WHITE PAPER

superna

A VISION FOR THE FILM AND DIGITAL MEDIA INDUSTRY

Securing the future of Media Production, Post, and Creative Technologies with Superna

Andrew MacKay, Superna

OVERVIEW

This paper will discuss how Superna can help modernize content creation and management for the hybrid cloud era.

Any business process or workflow that moves documents or transforms data can benefit from a secure hybrid cloud architecture. The creation of digital media content involves many steps to produce a final product. While this paper focuses on the media industry, the same principles apply to any business process/workflow that benefits from a hybrid cloud architecture.

In white papers published by MovieLabs — *The Evolution of Media Creation* and *The Evolution of Production Security* — they present a forward-looking vision for the movie industry and how to think about securing collaborative content creation efforts in the era of the hybrid cloud.

Superna's core tenet of product development centers around a long history of innovation and data-centric product solutions that offer a "data first" approach to security, orchestration, management, and analysis. A data first security strategy places data (files or objects) at the center of any feature or function that's designed to secure a workflow. In this paper, we'll explore how Superna solutions address the key design principles outlined in *The Evolution of Production Security*.

In this paper, we'll highlight the benefits of a "data first" security strategy, how it aligns with the MovieLabs 2030 vision, and how Superna is delivering solutions *today* that move customers to this vision.

MOVIELABS' CORE PRINCIPLES

- **Security Principle 1**: Security is intrinsic to every component of every workflow, and does not inhibit creative processes.
- **Security Principle 2**: The security architecture addresses challenges specific to cloud-based workflows.

  a) Security is centered on workflows, rather than on infrastructure.

  b) Security is centered on assets, rather than their storage and transport.

HIGHLIGHTS

A data-first security strategy is the foundation for secure collaboration among content creators in the film and media industries.

- Security is essential to every workflow, and cannot inhibit creative processes
- The security architecture must address challenges of cloud-based workflows.
- Production workflows, processes, and assets must remain secure, even on untrusted infrastructure
- Content owner must control security and workflow integrity
- Security must scale to appropriate levels and integrate with existing security policies and management systems
- Security architecture must limit spread of any breach, while being adaptable to evolving threats

c) The integrity of assets, processes, and workflows is protected.

- **Security Principle 3**: Production workflows, processes, and assets are secure, even on untrusted infrastructure

  a) For our purposes, Zero Trust security is an ideal approach. It places the control of security in the hands of the content owner, it is part of the security architecture, and it can be implemented through specification of the security of the applications and other components that are used in the workflow. Granularity is controlled by policy, not technical constraints

- **Security Principle 4**: The content owner controls security and workflow integrity
- **Security Principle 5**: Security can be scaled to appropriate levels and can integrate with existing security policy and management systems
- **Security Principle 6**: The security architecture limits the spread of any breach and be adaptable to an evolving threat and response landscape

SECURITY PRINCIPLE 1

Security is intrinsic to every component of every workflow, and does not inhibit creative processes. By design, security is the key to this principle which means the data used within a workflow determines how security should apply to each file within the workflow. As files are created, metadata is created by all applications, additional custom metadata can be applied to files or objects that signals to the infrastructure how this file should be treated.

Let's use the example of a video application creating a new video clip and saving the video file. The application typically adds metadata about the clip such as project name, scene, resolution or bitrate etc… Many applications allow adding custom metadata to the file. This allows content creators to specify the security of the asset. If this asset was a new ending to a movie the file could be marked as "confidential" and this metadata attribute could be read by the infrastructure to apply security policies automatically based on this metadata.

How does this approach integrate security without inhibiting the content creator?

Content creators would be adding metadata to assets as a normal part of the workflow. Integrating security at the source of asset creation simplifies the integration.

What happens if the content creator forgets to tag the security of the asset?

Security falls down when manual processes are required. This is unacceptable when security of data is at stake. The data first security strategy handles this scenario easily, because the infrastructure will see that no security handling metadata exists and can quarantine the file and lock it down (apply ACL's to mark as read-only) and notify administrators in real-time that an asset has been quarantined. The infrastructure could even send an email to the content creator to apply the security tag to the asset so that it can be used in the workflow.

Problems addressed by a Data First Security Strategy

This approach avoids securing the infrastructure and allows the data to signal to the hardware and software that processes or touches the data in the workflow to react to the data inside the workflow. It allows metadata applied at source (application or content creator) to:

1. Determine how the file/object is handled inside the workflow.
2. Data carries its metadata with it, allowing the security tag signal to be honored at each step of the workflow process
3. Data without a tag can also be rejected inside the workflow
4. Metadata within files can be extracted and moved between files and object storage systems allowing full transparency between file systems and object based systems
5. Hybrid cloud architectures can natively handle file and object and metadata transparency
6. Workflows that use the same infrastructure can change the security model without changing the infrastructure

7. Metadata allows full customization and extensions to any solution by defining customer metadata
8. Security is only one possible use case for custom metadata
    a) Number of copies of the data required for data protection
    b) Encryption status for at rest protection
    c) Life cycle status (active, archival, short term archive, scratch data)
    d) Disaster Recovery requirements for the data
    e) Data type (pre-production, RAW, Final Cut, etc..)

How does Superna help?

Superna Golden Copy™ is a metadata-aware data orchestration platform for hybrid cloud media use cases.

1. Transparent metadata extraction from file to object feature allows rich metadata on files to be exposed when data is synced, moved to object. This functionality is enabled with support for over 1500 different file types.
2. Open data format allows all data copied to object to be read, written and updated with an native tools
3. Secure data once. This means leveraging SMB share permissions to limit users access to object data synced to object stores in the cloud. Who wants to manage two different security frameworks SMB, S3? No one!
4. Self service is *always* the answer.
    a) Storage administrators have no time to manage each and every users data access request or data movement from production file system to completed storage in S3. Self service is mandatory to enable the future hybrid cloud software defined workflow model.
    b) Workflows require data to move between file and object and object back to file, content creators, editors or other production user roles need tools that allow simple secure access to data regardless of where it's stored. Golden Copy enables Cloud browser which integrates file system and object store browsing into a single user experience.
    c) This enabled self-serve data movement orchestrated by Golden Copy but secured once with SMB permissions overlaid onto the object store preserving the security integrity of the data stored in the object store.
    d) Active Directory login allows seamless presentation of data to end users to view versions, metadata, backup status

SECURITY PRINCIPLE 2

The security architecture addresses challenges specific to cloud-based workflows. Let's review the key principles for workflow-based security architectures:

1. Security is centered on workflows, rather than the infrastructure they run on.
2. Security is centered on assets, rather than their storage and transport.
3. The integrity of assets, processes, and workflows is protected.

These principles all place data at the center of the security but only suggest moving encryption to the application layer with a Key management system to control the DRM (Digital Rights Management) of the assets. Encryption in the transport and at rest within the infrastructure can still be used but becomes redundant if the application layer owns the encryption/decryption of the assets.

The second aspect of securing the workflow is the hand off from one process to another. Let's reference this diagram from the MovieLabs Evolution of Production Security paper.

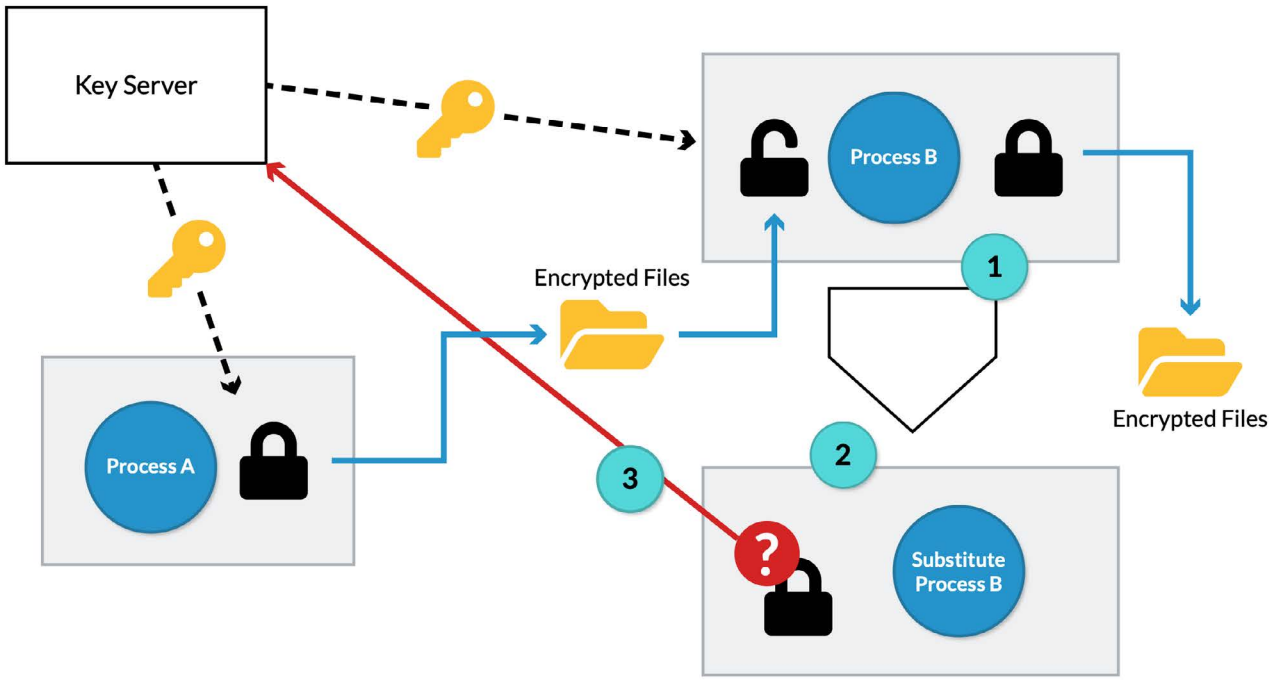

The diagram above shows an application and user-initiated workflow that requires both the user and the application to use the same key server when data moves between processes. Remember, in a hybrid cloud application architecture, Process A and Process B could be located anywhere: on premise, in the cloud, or across 2 different cloud providers.

The movement of data between Process A and B is not addressed in this example. Let's explore how the infrastructure can improve the security of the workflow.

Process A produces an application-encrypted file; the file has custom metadata added in this example <workflow>workflow-1</workflow>. The metadata can be applied by the application or, in a hybrid cloud use case, the metadata can be applied by the data orchestration layer (such as Superna Golden Copy™). As the data is placed on a NAS SMB share it's been configured to sync data to a Cloud S3 bucket where it's processed by Process B.

In the diagram below, we can see that Golden Copy reads metadata and only allows copying data to the Process B S3 bucket if it matches the data movement policies configured within Golden Copy.

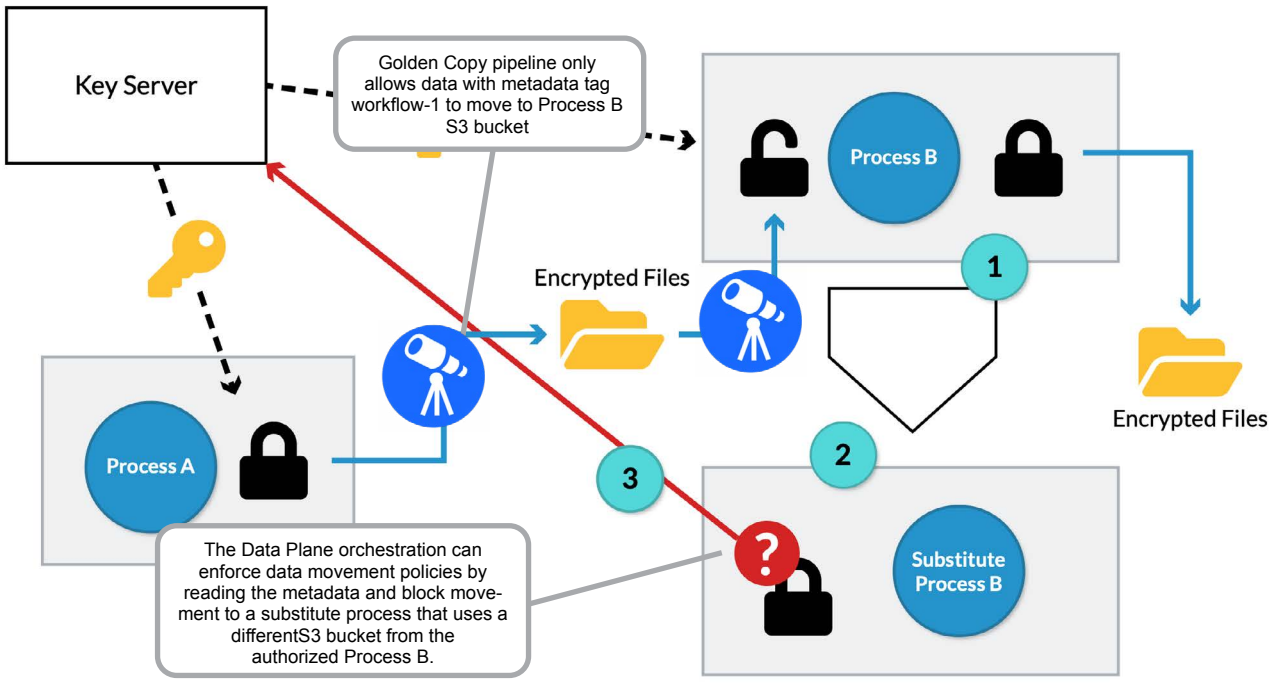

How can we increase the security of this workflow even further?

The garbage-in/garbage-out principle applies to security. If unauthorized data enters the workflow, this can lead to a potential data breach or denial-of-service attack on the workflow.

The workflow needs to trust the data it receives; this means the application needs to trust that the file that's being opened was created by an authorized user in the system. Application security and file system security are separate systems. Superna offers Easy Auditor that enables file system auditing and can identify who, what, and when data is created, deleted, or modified.

What would this extra layer of security look like?

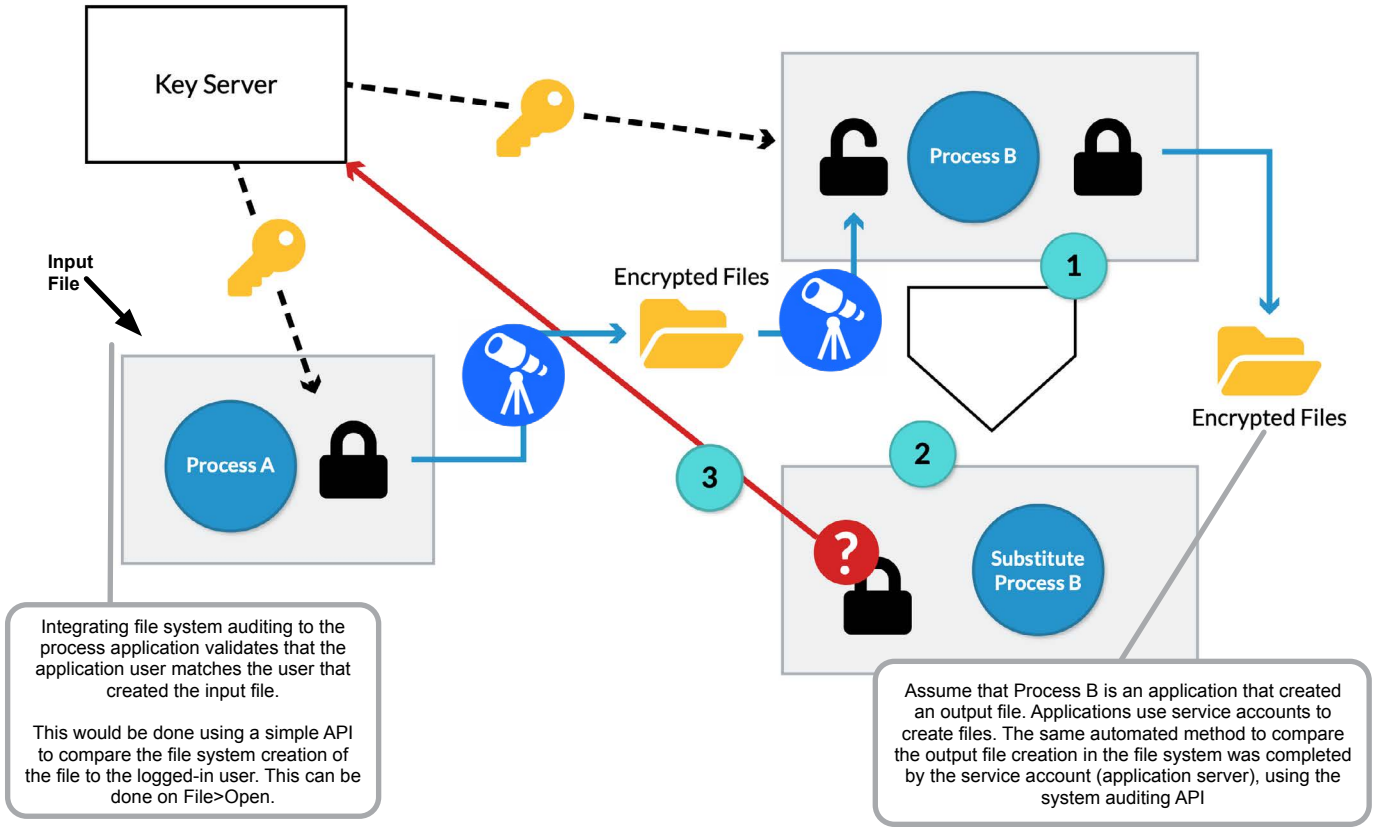

SECURITY PRINCIPLE 3

Production workflows, processes, and assets are secure, even on untrusted infrastructure

This principle is under the umbrella of Zero Trust architecture. The core premise is each component of the workflow: users, applications, data has Zero Trust applied at each step. What this means is that data flowing through this Zero Trust workflow design allows checking and validating that the data passing through was:

1. Created within the file system by an authorized user in the workflow (enforced by Superna Easy Auditor validation)
2. Created by an application service account (Enforced by Superna Easy Auditor validation)
3. Has valid metadata encoded into the file before it is allowed to move between processes (Enforced by Superna Golden Copy Data Orchestration policies)
4. Has data integrity checked when moving between processes (Enforced by Superna Golden Copy MD5 checksum validations)
5. Has valid encryption by using centralized Key Management server to decrypt the data at each step in the process

SECURITY PRINCIPLE 4

The content owner controls security and workflow integrity

The Security paper focuses on the content owner controlling how their content is secured at the application level but this can easily be extended to how the data moves through workflows and other downstream applications. Here are some examples of additional use cases that a content owner can control.

Content Owner Use Cases

1. Key management with Studio, Production or delegation of keys to a VFX facility to complete a process

2. Superna Extended Data Orchestration Use Cases

   a) The content owner could apply metadata to a file that signals the infrastructure to apply data centric policies to the content. Refer to examples below.

      - If the content is allowed to be moved, synced to the cloud
      - How many copies of the asset should exist? Example: 3-2-1... meaning 3 copies total: 2 different storage technologies and 1 offsite copy
      - Versioning requirements would allow the content owner to specify how many versions of the asset should be maintained, as well as if old versions of the asset should be archived for long-term storage.
      - Disaster Recovery requirements would mean the asset needs to be replicated to a secondary location that can present the data to applications, should a site or equipment failure be detected.
      - Data lifecycle management would allow the content owner to apply metadata that signals the infrastructure to move the data to long term storage
      - Immutability requirements for finished assets that need to be protected from modification

3. Superna Extended Asset Security Use Cases (Enforced by Easy Auditor)

   a) **DLP (Data Loss Prevention)** of sensitive data assets that need to be monitored for data copy and read requests from users or applications. This can be auto applied based on metadata that signals the infrastructure this data should be real-time monitored for data exfiltration attempts by authorized users or unauthorized users and applications. This requires file system auditing to be integrated with the workflows.

   b) **Failed data operations on sensitive data assets**. This is another file system monitoring requirement that alerts administrators or content owners that a file operation (open for read, write, delete or other data operation) failed due to insufficient permissions to the data.

   c) **Network-aware data operations**. This concept places workflows inside a trusted network (Zero Trust principle) and treats all networks outside the trust zone as untrustworthy.

      - A network is a group of IP subnets and can be used to provide real-time notification of access to workflow data from untrusted IP subnets. If application servers or users operate within a trusted network (i.e., a VPN IP address space), then all data requests outside of this can be flagged as a potential breach

   d) **Trusted Users**. The list of authorized users would be known in advance for all workflows, so applying a real-time trigger of any file system data operations attempted by a user or application service account that is outside the trusted user list can be flagged to administrators or content owners.

Summary

By extending the use cases to include a “data first” security strategy controlled by content owners, the security of the workflow can be integrated while still maintaining the core principles of who should control how assets are secured, accessed, protected or moved within the workflow components. The solution based on Superna’s Easy Auditor and Golden Copy hides the complexity of securing, protecting data from the Content Owner. This is a core objective when designing for security. Your data is unmodified using metadata features of OneFS and S3 to store critical information about how your data should be treated in a hybrid cloud environment.

SECURITY PRINCIPLE 5

Security can be scaled to appropriate levels and can integrate with existing security policy and management systems

In short, this security principle stipulates that security and risk management should be scalable to meet basic requirements all the way to the highest level of data/asset security. Studios have different departments that have different security requirements and all of these different requirements need to be met with a single solution. It is also important to allow studio's to use different tools, applications, IAM systems and that SCIM (System for Cross-domain Identity Management) can help address centralized security.

Superna's Data First Security Strategy

The key to Superna's strategy of "data first" security is based on metadata that's carried through a workflow to signal which policies should be applied at each point in the workflow. File system to object and object back to file system needs to preserve metadata so that it can be inspected anywhere within the workflow. The other key benefit of using metadata is that different vendor solutions can read the metadata and apply policies or extend value by simplifying defining their own metadata to differentiate or extend the workflow without concern for managing a centralized SCIM schema.

While the likelihood of getting agreement on a multi-studio schema or standard is understandably low, a strategy that allows vendors to innovate and allows for a multi-vendor technology ecosystem is far more viable.

SECURITY PRINCIPLE 6

The security architecture limits the spread of any breach and Is adaptable to the evolving threat and response landscape

This principle discusses the *architecture*, not the specific technologies used within a Zero Trust architecture which determines the longevity and ability to adapt to new threats. In order to build a system that can evolve and meet this goal, some building block technologies need to be selected. The MovieLabs 2030 vision is based on this diagram.

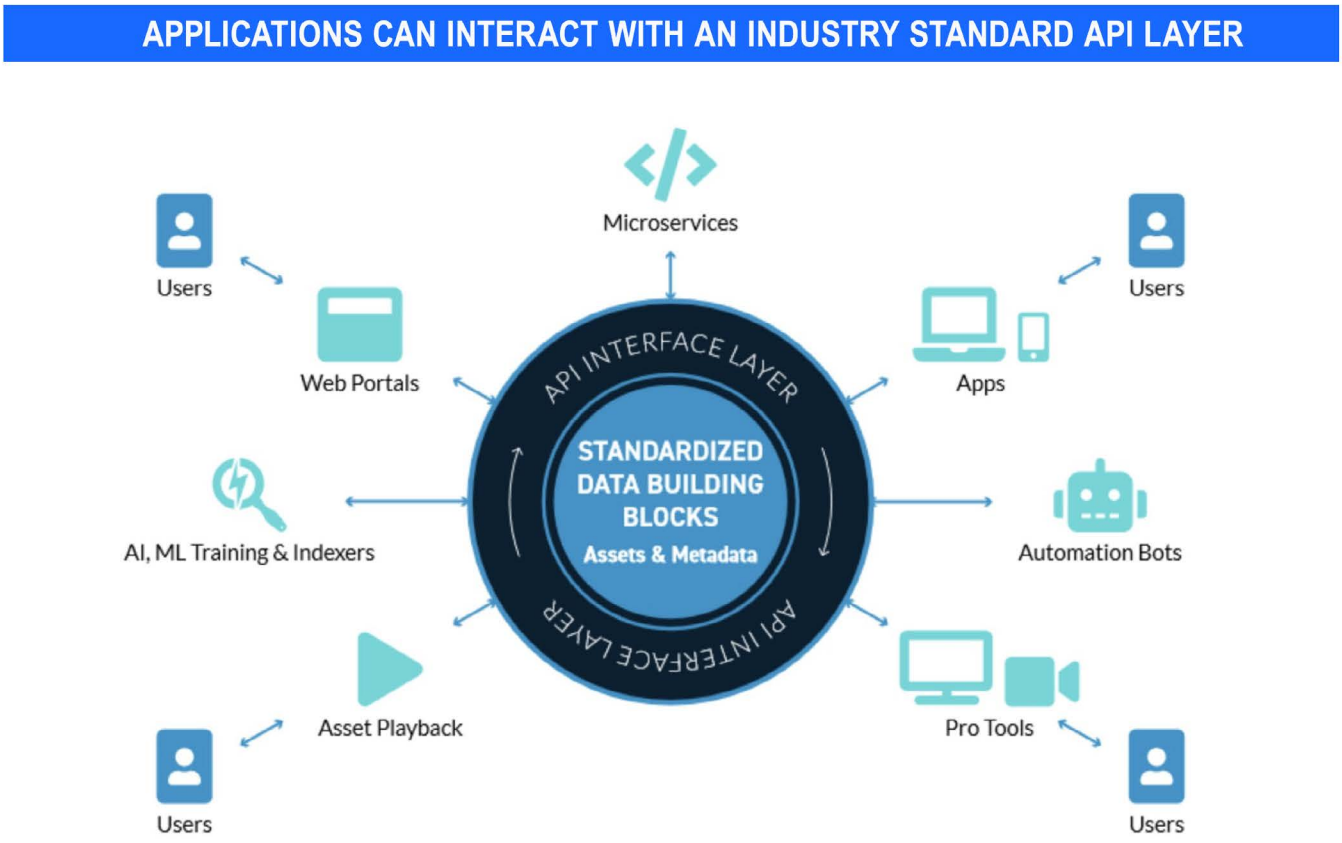

Standardized building blocks contain agreed-upon common data and metadata for key processes in production. The API layer abstracts that information so that any appliction, tool, portal or service can interface through the API layer with any other tool without needing to know about it in advance.

Placing content owners at the center of the process allows data to carry policies with it that dictate how an asset should be secured, accessed, protected, played back or processed.

As you can see from the diagram above, specific API's need to exist with data assets and metadata. These API's need to execute within software that interacts with the data assets. The API's need to exist across multiple vendors to allow seamless implementations that offer consistent behavior and more importantly, consistent security across all of the use cases shown in the diagram above.

Therefore, the strategy needs to allow multi-vendor API's to be built on top of metadata that all vendors can extract to enable the use cases, workflows, and security requirements of a modern media creation ecosystem.

SUPERNA'S ROLE IN THE JOURNEY TO THE MOVIELAB 2030 EVOLUTION OF MEDIA CREATION

The shift to a hybrid cloud or fully cloud-enabled architecture needs to address the storage technology choices that limit realization of vision. Specifically, the file vs object choice and the selection of the best storage for the use case. This paper is not going to answer which storage technology should be used for each workflow or use case but, more importantly, highlights that either file or object or both need to co-exist to enable the vision.

What problems exist today with file and object storage?

Some background on the issues that exist when mixing and patching file and object storage within your workflows.

Function	File	Object	Issue
Byte range read and write	Yes	No	Applications need to be writter differenty for object; many Use Cases require byte storage range IO, and eliminates object storage
Random IO	Yes	No	Same as above
Partial updates to data	Yes	No	Same as above
Symbolic links and hard links	Yes	No	1. File systems leverage symlinks nd hard links for rapid cloning file system paths and reducing duplicate file issues for media workflows 2. S3 object storage has no functional equivalent
High-performance streaming	No	Yes	Object supports parallel reading and writing with muti threads and cn outperform file systems for some use cases.
Embedded metadata	No	Yes	1. SMB and NFS don't allow creation of custom metadata on files, this is built into the S3 object protocol 2. Dell PowerScale supports an API to attach metadata to files; Superna Golden Copy™ uses this API to track data movement history as a property of the file folders
Storage Tier metadata	No	Yes	S3 has the concept of storing data based on a lifecycle stage at the object level; this does not exist ntively within file systems
Data Versioning	No	Yes	S3 has the bility to store versions of data and provide read/write access to each version. File systems have no versioning cpabilities.
Data integrity	No	Yes	When data is moved, copies in S3 MD5 can be used by the sender and receiver to ensure that data arrives error-free. File systems don't have this type of built-in integrity checking
Authentication and metadata checking and validation	No	Yes	S3 allows the metadata of an object to be included in the authentiction process between sender and receiver. This means that custom properties are part of the vlidation process. FIle systems have no equivalent capability.
Access Control Lists	Yes	Yes	Used frequently in file systems, rarely used in S3 but also incompatible with no simple mapping for ACLs on files and folders to S3 prefixes and buckets.
In-flight encryption	Yes	Yes	1. FSMB uses SMB3 and is not supported by all devices and is disabled 2. NFSv4 yes, but bery complex configuration, and rarely used 3. S3 https works over any network, enabled by defaut 4. S3 allows control over certificates for sender and receiver authentication and TLS level selection is easily managed

What does it all mean? To realize the MovieLabs 2030 vision, hybrid cloud architectures need a robust solution for file and object transparency. While storage is key for supporting modern workflows, security should start at the data layer and move out to compute, network and endpoints. This is the core tenet of Superna's product innovation, to deliver full transparency for the hybrid cloud. It demonstrates that Superna is well-positioned to help the entertainment industry achieve its ambitious objectives for technology and collaboration, as outlined in MovieLabs' *The Evolution of Media Creation*.

Criteria	Superna Product	Compliance with Secure Hybrid Cloud
File to object and object-to-file data orchestration	Golden Copy	✓
Automation API	Golden Copy	✓
Partial updates to data	Golden Copy	✓
Carrier-grade upgrade (no down-time upgrade, containerized, built for cloud scaling)	Golden Copy	✓
Broadcast publish (single source to multiple cloud targets)	Golden Copy	✓
Contribution mode (many S3 buckets to single file system)	Golden Copy	✓
File auditing for integration to advanced workflows (files moved, synced, File-to-Object of Object-to-File available in JSON files for developers	Golden Copy	✓
Transparent security (POSIX permission transparency to object and back again)	Golden Copy	✓
4K long path names (File system to S3 URL transparency)	Golden Copy	✓
Integrated security for data movement orchestration	Ransomware Defender	✓
Symlink and hardlilnk transparency for media	Golden Copy	✓
High-performance File-to-Object copy, sync, or move	Golden Copy	✓
Single Security Model (Transparent File System to Object data browsing)	Golden Copy	✓
Data Lifecycme Support (S3 tiering)	Golden Copy	✓
Open Object format with cross-cloud support	Golden Copy	✓
Cloud ingest with move, sync to File Systems	Golden Copy	✓
Full File System Auditing for MPA Content Security Compliance	Easy Auditor provides compliance with DLP Detection, User Access Reporting, User Data Access Auditing, Honeypots	✓
Full Object Data Access Auditing for MPA Content Security Compliance	Easy Auditor (AWS S3, Dell ECS)	Planned
Hybrid cloud security	Superna Defender for AWS S3, Defender for PowerScale, ECS	✓
Metadata-aware to expose asset	Golden Copy	✓

For more insight into how Superna® can help solve your organization's unstructured data security challenges, visit us at superna.io.

superna
Securing unstructured data, wherever it resides

superna.io
marketing@superna.net

20230228