. In this model, Rally manages job submission and status tracking, but infrastructure is owned and managed by the third-party vendor.

Design consideration: data locality and processing

Supply chains that process on-premises content using cloud-based tools will incur network latency and potential data transfer costs. Data transfer costs will vary based on the cloud provider. Organizations should account for this in their architecture decisions — for example, by staging content in cloud storage before executing multi-step processing workflows, or by using on-premises processing for content that remains on-premises. The right choice depends on content volumes, workflow structure, and the organization's infrastructure.

Trade-offs Summary

The following tables summarize the key operational and technical trade-offs across the storage and processing integration models described in this paper. They are intended to help architects and operations teams evaluate which combination of models is appropriate for their organization's current state and deployment strategy.

Storage Integration Options

Integration Model	Asset Inventory Tracking	Media Analysis	Operational Metrics	Implementation Complexity
Cloud Managed Storage	Full	Full	Full	Low
On-Prem Managed Storage	Full	Full	Partial	Medium
On-Prem Unmanaged Storage	None	None	None	Medium
Cloud Unmanaged Storage	None	None	None	Low-Medium

Processing Integration Options


Integration Model	Scalability	Operational Metrics	Implementation Complexity
Cloud Processing (SDVI-managed)	Elastic	Full	Low
Cloud Processing (SaaS integration)	Vendor-managed	Partial	Medium
On-Prem Processing (User Defined Provider)	Fixed cluster	Limited	Medium-High

Deployment Options

Rally is designed to support organizations wherever they are — and wherever they choose to go. Because the Rally control plane is always cloud-hosted and its integration mechanisms are designed to be incrementally adjusted as infrastructure evolves, the platform provides a consistent operational foundation regardless of the underlying technology deployment.

For some organizations, the priority is operational modernization: gaining the benefits of advanced, intelligent orchestration across existing on-premises infrastructure, with no immediate need to change what sits underneath. For others, the goal is a gradual evolution toward greater cloud utilization — at a pace and in a direction that reflects their own operational and financial reality. In either case, Rally accommodates the full range of approaches without prescribing a specific path.

The deployment models below illustrate a spectrum of typical configurations. They are not intended as sequential stages to move through, but as reference points that reflect the diversity of approaches organizations take in practice.



Throughout these options, Rally's supply chain definitions, asset inventory, and workflow logic remain constant. And, if appropriate, as on-premises components are retired and replaced by cloud-native equivalents, the integration mechanisms change — but the supply chains themselves do not need to be rebuilt. This continuity significantly reduces the operational risk of any infrastructure transition and allows organizations to move forward incrementally, validating each step before proceeding to the next.

Summary

Hybrid is a valid, well-supported operating model for the SDVI Rally platform — not a workaround or a temporary state. Rally provides structured, well-defined integration mechanisms for both on-premises storage and processing, with transparent trade-offs at each level of integration. Organizations can begin realizing the benefits of cloud-based orchestration and workflow intelligence today, against their existing infrastructure, and use Rally as the consistent operational foundation for their deployment strategy — whether that means a gradual evolution toward cloud-native operations, a long-term hybrid model, or something in between, if desired.