

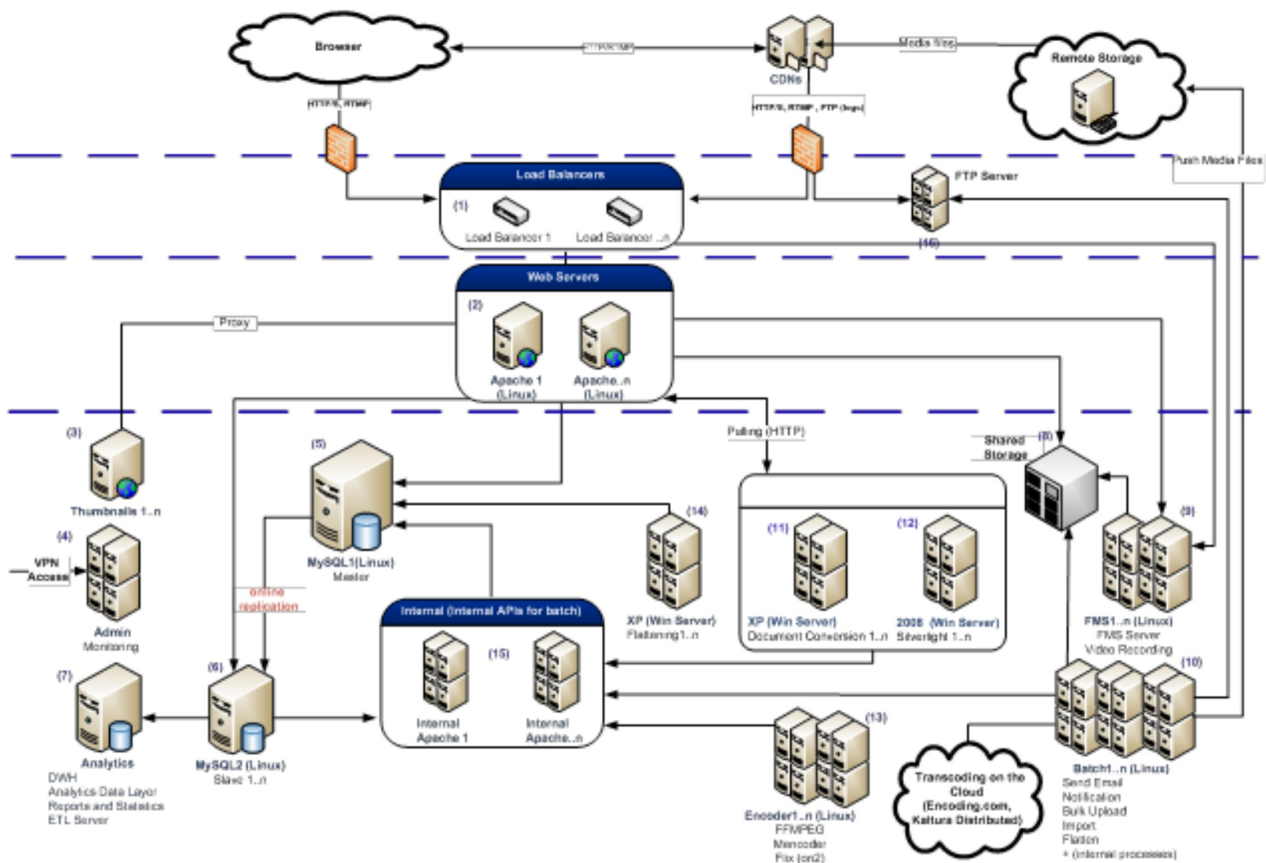
Overview of Physical Setup

Last Modified on 11/18/2019 12:16 pm IST

The physical deployment of the Kaltura platform is flexible and should be tailored according to the size and applicative needs of the specific site deployment. For optimal resource utilization as well as security reasons, Kaltura platform deployment should be separated to a front-end layer and a back-end layer.

Sample Deployment Architecture

The following diagram illustrates a sample deployment architecture:



Category	Component Name	Description
	Load Balancers	LVS based, running on Linux server
API	Web Servers	Kaltura API and access to shared content
API	Thumbnail Servers	Dedicated API servers for generating thumbnails from images or videos
API	Internal API Servers	Internal API servers for Kaltura batch servers
Data	Database master	Primary MySQL server for write operations
		MySQL servers for read-only

Category	Component Name	Description
Data	Sphinx servers	Search engine for text-base search
Data	Memcache servers	Caches database query results
Data	Shared storage	Stores all Kaltura media assets, accessed by all relevant servers over a distributed file system
Data	Analytics database	Runs analytics ETL software and a data warehouse database
Batch	Batch servers	Kaltura dedicated batch servers
Batch	Document conversion servers	Windows based server that performs document conversion tasks
Batch	Encoding servers	Kaltura generic Linux-based encoding server
	FMS servers	Servers for webcam recording
	FTP servers	FTP server for pushing logs from CDN
	Admin server	Monitoring and application administration

Virtualization

All server components listed in this section can be set on a physical machine or run on top of a virtualization infrastructure. Kaltura provides a standard VMware image to be deployed and run on a VMware ESXi Hypervisor setup. In this case dedicated hardware should be [VMWare ESXi compatible](#).

Standard Setup Sizes

The following sections describe Kaltura's recommendation for required hardware to be set for initial setup of an instance of the Kaltura platform. Specific site assessment can be made by Kaltura's engineering and IT teams upon request.

Small - Single Server Setup

Single server setup may be suitable in the following cases, all assuming low traffic and CPU utilization:

- Small deployment of a Kaltura Community Edition instance.
- Platform deployment dedicated for providing limited usage video capabilities to a CMS/LMS platform while using Kaltura CMS/LMS extensions
- Platform deployment used for development or platform evaluation purposes.

In this setup, all required components are deployed on one server. Document conversion module and workflow cannot be utilized in this topology.

Medium - Distributed Site Setup

This physical setup is suitable as an initial setup for small to medium-size deployments of the Kaltura platform, that are expected to serve medium to high traffic/streaming volume and medium to high number of video uploads and transcoding operations. With growth in service utilization this initial setup can be scaled up as needed, to add the required server type (e.g. additional transcoding servers, additional database servers for replications, additional front end server etc).

Server	Recommended Hardware Specs.	Operable Components
Load Balancing Server/s	Site Admin responsibility	Site Admin responsibility
Platform Shared Storage	<p>Can reside on one of the platform servers. Storage size to be set according to the expected volume of media assets.</p> <p>On average - 3X size of the expected video source files (assuming source should be stored as well).</p> <p>To be accessible by the Kaltura platform servers via NFS.</p>	Storage
2 Front-End Servers	64bit based Linux OS with 4-8 GB RAM and 1 quad core CPU	<p>Kaltura Web Services Module, on top of an apache server</p> <p>Kaltura Video Recording Module (if applicable)</p> <p>Connected to Shared Storage (NFS)</p>
Back-End Batch Server	64bit based Linux OS with at least 4GB RAM and 1 quad core CPU	<p>Batch Jobs Module</p> <p>Connected to Shared Storage (NFS)</p>
Back-End Transcoding Server	64bit based Linux OS with at least 8GB RAM and 2 quad core CPU	<p>Transcoding Module</p> <p>Connected to Shared Storage (NFS)</p>
Back-End DB Server	64bit based Linux OS with at least 4GB RAM and 1 quad core CPU	<p>Operational Database</p> <p>Sphinx – full text search engine</p> <p>Connected to Shared Storage (NFS)</p>

Server	Recommended Hardware Specs.	Operable Components
Back-End DWH Server	64bit based Linux OS with at least 8GB RAM and 1 quad core CPU	Video Analytics Module Connected to Shared Storage (NFS)
Site Admin Server	Linux based OS. No minimal specifications on hardware	Site Admin Module
Document Conversion Server (optional)	Windows server 2008 With at least 2GB RAM. 1 CPU	Kaltura Document Conversion Module (if applicable)

Large - Full Distributed Site Setup

This physical setup may be suitable as an initial setup for large-size deployments of a Kaltura platform, that are expected to serve high traffic/streaming volume and high number of video uploads and transcoding needs. With any growth in service utilization this initial setup can be scaled up as needed, to add the required server type (for example, additional transcoding servers, additional database replication server, additional front end server and others).

This setup may include the following servers:

Server	Recommended Hardware Specs.	Operable Components
Load Balancing Server/s	Site Admin responsibility	Site Admin responsibility
Few Front-End Servers	Several 64bit based Linux OS with at least 8GB RAM and 1 quad core CPU	Kaltura Web Services Module, on top of an apache server. Kaltura Video Recording Module (if applicable) Connected to Shared Storage (NFS)
Few Back-End Batch Servers	Several 64bit based Linux OS with at least 4GB RAM and 1 quad core CPU	Batch Jobs Module Connected to Shared Storage (NFS)
Few Back-End Transcoding Servers	Several 64bit based Linux OS with at least 8GB RAM and 2 quad core CPU	Transcoding Module Connected to Shared Storage (NFS)

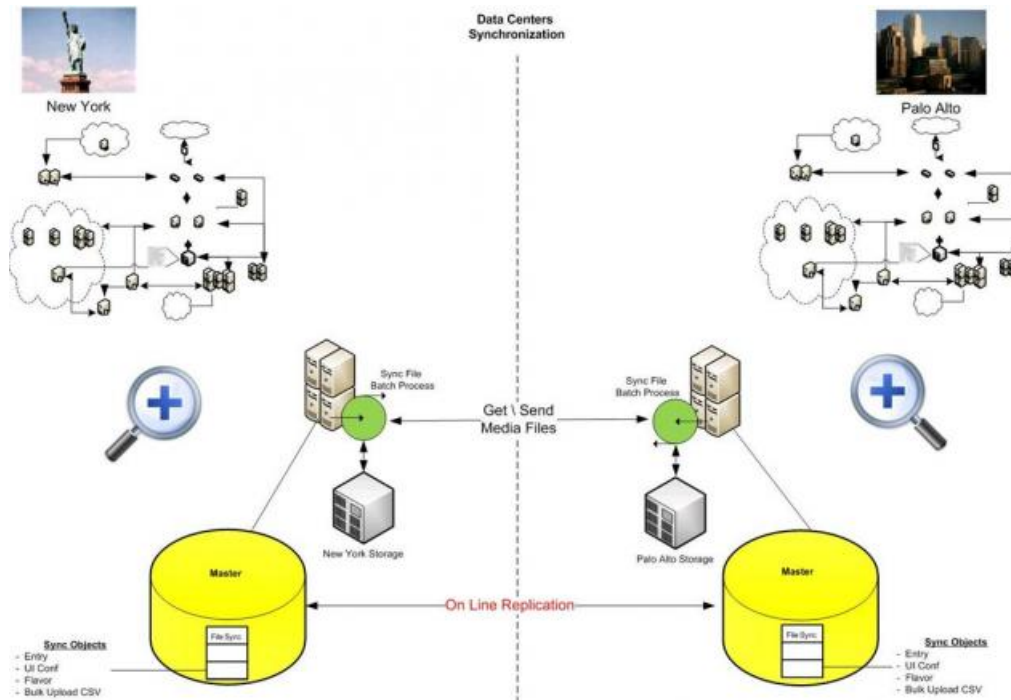
Server	Recommended Hardware Specs.	Operable Components
Few Back-End DB Server/s (master + replications)	64bit based Linux OS with at least 4GB RAM and 1 quad core CPU	Connected to Shared Storage (NFS)
Back-End DWH Server/s	64bit based Linux OS with at least 16 GB RAM and 1 quad core CPU	Video Analytics Module Connected to Shared Storage (NFS)
Site Admin Server	Linux based OS. No minimal specifications on hardware	Site Admin Module Connected to Shared Storage (NFS)
Document Conversion Server/s (optional)	Windows server 2008. With at least 2GB RAM. 1 CPU	Kaltura Document Conversion Module (if applicable)
Shared Storage Dedicated Server	64bit based Linux OS with at least 4GB RAM and 1 quad core CPU On average - 3X size of the expected video source files (assuming source should be stored as well) Storage size to be set according to the expected volume of media assets	Shared Storage

Multiple Data-Centers Site Setup

For enabling full redundancy and better performance for end users, utilization of multiple data centers may be needed. The geographic location of the data center should be chosen according to the geo- location distribution of the end users.

The synchronization of multiple data centers should be performed on both database and storage. For database synchronization a native MySQL Master-Master replication can be utilized.

For media files synchronization, a batch process can copy media items from one datacenter to the other datacenter, in order to maintain full online redundancy. The following diagram illustrates the architecture of a multi-datacenter setup of the Kaltura Platform.



Self-Streaming Considerations

In most cases, when planning for a new deployment of the Kaltura online video platform, a use of external CDN functionality is recommended. When this is not possible and/or not required, it is recommended to add additional front-end hardware to be dedicated for playback streaming.

Kaltura recommends the following addition to site setup in such cases:

Server	Recommended Hardware Specs.	Operable Components
Streaming proxy servers	2 servers with 64bit based Linux OS, at least 8GB RAM and 1 quad core CPU	Media streaming (relevant for self-streaming deployments)

Kaltura eCDN Plus is the newest addition to Kaltura’s wide range of deployment options – including on-premise, hybrid, and full SaaS deployment. It is based on the robust Kaltura API and leverages streaming technology by Wowza Media Systems as well. Kaltura eCDN Plus is tightly integrated with Kaltura’s market-leading video products, such as the company’s CorporateTube video portal, and integrations with social business platforms (SharePoint, IBM Connections, and Jive), as well as Kaltura’s products for the education market (including integrations with Blackboard, Sakai, Moodle, Desire2Learn, and Canvas) providing a seamless user experience, and centralized management and control.

Additional Scalability Considerations

Factor	Estimated Sizing Information
	<p>As a general rule of thumb, in order to have multiple flavors generated for each video entry, including mobile flavors, - storage of 3X the size of the expected video source files should be reserved (assuming source should be stored as well).</p> <p>Example of storage calculation:</p> <p>Original video clip at backup storage:</p>

Factor	Estimated Sizing Information
Storage Resources	<p>Storage Size = ~60 Mb for a one minute duration</p> <p>Normal Web Quality Flavor (transcoded files)</p> <p>Bit rate = 600 kbps</p> <p>Storage Size = 4-8 MB for a one minute duration = ~ 7.5% of Original</p> <p>High Web Quality Flavor (transcoded files)</p> <p>Bit rate = 1200 kbps</p> <p>Storage Size = 8-16 MB for a one minute duration = ~ 15% of Original</p>
Transcoding Resources	<p>Each CPU allows for 360-720 video hours of transcoding per month</p> <p>One CPU can handle one transcoding task at a time.</p>
Playback Streaming	<p>The actual streaming volume can be deduced from the single clip storage size and site usage information.</p> <p>if CDN is in use - most streaming volume is handled by CDN provider (excluding cases of long tail media type)</p>
Video Recording Streaming (if applicable)	<p>Each FMS server can hold hundreds of concurrent recording connections</p> <p>30 days * 24 hours = 720 streaming hours a month per connection</p>
Web Services	<p>One front-end server can support ~75,000 daily unique users (API calls and playback streaming to CDN when CDN is in use)</p>
Operational Database	<p>2 quad core DB Machines running in a master/slave topology can hold tens of millions of media items</p>
Video Analytics	<p>1 quad core DB Machine for storing 6 months of historical data for tens of millions of media items</p>

The following aspects should be taken into account as a basis for a scale-up decision:

Scale up Reason	Scale up Activity
Back-end server/s CPU's are repeatedly overloaded due to increase in number of transcoding operations	Move transcoding to a dedicated server or add a dedicated transcoding server

Scale up Reason	Scale up Activity
Apache web services are repeatedly overloaded	Add a dedicated front end server (and load balancing hardware if needed)
System is repeatedly slow due to db related operations, with no change after performance tuning	Move DB to a dedicated server or add a db server
System is intolerable slow during daily ETL processes or reporting executions	Move ETL or DWH& ETL to a dedicated server
Back end server/s CPU is overloaded (no increase in number of transcoding operations)	Add a dedicated back end server for batch processes