

2015

# Expediting Digital Workflow with OpenStack

OPENSTACK WHITE PAPER

## Contributors:

**Guillaume Aubuchon**, *CTO & Managing Partner, DigitalFilm Tree/DFTi*

**Kathy Cacciatore**, *Consulting Marketing Manager, OpenStack Foundation*

**Morgan Fainberg**, *OpenStack Identity Project Team Lead (PTL), HP Cloud Engineering*

**Chris Hoge**, *Interop Engineer, OpenStack Foundation*



This work is licensed under the **Creative Commons Attribution-NoDerivatives 4.0 International License**.  
To view a copy of this license, visit <http://creativecommons.org/licenses/by-nd/4.0/>

## Executive Summary

Supply chains used to be a matter of assemblers, conveyer belts, and loading docks. But in today's software-defined economy, that's not the only kind of supply chain. Most industries, including banking, medical, retail, and entertainment are being transformed by online and mobile applications that deliver secure, on-demand, self-service capabilities directly to customers.

These capabilities aren't the product of a single inventor or creator. The delivery of an application or service often depends on many contributors, possibly in multiple companies, in an end-to-end workflow.

Imagine this scenario: with the right technology, could a business create a workflow that moves a digital project across many contributors in different organizations? Could the workflow be so adaptive as to allow each contributor to add their contributions efficiently – regardless of the size, complexity, and location of the project and people?

Yes, it's possible. Adaptive workflows benefit from on-demand, ubiquitous access to private and public clouds that offer widespread and virtually unlimited compute power, storage capacity and network bandwidth. These digital workflows require projects to be portable and secure across many clouds to align scale, capacity, and performance to unique project requirements. Cross-cloud identity services ensure contributors or providers have access to required resources for their portion of the workflow regardless of which cloud it resides on. Massive cloud capacity has the potential to accelerate existing workflows as well as enable innovative future workflows.

The multitude of OpenStack clouds around the world provide portability and interoperability through open APIs, a consistent set of core services and federated identity. Applications can run on interoperable private, public and hybrid OpenStack clouds without modification – allowing organizations to create workflows to fit ever-changing requirements so they can evolve and innovate.

- Interoperability testing and branding allows applications to be written once, for execution on many clouds
- Federated identity allows choice of the best infrastructure for applications based on price, performance, geography and other requirements

This paper explores a real-world example from the entertainment industry. DigitalFilm Tree, a leading post-production company for television and film, moves digital footage through a workstream that includes initial filming, editing, special effects creation, post-production, and review across many collaborators. Interoperability among OpenStack clouds allows DigitalFilm Tree workloads to flow seamlessly, using the resources needed at each stage, while maintaining security from end to end with a single point of identity control.

DigitalFilm Tree is using their post-production work on several shows, including UnREAL (launching on Lifetime in June, <http://www.mylifetime.com/shows/unreal>) to demonstrate how IT organizations can leverage the portability and extensibility of OpenStack by integrating it into existing infrastructure and processes.

By tailoring cloud capabilities to the project, an enterprise like DigitalFilm Tree can maximize business benefit by minimizing time, cost, and risk. In addition, they are building their reputation as the leader in the post-production field by demonstrating the opportunities and benefits any software defined enterprise can experience.

## From Camera to Couch

DigitalFilm Tree demonstrated how OpenStack clouds help fill the gap between existing workflows and emerging requirements. For production and post-production processes, they uncovered ways to use OpenStack (specifically Swift, Nova, Keystone and Neutron) to maximize efficiency while providing better services to studios.

DigitalFilm Tree needs to scale for a 2-3 month period for the pilot season, then scale back down. Some projects are much larger or more intensive than others, due to format (4K vs 8K video) or complex special effects requirements. DigitalFilm Tree also doesn't always know when shows will start or stop. Leveraging public and hybrid cloud allows DigitalFilm Tree to adapt to change at will, effectively creating a new business model by instantly augmenting their workflow.

Originally their workflow spanned multiple private cloud locations, with one node in Vancouver, one node in Los Angeles, and a third editorial node. Now they're extending their workflow beyond private cloud, demonstrating that moving part of their workflow to HP Public Cloud and hosted private cloud provider Blue Box offers additional benefit.

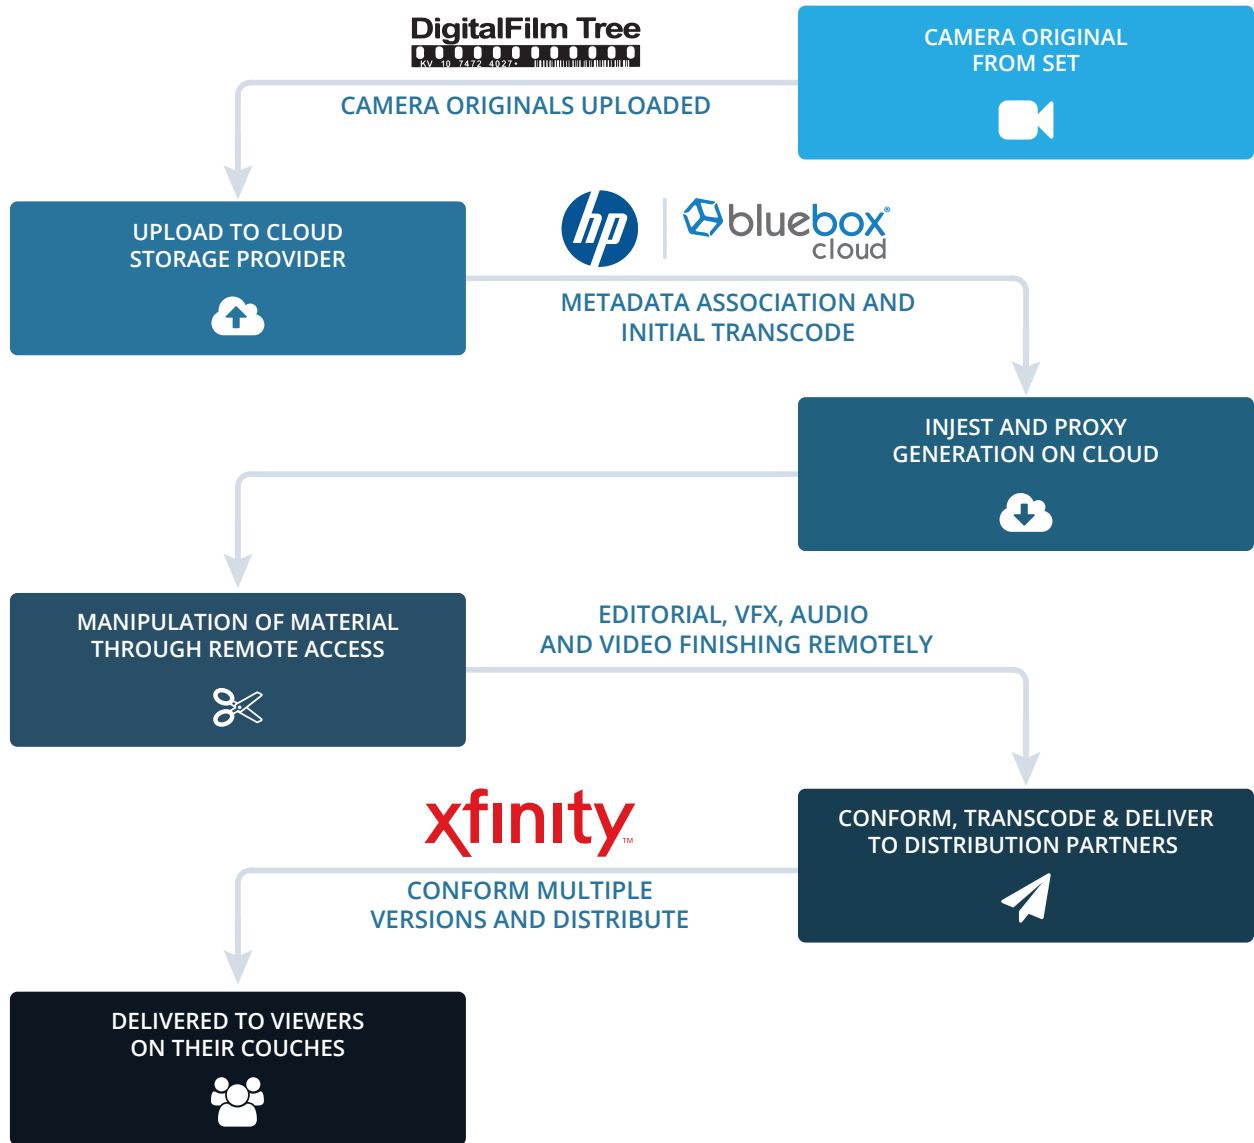


Figure 1. DigitalFilm Tree Post-Production Workflow Through Multiple OpenStack Clouds

Their workflow is as follows:

- Directors shoot raw footage for the UnREAL TV show on digital video in Vancouver. Filmmakers upload video to temporary storage in the HP Public Cloud and Blue Box Cloud.
- Raw footage is transcoded from high resolution to proxy resolution on the external clouds, moved to Swift object storage and transported between Vancouver and Los Angeles.
- Collaborative process until footage is approved as final:
  - The DigitalFilm Tree team in Los Angeles cuts and processes the raw footage, performing editing, color correction, and digital effects.
  - Raw and post production footage is distributed to directors and producers in the field for rapid editorial feedback.
  - Reviewers utilize mobile devices and an app called Critique running on OpenStack.
- Final video footage is assembled. Director/producer feedback is incorporated, outsourced special effects are rendered and added.
- The show or movie is transferred from the DigitalFilm Tree cloud to Comcast or another cable provider. This expedites on-demand delivery to viewers' living rooms via set-top boxes.

Essentially, DigitalFilm Tree offloaded compute of the transcode and rendering processes to HP Public Cloud and Blue Box while using both for storage. This gave them efficiencies of scale and performance.

However, there are potential obstacles to augmenting a workflow with public or hybrid cloud capacity. When DigitalFilm Tree planned this workflow, they immediately recognized two fundamental requirements that had to be met for success.

## Workload and Data Portability Across OpenStack Clouds

The first requirement is 100% reliable workload and data portability – in other words, any OpenStack cloud or distribution has to provide identical foundational capabilities so developers are confident that apps developed on one cloud will work perfectly on another. That's a function of OpenStack's REST APIs, which are easy to learn and program against, and ensure that when a developer writes an app for one OpenStack cloud, it will work across many. An open, stable API means application portability.

OpenStack values interoperability because it enables optimal IT decision making. To guarantee interoperability, the DefCore committee was formed to define the OpenStack core. This definition uses community resources and involvement to drive interoperability by creating the minimum standards for products labeled "OpenStack." DefCore sets base requirements by defining capabilities - encompassing must-include code and must-pass API tests for all OpenStack products. The interoperability requirements ( <http://www.openstack.org/interop/>) enable workload portability, delivering application consistency across a multitude of public and private clouds, so workflows can function seamlessly with the right capabilities at the optimal cost.

It's not just the computational aspects of workloads that need to move but also the data the computation runs against. OpenStack software is built so that data can be easily shared internally within an OpenStack cloud and also with other OpenStack clouds. OpenStack Object Storage (code-named Swift) facilitates building redundant data distributions across wide geographic regions and multiple data centers. It transparently replicates data for redundancy and performance within one cloud, and can also expose data through public APIs for other users and clouds to consume. In the DigitalFilm Tree example, footage is uploaded to a Swift container in a data center in Vancouver. This data is replicated across geographic regions where it can be accessed locally by production staff in Vancouver and Los Angeles. The data can also be shared publically with any other vendor that may need to access that resource, such as a special effects company.

# Federated Identity, Simple Security

Workload and data portability, though essential, isn't enough. In addition to workload portability, OpenStack Identity Service (code-named Keystone) provides DigitalFilm Tree's OpenStack cloud with federated identity capabilities to streamline secured access across many clouds.

## Why is that important?

DigitalFilm Tree's cloud, running Keystone, acts as an identity provider for all the collaborators in the process who need to access compute and storage resources for their portion of the workflow. With Keystone, collaborators can easily use the best-fit cloud for their tasks. Though DigitalFilm Tree uses their private cloud, HP Public Cloud, and hosted private cloud by Blue Box for various parts of the workflow, theoretically DigitalFilm Tree could use any OpenStack cloud to augment their workflow. With OpenStack, a workflow owning enterprise can use one credential set across multiple clouds. The enterprise maintains a single point of access for everyone involved in the workflow, simplifying and enhancing security.

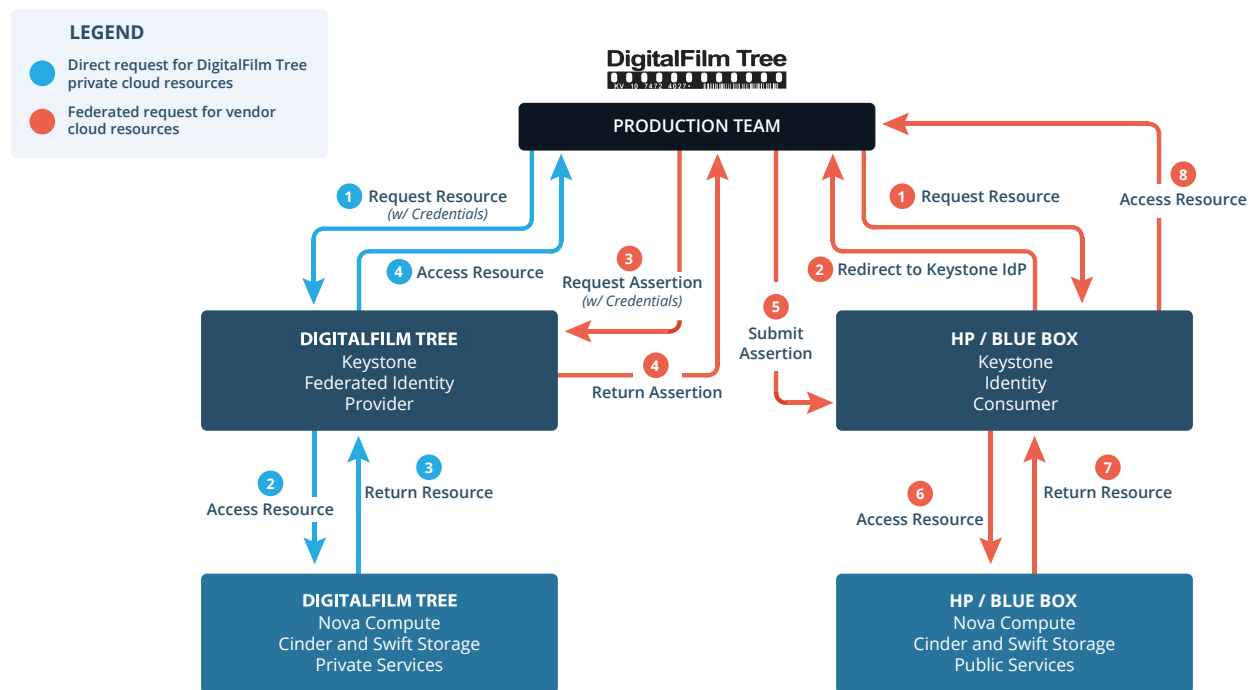


Figure 2: DigitalFilm Tree Provides Federated Identity Across Many OpenStack Clouds

How does that work? With Keystone identity services, there are at least two parties: an identity provider and one or more identity consumers (service providers) all running OpenStack Keystone. To put it simply, Keystone at the identity consumers confirms incoming requests by validating a set of user credentials (userid and password) with the identity provider. The identity consumers, usually cloud service providers, give access to secured applications, resources, and content once credentials are approved.

Keystone provides IAM: Identity (authentication) and Access Management (authorization). DigitalFilm Tree maintains the identity for each employee, vendor or contractor. With this central source of identity, DigitalFilm Tree can sign up for another cloud, all the identities stay the same, the admin just sets up another identity consumer. XML files are exchanged with cloud identity credentials. Users now point to another cloud, transparently allowing use of additional resources.

To help understand this, there's a familiar example: the interaction between Instagram and Facebook. When someone accesses Instagram for the first time, the application asks to login using Facebook; Facebook is the identity provider and Instagram is the identity consumer. Facebook has become a default source of identity, facilitating access for many services by maintaining a single username and password, allowing those services to alleviate maintenance of usernames and passwords of their own. DigitalFilm Tree implemented Keystone for the same basic reason: to provide an identity service to simplify access and security across their multi-cloud digital workflow.

OpenStack's federated identity capability is based on open federated identity standards. The first of these is Security Assertion Markup Language 2.0 (SAML 2.0), a version of a standard for exchanging authentication and authorization data between security domains. OpenStack uses another open source project, Shibboleth, instead of re-inventing a standard implementation. OpenStack tends to leverage existing open source approaches to ensure commonality and consistency.

With federated identity, the basic idea is that one identity unlocks services, no matter which cloud provides the services. Single-sign-on (SSO) like technology serves as the basis of federation so that an OpenStack user doesn't need to create a userid/password for every cloud they want to access. The user authenticates (makes a request) against Keystone and gets a standard OpenStack token, with the request kicking off an internal check of the token. Keystone plugs into the LDAP or Active Directory already in place. With a token, trusted users have the ability to spin up not only private cloud, but also public cloud resources. Federated identity enables user or group level authorization and can set up resource restrictions and access to specific services like Nova, Swift, etc.



First, external clouds have to be set up as part of DigitalFilm Tree's trust. Administrators add public clouds to the trust and ensure federation is enabled, as do administrators of the public clouds, who must verify the trust of DigitalFilm Tree. One of the major criteria for trusted status is robust security, but interoperability is the other relevant criteria. DigitalFilm Tree's IT admins verified both before entering into a business relationship with HP Public Cloud and with Blue Box.

## Benefits

Workload portability and federated identity are important features that convey technical benefits, but it's important to realize that they both support a particular kind of adaptability that's a result of widespread OpenStack availability. That adaptability is the core reason why DigitalFilm tree, its vendors, and its clients enjoy substantial business benefits, including:

- **Instant-on capacity:** Rendering and editing are highly compute and storage intensive tasks. Digital resources can be brought online during peak times to support and accelerate their color correction, grading, special effects and rendering processes.
- **Granular utilization:** they only consume those resources that they need, then release them – controlling costs.
- **Complexity reduction:** in the past, to cope with pilot season, DigitalFilm Tree would add extra servers and storage to their private clouds, then decommission that infrastructure when it wasn't required any longer. Public cloud consumption minimizes that complexity during peak times.
- **Latency:** by utilizing third-party cloud resources close to their work locations and data centers, they benefit from lower latency and improved bandwidth.
- **Resiliency:** data exists in multiple clouds and users can access seamlessly. If one cloud goes down, users access copies with no loss of service.
- **Client requirements:** some studios want data stored in specific regions, so widespread availability of OpenStack clouds in different geographies facilitates these requirements.
- **Retention:** DigitalFilm Tree has ample capacity to store and archive all content, from pre-production to finished project with as much metadata as possible.
- **Workflow acceleration:** Prior to this approach, video was shipped on disk or tape via truck. Typically, this meant a simple editing or post-production effort took 2-3 days for a review cycle. With OpenStack software, their workflow now takes 2-3 hours.
- **Avoiding media:** in the past they used hard drives or tapes. Now they avoid the costs and risks of physical media transport.

- **Workflow simplification:** In the past, final content was written to tape, and sent to Comcast on a truck. Now DigitalFilm Tree or one of their cloud providers can share the final video to Comcast's directly.
- **Enhanced security:** Since DigitalFilm Tree acts as an identity provider, maintaining security and access is streamlined because they are the only party responsible for maintaining authentication directories. Changes (new hires, resignations) are effective across every service provider instantaneously. Also, admins can add federation enabled external clouds as part of DigitalFilm Tree's trust as needed, minimizing ad hoc reliance on untrusted public cloud resources.
- **Planning for the future:** archiving metadata enables new capabilities such as "second screen". Imagine using a tablet to order the dress an actress is wearing on a television show.

These advantages align to fundamental business benefits.

### Reducing time

Overall DigitalFilm Tree reduced the amount of time workflows consume. Processes that took days now take hours. Bringing up extra capacity takes minutes instead of weeks. Administrative time is reduced and projects are completed more quickly.

### Reducing risk

Federated identity reduces risk by maintaining a single administrative point for access. But risk reduction goes beyond federated identity. DigitalFilm Tree now has a massively resilient architecture for their workflows that minimizes the risks of data loss and performance degradation.

### Decreasing cost

All the advantages and benefits come down to decreasing costs.

DigitalFilm Tree now has the ability to scale up and down as requirements change. Using OpenStack public and hybrid cloud resources to right-size their capacity requirements gives them unparalleled granularity – which translates into controls on cost.

Beyond that, DigitalFilm Tree's approach allows them to meet customer requirements more efficiently while reducing capital expenditures and operating expenses.

### Enabling innovation

These benefits are a direct result of many OpenStack clouds working together to deliver seamless compute, storage, and networking. Their hope is that as more entertainment companies standardize on OpenStack technology, new workflows can be developed to accelerate business transformation across the industry. More innovations will emerge as more entertainment companies rely on OpenStack software.

## Summary

DigitalFilm Tree – and its vendors -- can leverage available public clouds with internal, private cloud security. With Keystone identity services running on all systems in the workflow, users have hybrid cloud capability that's not complicated or compromised.

Any enterprise that uses intensive, collaborative digital workflows in a software defined economy could benefit from this approach. Putting a workload on “best fit” clouds to optimize resource availability and cost-efficiency is a huge opportunity enabled by an OpenStack-powered planet. Many use cases, including scientific, medical, insurance and industrial applications could benefit from the approach that DigitalFilm Tree successfully demonstrates.

*“Federated identity and interoperability are features that speak directly to the enterprise. Economics, proximity and performance requirements change over time, and enterprises need to move workloads and data to the infrastructure that makes the most sense for that moment in time. DigitalFilm Tree’s workflow illustrates perfectly a use case we see from customers on an increasingly frequent basis. For companies running applications built for OpenStack, it makes sense to leverage the global ecosystem of OpenStack providers.”*

**André Bearfield**, Senior Director of Product, Blue Box Cloud

For more information on OpenStack products and the vendors which have achieved the OpenStack Powered logo of interoperability and are supporting federated identity, visit the OpenStack Marketplace at <http://www.openstack.org/marketplace/>.