

Choosing a Cloud Encoder


Jan Ozer

www.streaminglearningcenter.com

@janozer

[jozer@mindspring.com/](mailto:jozer@mindspring.com)

[276-238-9135](tel:276-238-9135)





Questions

- For more information, check out the book
 - http://bit.ly/Ozer_multi
 - “A beacon of light in a valley of half-knowledge”
 - Mustafa Isik, on Twitter





Agenda

- What is cloud encoding
- VOD cloud encoding
 - Business models
 - How cloud encoding works
 - Who should switch
 - Choosing a VOD encoder
- Live Cloud Transcoding
 - Business models
 - How it works
 - Choosing a live cloud transoder



What is Cloud Encoding

- Someone either:
 - Builds their own encoding center and rents it to you (encoding.com)
 - Or, integrates with cloud computer vendor (Amazon, Azure) and sells service to you
- High level benefits
 - Lower CAPEX
 - Elastic scalability



Cloud Encoding Models

- Lease your own cloud, install your own encoder
 - Every encoding can be a cloud encoder
- Software as a Service (biggest category)
 - Direct (Encoding.com, Zencoder)
 - Via Amazon (Vantage)
- Platform as a Service (Elemental)

Roll You Own/Do It Yourself



- Model
 - You lease the cloud facilities as needed
 - You buy software, install and manage
- Pros
 - You have dedicated resources
 - Can run as VPN firewalled from outsiders
 - You don't buy hardware
- Cons
 - You have to run (need technical resources)
 - You buy software



Software as a Service

- Model
 - Service provider makes all cloud hardware decisions
 - You upload, they encode
 - Pricing based upon:
 - Commitment
 - Encoding speed (some)
 - SD/HD
- Vendors
 - Encoding.com, Zencoder, Amazon, Harmonics, Heywatch
- Pros
 - Simple, easy, usually cheaper
- Cons
 - You don't have dedicated resources; can't reserve capacity
 - Your files are encoded in shared resources
 - Counterpoint:
 - Most SaaS vendors are extremely fast (more later) and secure



Platform as a Service

- Model
 - You reserve (and pay for) specified resources from vendor
 - You upload, they encode
 - Pricing based upon:
 - Commitment
 - Encoding speed (some)
 - SD/HD
- Venders
 - Elemental
- Pros
 - Resources are yours; always available
 - Run as VPN
 - Can control cloud and appliance encodes from single common interface
 - Cloud is extension of appliance
- Cons
 - More expensive buy in
 - More administration required



How's it Work – VOD?

- Just like YouTube (except you get to choose the encoding profiles)
 - So, you upload your files
 - Choose your encoding profiles
 - Encode away
 - Larger customers use API, not UI
- Biggest negative
 - Upload time (for local files)
 - Reduced by upload accelerator tools like Aspera

Who Should Use Cloud - VOD?



- Who should go exclusively to the cloud?
- Who should stay on premise?
- Who should augment with the cloud?

Who Should Go Exclusively to the Cloud?



- Companies that:
 - Receive bulk of videos from third-party sources
 - Upload lag isn't a factor
 - Deliver from the cloud (either CDN or Amazon)
 - No upload lag for delivery
 - Video encoding isn't core business

Changing Math of Cloud Encoding



- 2009
 - Upload one 30 mbps mezz file to produce one 600 kbps distribution file
- 2014
 - Upload one 30 mbps mezz file to produce 24 distribution files (in my test group) with total bandwidth of 70 mbps
 - If distribute from the cloud, more efficient to encode in the cloud



Who Should Stay on Premise

- Companies with fixed and consistent encoding needs (and CAPEX)
- Why? Cost
 - Our test project (ten 50-minute videos to 24 presets) cost \$270 with Amazon
 - At 2X/week, annual cost would be ~ \$28,080 for Amazon, or around \$85,000 for three years
 - Lots of high end encoders from Elemental, Telestream, and other vendors could produce those files at much less cost



That Said

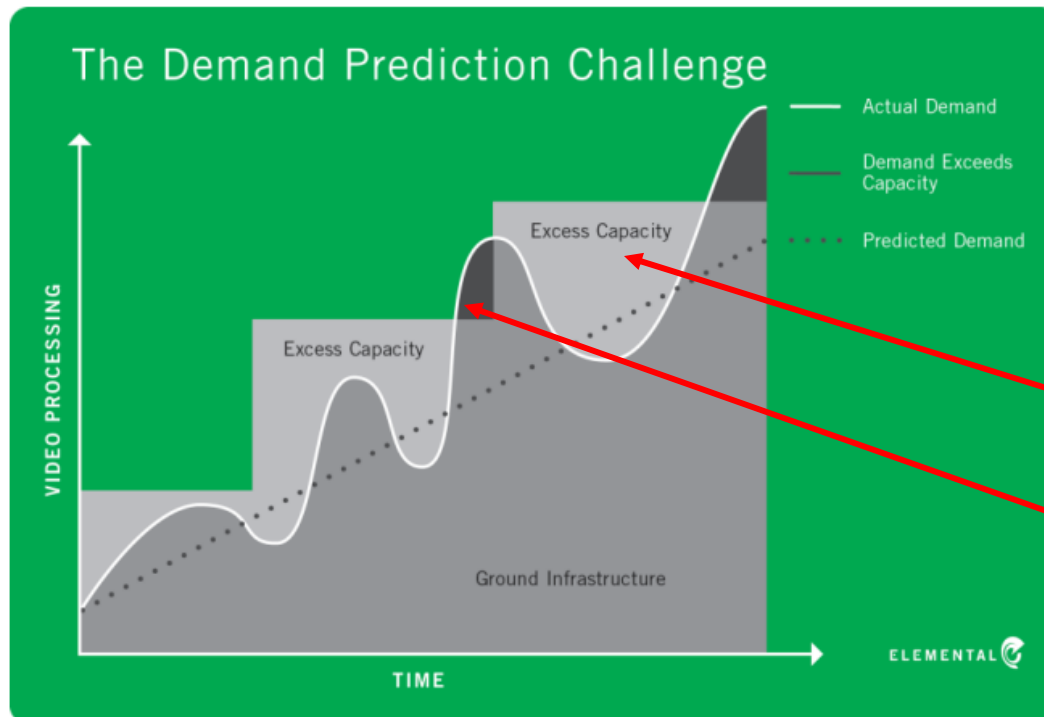
- Univision (Hispanic Broadcast Network)
 - Anvato hardware at video router captures cable broadcasts, encodes to mezzanine format and uploads to the cloud
 - Videos from distribution partners submitted in ProRes/Avid DNxHD converted to same mezz format and uploaded to cloud
 - No in-house distribution type encoding (e.g. to HLS/HDS)

http://bit.ly/ozzer_cld1



Who Should Augment with the Cloud

- Companies with fixed encoding capabilities with burst requirements



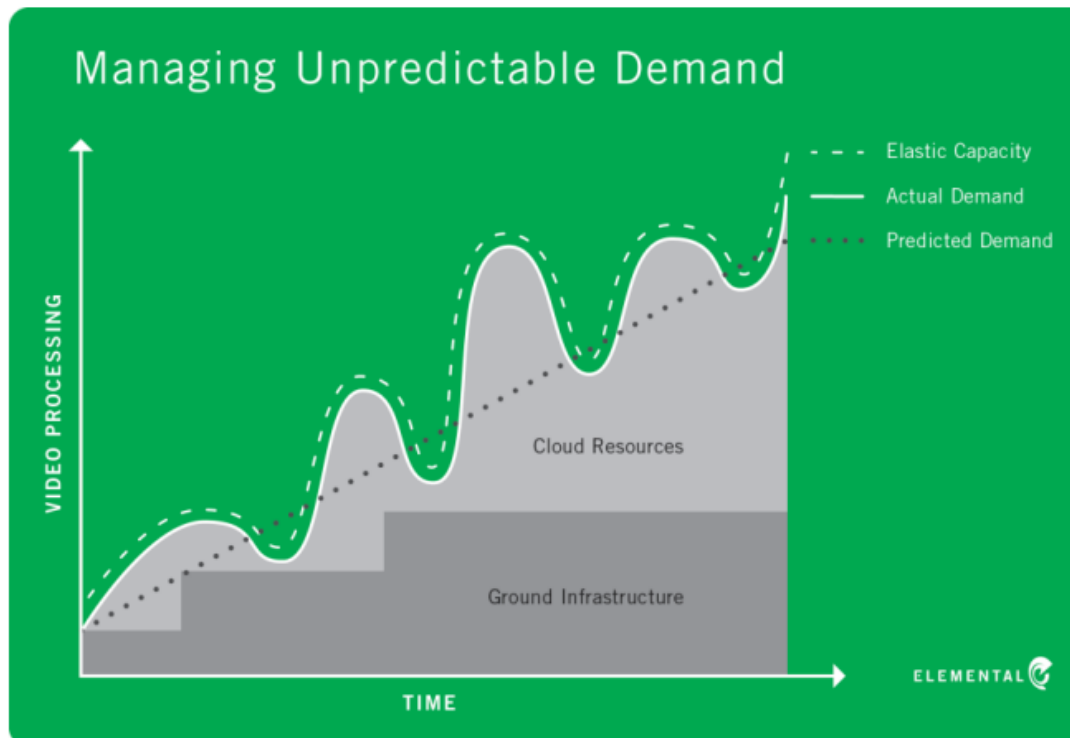
In the past, would add capacity via on-premise gear, which caused two problems:

- periods with excess capacity, money spent on gear you weren't using
- also times where capacity was insufficient



Who Should Augment with the Cloud

- Companies with fixed encoding capabilities with burst requirements



Using the cloud, you can buy what you need, saving money and never running out of capacity

Who Should Augment with the Cloud



- Scripps Network
 - Brands like HGTV, DIY Network, Food Network, Cooking Channel, Travel Channel and Great American Country
 - Consistent encoding needs, plus burst:
 - Encode library to new format (iPad, HEVC)
 - On-board new distribution partner who needs custom library

Fixed With Burst Capabilities



- T3 Media
 - Content management, delivery and monetization solutions; encodes 10,000+ files per day
 - Went to the cloud for:
 - Burst requirements (“housekeeping” encodes not billable to client)
 - Clients with very erratic demands (UGC-based shows), 2K videos one day, a dozen the next



Obtaining Burst Capabilities

- Simplest method; use cloud service provided by supplier of on-prem system (Elemental, Telestream, Harmonic)
- Why?
 - Can use same presets (so no debugging)
 - Can control all jobs from same interface



Lessons from Scripps/T3

- Both use Zencoder (not on-premise provider) and use MAM to distribute work
 - Sample of 2 doesn't disprove previous slide
- Initially, both used cloud for specific types of jobs (e.g. 3 minute videos or less)
 - Later, reconfigured infrastructure to enable all jobs on on-prem/cloud
 - Scripps can upload 45-minute mezz file to the cloud in “seconds”



Choosing a VOD Cloud Encoder

- Integration with existing encoders
- Can it get the job done?
 - Retrieval/delivery
 - Format support
 - Interface paradigm (UI/API/Both)
- Is it the best option for you?
 - Output quality
 - Performance
 - Service models
 - Service Level Agreement
 - Cost/pricing model

Integration with Existing Encoding Infrastructure



- On-prem encoders with cloud capabilities
 - Elemental Cloud
 - Harmonic ProMedia Carbon MP
 - Telestream Vantage Cloud
 - Sorenson Squeeze Server (installed in the cloud)
- Key benefits
 - Single interface for control
 - Same presets/workflows
- May not be the “best” solution, but you should consider it



Retrieval

- Can cloud encoder retrieve/deliver files?
 - Desktop folder(s), FTP, SFTP, Amazon S3, RS Cloud Files, or MS Azure storage locations
 - Amazon Elastic Encoding is to/from S3 buckets only
- Does the encoding service offer upload acceleration via:
 - Aspera or similar service
 - Multi-threaded FTP ingest



Delivery

- Delivery options – should be as extensive as input
 - Desktop folder(s), FTP, SFTP, Amazon S3, RS Cloud Files, or MS Azure storage locations
- Syndication integrations
 - YouTube
 - The Platform
 - WordPress
 - Joomla
 - Drupal



Format Support

- Can the cloud encoding service input files in the existing format
 - Amazon Elastic Encoding doesn't input ProRes (currently)
 - Don't assume, check for formats like MPEG-TS, ProRes, AvidDNxHD, Red, etc.



Output

- Single file
 - H.264
 - WebM
 - HEVC (when?)
 - MPEG-2 (for playout servers)
- Adaptive formats
 - HLS
 - HDS
 - Smooth
 - DASH
- Caption-related
 - CEA-608
 - CEA-708
 - Teletext (and other Euro standards)
 - Sidecar formats
 - SCC
 - SRT
- DRM packaging
 - DivX/Widevine
 - Playready
 - Flash Access



In General

- Most cloud encoders save Amazon offer very full format support
- Check details for captions/DRM
 - Recent check, encoding.com only cloud vendor that could package in Widevine Classic
 - Some vendors didn't offer DRM



Interface and API

- Should match technical capabilities of user
 - Low volume/non-technical
 - All about UI
 - Not all vendors offer a UI (zencoder)
 - High volume production
 - Typically 99% of production is via the API
- API – check for
 - Documentation
 - Typically XML based
 - Wrappers for other programming languages
 - PHP, Java, Cold Fusion, C#
 - Rich progress reporting




Quality

- Tested Amazon, Elemental, Encoding.com, QuickFire, Zencoder
 - Amazon (single pass only) is slightly behind (not commercially relevant)
 - All the rest are very, very close
- Quality comparison
 - 640x360@600 kbps (lowest quality in adaptive group)
 - 720p@800 kbps

Comp 1



Comp 2



Comp 3





Quality Perspective

- Many cloud vendors use x264
 - Same codec, quality should be close
- Others use their own codec (Elemental)
 - Have been competing with x264 for many years
- All companies are reputable
 - Despite claims of superior quality, there's just not that much difference
 - I wouldn't make a vendor selection based upon quality



Performance

- Two tests
 - 10~50 minute files to 24 presets (12 Flash, 12 HLS)
 - 50 30-second files to 24 presets

Presets



Preset Name	Duration	FILE FORMAT	CODEC	Profile	Level	CABAC	Width	Height	Frame Rate	Key Frame	Video Bit Rate	Maximum	Bits per pixel
Adaptive Group													
1080p_high	190	mp4	h.264	High	Auto	Yes	1920	1080	29.97	90	7000	8750	0.113
1080p_med	190	mp4	h.264	High	Auto	Yes	1920	1080	29.97	90	6250	7812.5	0.101
1080p_low	190	mp4	h.264	High	Auto	Yes	1920	1080	29.97	90	5000	6250	0.080
720p_high	190	mp4	h.264	Main	Auto	Yes	1280	720	29.97	90	4400	5500	0.159
720p_med	190	mp4	h.264	Main	Auto	Yes	1280	720	29.97	90	3600	4500	0.130
720p_low	190	mp4	h.264	Main	Auto	Yes	1280	720	29.97	90	2900	3625	0.105
540p_high	190	mp4	h.264	Main	Auto	Yes	960	540	29.97	90	2200	2750	0.142
540p_low	190	mp4	h.264	Main	Auto	Yes	960	540	29.97	90	1600	2000	0.103
360p_high	190	mp4	h.264	Baseline	Auto	No	640	360	29.97	90	1200	1500	0.174
360p_low	190	mp4	h.264	Baseline	Auto	No	640	360	29.97	90	600	750	0.087
270p	190	mp4	h.264	Baseline	Auto	No	480	270	15	90	350	437.5	0.180
234p	190	mp4	h.264	Baseline	Auto	No	416	234	10	90	200	250	0.205



Performance Notes

- Implementations vary
 - Elemental/QuickFire – automated tests supplied by vendor (instantaneous trigger of all tests)
 - Encoding.com/Zencoder – XML provided by vendors; manual copy and paste to create jobs
 - Amazon – I created all jobs via UI and manually entered

Comparisons



Total - 10 ~40 minute	4:37:12	1:48:12	1:53:04	3:01:44	0:23:46
Job entry time	6 minutes	4 minutes	4 minutes	0 minutes	0 minutes
Net encoding time	4:31:12	1:44	1:49:04	3:01	23:46
50 30-second files	NA	9:37	19:07	19:22	4:41
Job entry time		6 minutes	2 minutes	0 minutes	0 minutes
Net encoding time	NA	~ 3:37	~17:07	19:22	4:41



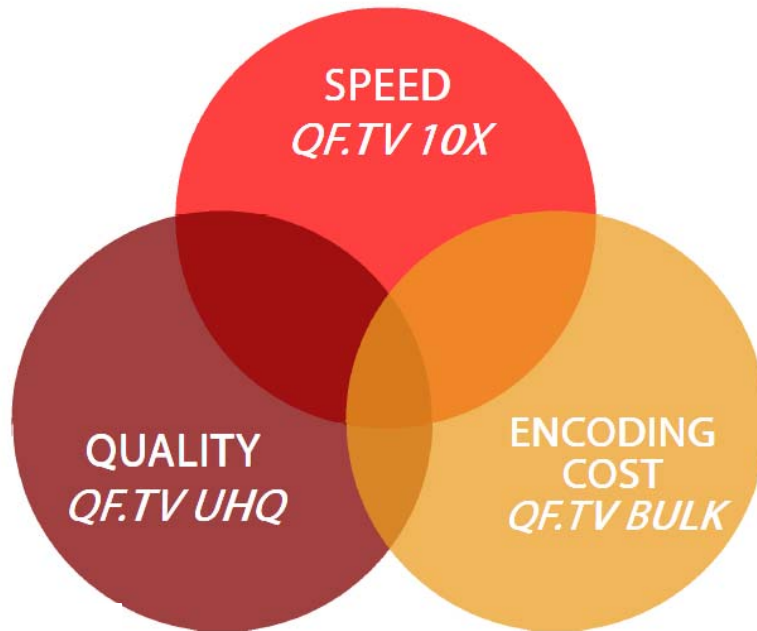
Performance Conclusions

- Performance notes
 - QuickFire very interesting new technology
 - Amazon – OK for companies where encoding time isn't that important
 - No SLA
 - Zencoder/Encoding.com – very close on longer files
 - Elemental – could have boosted performance by throwing more cores at the job
- Overall, faster is always better, but differences here are probably not a dispositive factor in most decisions

Service Models – QuickFire



Three Offerings for Different Pain Points



- ⊙ QF.TV 10X Encoding Service
 - > Encode at 10X realtime with same quality as other services
- ⊙ QF.TV Ultra High Quality Mode
 - > 10-30% Higher Quality Encodes
 - > Save money on CDN delivery or improve user experience
- ⊙ QF.TV Bulk
 - > Bulk encoding service at $< \frac{1}{2}$ cost of other services



Pricing Model

Vendor	Amazon	Zencoder	Encoding/ Heywatch	Elemental	Quick Fire
Model	Cost/ minute	Cost/per minute	Cost per GB I/O	Based on GPU/CPU instances, usage and commitment	Multiple models <ul style="list-style-type: none"> • Bulk .5¢/min • 1X speed 1¢/min • Ultra HQ – 3¢/min • High Speed 5X or 10X I/O dependent
Prices	Varies by region	Varies with commitment 1.25¢ - 5¢ /minute			
Details	SD – 1.5¢ /min HD – 3 ¢ /min	Transmux is 25% of standard pricing			



Bottom Line

- Apple-to-apples pricing comparisons are tough to formulate
 - Variables differ from company to company
 - All – monthly commitment will reduce price
 - Price per GB – size of input and outputs
 - Price per minute – minutes of output
 - Other
 - Transmux vs. transcode
 - Encoding modes/queue times



Overall

- ID architectural factors
 - Does the system need to integrate with existing encoder
- ID inputs and outputs
 - Make sure vendor can:
 - Input your source; deliver required outputs, particularly DRM
 - Retrieve and deliver files as necessary
- ID interface requirements
 - UI vs. API
 - Robustness and utility of API
 - Integration requirements
- Gather data for pricing estimates and estimate pricing
- Run some performance trials at different times
 - Make sure performance is reasonable

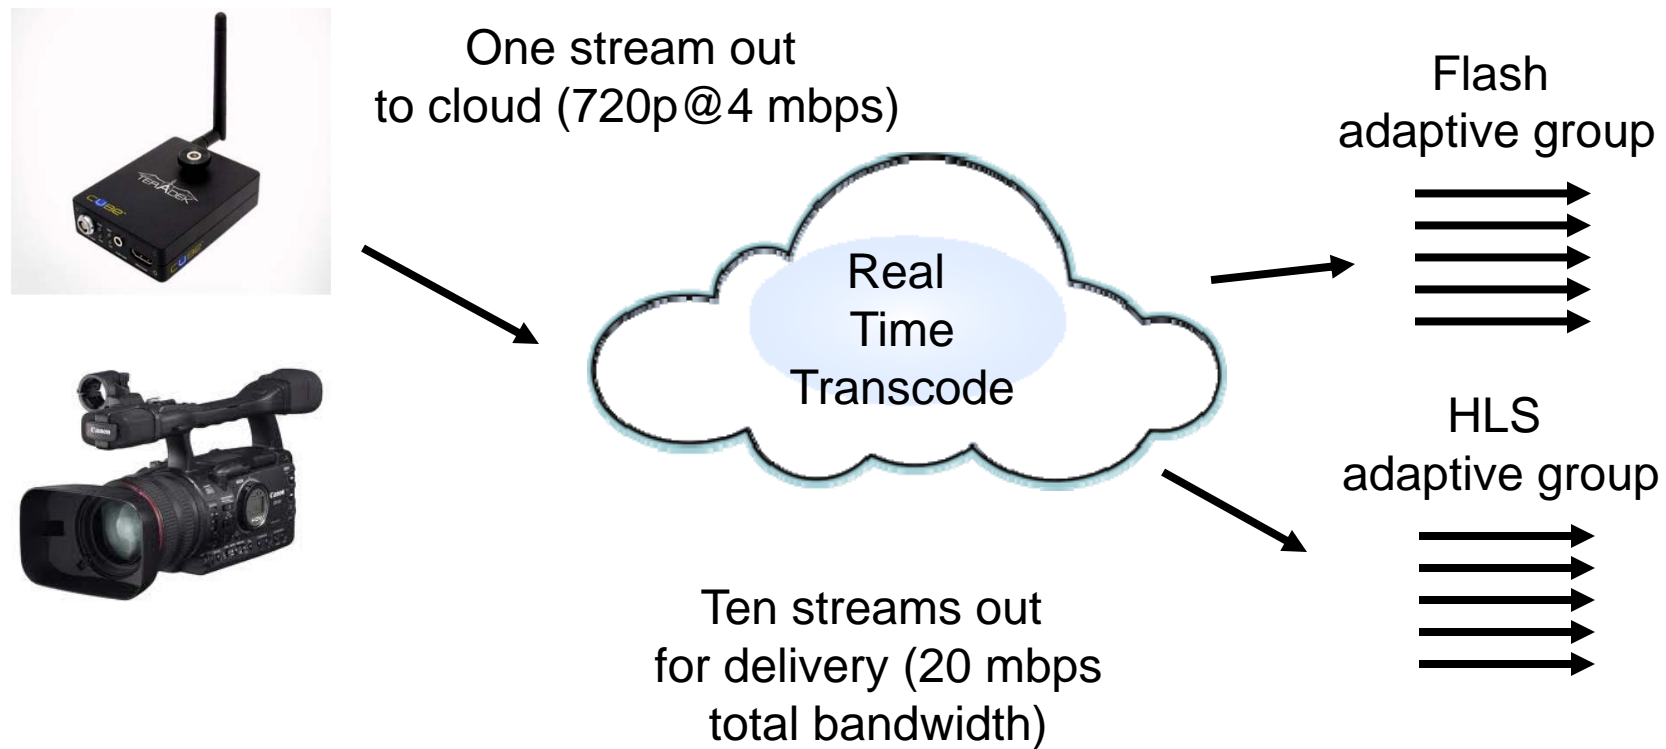


Choosing a Live Transcoder

- Workflow
- Key benefits
- Use cases
- Business models



Live Transcoding Workflow



Key Benefits – Live Transcode



- Traditional benefits
 - Lower CAPEX/Scalability
 - Don't need an encoder that can produce 10 streams simultaneously
- Reduces outbound bandwidth requirements
 - Encode to 10 streams on site, need 20 mbps outbound bandwidth
 - Encode to one stream, transcode in the cloud, need 4 mbps outbound bandwidth (or less)



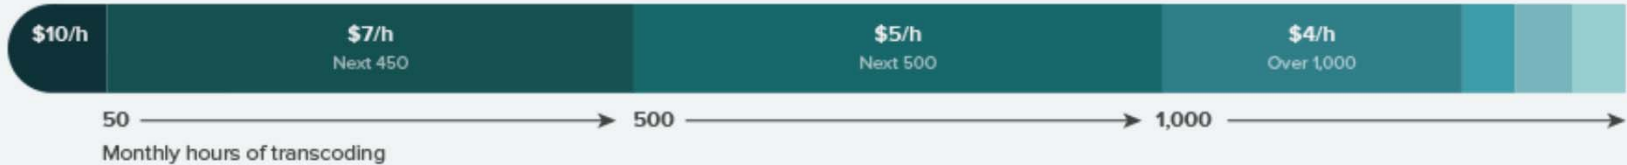
Use Cases

- Companies with infrequent events with multiple streams that would require expensive encoder
 - Live transcoding limits CAPEX
- Companies broadcasting from sites with limited outbound bandwidth
 - Live transcoding limits output bandwidth requirements
 - Could save having to bring satellite truck in to broadcast event



Pricing Models


- Zencoder



HOW ARE LIVE TRANSCODING HOURS CALCULATED?

Both input* and output streams are used to determine hours used. For example, a one hour event encoded in 4 renditions (2 RTMP + 2 HLS) counts as 5 hours of encoding (1 input + 4 output).





Pricing Models

- Livetranscoding.com

Plans & Pricing							
	PRO Multi \$1299 /MONTH	PRO Single \$599 /MONTH	Flash \$499 /MONTH	iOS \$499 /MONTH	MS Smooth Streaming \$499 /MONTH	Windows Media \$499 /MONTH	REQUEST TRIAL 24-Hours FREE
Features	SIGN UP	SIGN UP	SIGN UP	SIGN UP	SIGN UP	SIGN UP	SIGN UP
Multi-bitrate	✓	•	✓	✓	✓	✓	✓
Video input source	1	1	1	1	1	1	1
Video outputs							
Adobe Flash RTMP Streaming	3	1	3	•	•	•	1
Apple HTTP Live Streaming	3 (+1 audio)	1 (+1 audio)	•	3 (+1 audio)	•	•	2
MS Live Smooth Streaming	3	1	•	•	3	•	1
Windows Media Video	3	1	•	•	•	3	2
Devices							
Desktop (Windows, Apple OSX, Linux)	✓	✓	✓	•	✓	✓	✓
Apple devices (iPhone, iPad, iPod)	✓	✓	•	✓	•	•	✓
Android 2.2+	✓	✓	✓	•	•	•	✓
Windows Phone 7	✓	✓	•	•	✓	✓	✓



Business Models

- Roll Your Own
 - Buy Wowza Transcoder, install in the cloud
 - Send stream in; distribute streams as desired
- SaaS
 - Zencoder/Elemental as standalone service provider
 - Send stream in/they send multiple streams out
 - Also integrated component of Brightcove Live
 - Review of service coming up in Streaming Media Mag



Choosing a Live Transcoder

- Market fairly nascent, still developing
 - Wowza obvious choice for DIY
 - SaaS
 - If VOD provider also does live (Elemental, Zencoder) this makes decision easy



Choosing a Live Transcoder

- Identify stream requirements
 - Formats/target players
 - Stream count
 - DRM
- ID output points
 - Ensure encoder compatibility
 - Ensure service can connect to your outputs
- Check pricing and commitment levels
- New market so:
 - Check references
 - Run at least one trial event



Live Cloud Transcoders

- Zencoder
- Elemental
- Livetranscoding.com
- Streamguys
- Vantrix



Resources

- The Cloud Vs. On-Premises Encoding Dilemma, http://bit.ly/ozер_cld1
- Review: Brightcove Video Cloud Live, http://bit.ly/ozер_cld2
- Elemental Cloud: Streaming Media's Comprehensive Hands-On Review, http://bit.ly/ozер_cld3
- Encoding.com: Streaming Media's Comprehensive Hands-On Review, http://bit.ly/ozер_cld4
- How to Webcast with YouTube Live, http://bit.ly/ozер_cld5

Questions?

