

Research^{2.0}

Technology Analysis for Business & Investment

Thought Leader Interview with Jules Urbach, Founder & CEO

OTOY

with Stephen Waite, September 20, 2010

In this Thought Leader Interview, we sit down with Jules Urbach, Founder & CEO of OTOY. OTOY is an emerging cloud computing company that you will likely be hearing much more about in the months ahead. Some of the key takeaways from our interview with Jules include:

- Clients use OTOY to move their applications, no matter how large or complex, to the cloud. With OTOY, console quality video games or huge CAD applications show up on a web page and are instantly usable.
- A lot of the ‘magic’ in OTOY comes from how the company uses GPUs. Successfully virtualizing these high end applications makes OTOY the “VMware of entertainment applications.”
- OTOY, OnLive and Gaikai leverage a GPU on a server to render content. Each company takes a different approach to how it encodes and virtualizes content.
- OTOY encodes ORBX streams on a GPU - the same GPUs that are used to render the content they are serving. OnLive uses an h.264 encoder chip on a separate piece of hardware to grab the video output of each GPU running on their servers. Gaikai uses a CPU to encode h.264.
- We have yet to see a fully baked virtual world really take off on the web - something like Second Life, for example, with film quality visuals; a vast, persistent and connected world, running on Facebook’s social graph. This would be a compelling use of OTOY’s technology.

We hope you enjoy this interview with Jules Urbach of OTOY.

SW: Thanks for your time today, Jules. It's great to have a chance to speak with you. For those readers who aren't yet familiar with OTOY, can you give us a short description of the company and its mission in the marketplace?

JU: OTOY is a cloud computing company. We provide server side rendering and computing solutions that leverage GPUs (graphics processing units) rather than CPUs, which traditionally power today's web servers. Our clients use OTOY to move their applications, no matter how large or complex, to the cloud. Typically, you wouldn't expect to see a console quality video game or a huge CAD application (such as SolidWorks) show up on a web page and be instantly usable. We make that possible.

SW: What was the inspiration behind starting the company?

JU: I started my career programming video games. I loved the classic Laser Disc games (like Dragon's Lair) when I was a kid. The graphics in those games may have been pre-rendered, and the interactivity limited, but they offered a taste of the future: video games that looked just as good as a movie. In one form or another, my work has been inspired by seeing that promise fulfilled. I was accepted to Harvard in part because I recreated Dragon's Lair and sent them the source code. I never ended up going, because I felt the path I was on was best explored on my own.

Around 1996, after finishing up a 3D pool game for a video game publisher, I decided that the web was the future for games and software, not boxed products. I rebuilt the pool game, and got it running in a web browser. Anyone could play it in a URL through a web browser. That was a big breakthrough at the time. In July 1999, Macromedia (now Adobe) used that game to launch their first content site (shockwave.com). They also tried to buy my company, which we passed on. Around 2004, I improved our rendering technology to the point where it was good enough to be used in films. That helped prove the case for rendering content server side. From there, OTOY directly evolved into the company it is today.

SW: OTOY seems like a poster child for Arthur C. Clark's observation that any sufficiently advanced technology is indistinguishable from magic. Can you give us some insight into the advanced technology that lies behind OTOY?

JU: A lot of the 'magic' in OTOY comes from how we use GPUs. The tools for GPU programming are still in their infancy. Developers, even today, have a hard time tapping into the raw power that this hardware offers. Around 2006, we were able to get two key technologies working on the GPU: ray tracing (used in movies for realistic visuals) and high speed compression. We have been perfecting both these technologies over the last few years.

SW: What makes OTOY's video compression technology unique?

JU: We built the codec ourselves, specifically to run on GPUs, and to encode 3D models and other complex media streams in addition to video. Most of the digital video you see online, or view on devices such as BluRay players or iPhones, is encoded with h.264. h.264 has become ubiquitous because of its terrific quality compared to previous generation video codecs like MPEG-2 (used on DVDs). But you cannot encode h.264 on a GPU without sacrificing quality. Many of the improved algorithms in h.264, like CABAC, are serial operations. These do not lend themselves to a GPU's parallel programming model.

ORBX, OTOY's codec, is designed to deliver better than h.264 high profile quality, while encoding multiple HD streams at high speeds (>1000 fps) on an off-the-shelf GPU. This is important on a platform like ours, where supporting thousands of concurrent users is a necessity. ORBX will also be leveraging CPU/GPU hybrid chips as they mature. This hardware is just now coming to market, both from Intel and AMD. There are new benefits that can be derived from these chips. We won't lose the parallelism of the GPU, but we can safely introduce serial operations where it makes sense.

SW: Progressive downloading is a cool feature of OTOY's technology. Tell us about it and explain why it is important.

JU: OTOY is not limited to screen scraping a server. The OTOY client supports progressive downloading of media, assets, arbitrary code and even tiny virtual machines and file systems. This is a core component of the system. The ORBX video decoder is running in that same sandbox.

At GDC we presented World of Warcraft and other games running inside a web browser using this technology. These games were streaming instantly from a server, but they were being started on the client, as if they had been already installed on the local machine. Further, we could switch between rendering the game 100% on the server and rendering it on the client (once enough of the game had been downloaded). This means that server side rendering only needs to be used when it is absolutely necessary. The OTOY client is really not much different from a web browser in this scenario. We cache content and use the local client's graphics capabilities when possible. That reduces strain on the servers, and means zero latency when we render local content.

SW: We're seeing a flurry of activity and interest in 3D technology today and, at the same time, rising penetration of broadband communications, the proliferation of smart mobile devices, and TVs connected to the internet. Could you imagine a better time to launch OTOY's technology into the market?

JU: I agree that this is the right time to bring this technology to market. Software and games are quickly following music, movies and TV to the cloud. We've seen the first wave of this transition in the form of Apple's app store and the Facebook platform.

I would argue that you don't even need server side rendering to move most of today's software to the cloud. When 20-30 Mbs connections become mainstream, the difference between accessing a Blu-ray drive on your PC or a copy of the disc on a remote server won't be that noticeable. Any app that runs off your hard disk will eventually run just fine in the cloud. We used to access our CD-ROMs at 1.5 Mbs - which is now the low end of broadband connections in the U.S. DVDs stream at 6-9 Mbs, and that is getting more commonplace on DSL or cable modem connections. When we reach 20+ Mbs, I believe that will be the nail in the coffin of physical media as we know it today. I don't see a new disc format superseding Blu-ray before that happens.

On the flipside, at 20+ Mbs, we also have enough bandwidth to deliver uncompromised server side rendering that won't look much different than what you can render locally. We can encode and stream ultra HD at around 15-25 Mbs with ORBX. That is 4 concurrent 1080p Blue-ray streams. We can add 3D stereoscopic rendering with only 10-20% more overhead.

SW: OTOY has a disruptive technology that is viewed by some analysts as a game changer. What kind of reception and interest are you getting from software developers, computer, mobile phone and TV manufacturers these days?

JU: The majority of video game publishers, chip makers (e.g. Intel, AMD, Nvidia) and software publishers agree that cloud-based streaming is poised to replace packaged software sales (and eventually ‘download and install’ solutions as well). When you talk to companies like Apple, which have a walled garden around their platform, the response is more nuanced. Obviously, a service like ours might disrupt their ecosystem. When we first approached them about streaming applications on iOS devices, we weren’t sure what to expect. Their response was interesting: wrap up each application in its own individual app, ensure the developer creates a UI that is designed for the iPhone/iPad, and limit server side rendering to portions of the game or app that cannot be reasonably rendered locally.

That is actually pretty fair. I don’t think we’re going to face a tougher gauntlet than Apple. The game publishers that want to use OTOY to publish on iOS also agree with Apple’s approach (vs. throwing a PC game ‘as-is’ onto a server and streaming it to an iPhone). So when you see OTOY content running on iOS, it will likely be in this form. Apple fascinates me. If they wanted to they could have moved iTunes and the app store to an entirely browser based model. It is ironic that they stopped short of that, given that Apple has moved mountains to make the browser a viable application delivery mechanism through WebKit, Canvas (now embraced by Microsoft in IE9) and later WebGL. You almost have to wonder if the first plan for the app store wasn’t to build it on top of Safari.

Google is going to take this concept and see it through to its logical conclusion through Native Client. Native Client is what Flash and Java ought to have been, and what Apple’s app store might have been. Its purpose is to safely extend the browser with the same capabilities and performance characteristics that developers get from writing low level C or assembly code. You can write your code in the language of your choice, and get native speed. That is huge. That does not mean that it will become a standard (like Apple’s Canvas extensions). But, in my opinion, it should be given that chance when it gets out beta. Why not? It is open source, and unencumbered.

For our part, the ORBX decoder will run in Native Client. That means we won’t have to install or download anything on Google TV, Google Chrome browsers, Chrome OS or Android devices. That is how software development should be. Philosophically, that is also where our heads are at. We want to help developers get their software to end users with as few barriers as possible. It was surprising to see Mozilla back away from Native Client in favor of JavaScript. I believe it was a setback for developers. The day when we can run ORBX in JavaScript, I’ll change my mind.

SW: We’re seeing on-line games take off on social networking platforms. What kind of opportunities do you see for OTOY in the social networking space?

JU: Zynga’s success validates the purchasing power of the casual user on the web. It does so in a way that would have been difficult to imagine 5 years ago. There’s still a lot of room left unexplored in this space. We have yet to see a fully baked virtual world really take off on the web. By that, I mean something like Second Life, for example, with film quality visuals; a vast, persistent and connected world, running on Facebook’s social graph. That would be a compelling use of our technology. We already have developers exploring these kinds of applications on top of OTOY.

SW: Gaming in the cloud is coming of age and there are several companies in the space today. What distinguishes OTOY from its competitors, such as OnLive?

JU: I have a tremendous amount of respect for what OnLive has accomplished. I ran into Steve Perlman last year at the BMO conference. We had a friendly talk about where we thought cloud rendering was heading. I don't see Steve's goals as being all that different than ours. We both want to see unlimited rendering power in the hands of developers in the cloud.

OTOY, OnLive and Gaikai leverage a GPU on a server to render content. This is a primary distinction between our services and previous generation cloud offerings. After all, if you want to render something on a server without a GPU, Amazon EC2 and similar services can do that for you already.

So we are all aiming to provide GPU-style rendering for high end graphics in the cloud. That much we have in common. But how we encode and virtualize this content is where we begin to diverge from one another. OTOY encodes ORBX streams on a GPU - the same GPUs that are used to render the content we are serving. Onlive uses an h.264 encoder chip, on a separate piece of hardware, to grab the video output of each GPU running on their servers. Gaikai uses a CPU to encode h.264.

There is nothing preventing us from encoding h.264 streams on the CPU (or GPU) or adding a custom ASIC for encoding. We've tested our platform with those configurations. From an engineering perspective, any of these solutions would have been easier than creating a new GPU codec from scratch. ORBX has given us an order of magnitude more users per server than any alternative we've explored. That is why we're sticking with it.

We can run 48 first person shooters at 60 fps on a single 1U server through ORBX. That is with legacy games that have not been optimized for our service (i.e. games that we run out of the box, without any modifications). When developers target our platform (through tools such as our raytracing pipeline), concurrent usage gets closer to 100 users per GPU. That is before you factor in local rendering power offloaded on the client.

From a business perspective, we are not primarily a portal. We license our technology so that others can build or enhance their own software on top of our services. That is not the only valid approach one could take in this space. But from our perspective, it is the best way to satisfy the needs of the market we are serving. To that end, we are offering two types of licenses. The first is similar to Amazon EC2, where we host your application in a sandbox on our servers. The other is like a VMWare site license: you run OTOY's software stack on your own servers (purchased from OEM partners like SuperMicro).

SW: The rise of GPGPU cloud services like PEER 1 seems ideally suited to OTOY's technology. How do you envision OTOY leveraging these kinds of services in the future?

JU: We plan to support any vendor that offers a reasonably sized GPU cluster. PEER 1's service is just the tip of the iceberg. When Amazon EC2 adds GPGPU support, we expect we would build an OTOY AWS stack that could run on EC2's ecosystem. The same would be true for similar services from Google, Microsoft and others.

SW: OTOY has been working closely with AMD. What are the major advantages of AMD's technology relative to Nvidia and Intel?

JU: We were very deliberate in choosing to go down this path with AMD. We tested early versions of ORBX on Nvidia GPUs, x86 CPUs, and AMD GPUs. We settled on CAL as our core development platform (CAL is AMD's low level computing language). It was very challenging to program a GPU using CAL, which is not officially supported by AMD. But we were seeing amazing speeds which we could not replicate on other architectures.

As our company has evolved, so have our relationships with other major hardware vendors. This year, we've added Intel and Nvidia as partners. More will follow. We announced Intel this summer. Obviously, we need CPUs on our servers as well as GPUs. And, in that respect, Intel has a very compelling offering. Intel is also developing hardware cards made from densely packed x86 cores which we may use in the future.

We are officially announcing our partnership with Nvidia in a few weeks. We have been working with them on a version of ORBX that will be deployed on Nvidia hardware in 2011. This is not trivial, given that ORBX's speed has, up until now, come from functionality that is specific to AMD hardware. But, from a practical perspective, we would be ignoring a significant portion of the professional graphics market if we didn't support CUDA applications. Adobe Photoshop, CS5, and countless other apps only support CUDA.

SW: Dassault Systèmes' 3D modeling and design program SolidWorks is leveraging OTOY's capabilities. Do you see a lot of opportunity for OTOY's technology in the enterprise?

JU: Absolutely! It is a market we take seriously. From our perspective, it is bigger than games. Enterprise applications require a much different approach to virtualization than video games. We have to support numerous 3rd party technologies that may interact in opaque ways. This is much more complex than hooking the display, sound and input for a video game.

SW: Speaking of 3D modeling, your LightStage technology is one of the most innovative we've come across in that space. Being able to render life-like images in real time seems to blow a lot of people away. Tell us about the technology.

JU: LightStage is at the center of our graphics ecosystem. It is used by Sony, Weta, Digital Domain, ILM and others for high end visual effects. LightStage captures real life actors, scenes and objects with perfect fidelity in 3D. It is unbiased, and in that respect is about as close as you can get to what you see in real life.

A LightStage capture is like a holographic recording of a person or object - except we capture an extra set of dimensions beyond what is typically stored in a hologram. Specifically, we are capturing lightfields, which can be used to relight and manipulate an object, not just view it from any angle. I am very proud of the progress our team has made this year with LightStage. Tim Hawkins (LightStage's CTO) and Paul Debevec (LightStage's Chief Scientist) recently received an Academy Award for their work with LightStage in *Avatar* and *Benjamin Button*.

SW: What are your future plans for LightStage?

JU: LightStage has been used mostly for high end visual effects up until now. But we are working to extend LightStage technology to devices and cameras that can be used by anyone. On the software side, LightStage data is big - up to 4 Terabytes per capture. That is something we are addressing as well. We are using ORBX to store, render and stream live LightStage data from the cloud. That is in fact a major design goal of the codec. Digitizing the world is an inescapable component of modern media. What the video camera is to YouTube, LightStage is to OTOY. It is difficult to emphasize just how important it is from a graphics perspective.

SW: One last question for you today, Jules. Looking out over the next 3-5 years, where do see OTOY?

JU: We have a lot of goals for OTOY. Ultimately, we want to break down the barriers between content creators and consumers. It is not going to happen in a single step. We know what the key milestones will be. Getting any application to stream to a thin client (like a web browser) is where we begin. Enabling content that matches the fidelity of what the human eye can see is where we go next.

We also need to spread the computing power bound inside a GPU beyond graphics. We have started down this path with compression, compilers and physics frameworks for the GPU. AI, for example, is an area that is ripe for further development. I believe we can empower developers with a much greater range of tools, and perhaps convert more casual users into creators themselves. One way of enabling that is through more flexible natural language parsers and programming tool chains. Just as I see rendering moving towards an unbiased model that becomes as simple as photography, I can imagine high performance computing and software development becoming equally democratized.

SW: Thanks again for your time today, Jules. It was terrific speaking with you. We wish you and your colleagues at OTOY all the best in the future.

About Jules Urbach



Jules Urbach is the Founder and CEO of OTOY and LightStage. Jules was born and raised in the French countryside west of Paris (near Versailles) until age 9 when he moved to LA. In LA he attended Harvard-Westlake high school before being accepted to Harvard University. However, he decided to defer his acceptance to Harvard (indefinitely, as it turned out) to make video games. He made his first game, Hell Cab - Time Warner Interactive, at age 18. Six years after he created Hell Cab, he founded Groove Alliance. Groove ended up creating the first game ever available on Shockwave.com (Real Pool). A decade later, Jules is busy working on his latest two companies, OTOY and LightStage, which he hopes will revolutionize 3D rendering.

For more information about OTOY, go to: <http://www.otoy.com/>

Previous Thought Leaders:

Think Twice: An Interview with Michael Mauboussin, Chief Investment Officer, Legg Mason Capital Management (November 2009).

Into the Multiverse: An Interview with Corey Bridges, Co-founder, The Multiverse Network (December 2009).

2D-3D: An Interview with Craig Summers, Founder and CEO, 2D-3D Video Inc. (January 2010).

Global Solar Technology: An Interview with Dr. Alan Rae, Technical Director (February 2010).

China: An Interview with Arthur Kroeber, Founder, Dragonomics Research and Advisory (March 2010).

Web 3D: An Interview with Anita Havele, Executive Director, Web 3D Consortium (April 2010).

3D Technology: An Interview with Ami Dror, Co-founder and Chief Strategy Officer, XpanD (May 2010).

OpenIPO: An Interview with David Marcus and James Donohue, WR Hambrecht & Co (June 2010).

GPU Computing: An Interview with Ujesh Desai, VP of Product Marketing, NVIDIA (July/August 2010).

Research^{2.0}

ABOUT US

Research 2.0 specializes in positioning emerging technologies and companies in a business and investment context for institutional investors. The two founders each have over 30 years of experience in applying emerging technology, company research, equity capital markets and investment management.

Our clients consist of technology companies and institutional investors that rely on in-depth, independent, thematic research and advisory services.

Feel free to visit our website (www.research2zero.com) to learn more about us and [sign up for our free research update service](#).

Our main office is located at 1313 Washington St., #326, Boston MA 02118 U.S.A. and we have staff in New York, Connecticut and Paris, France. Our main telephone number is +1-617-848-9651.

OUR TEAM

Kris Tuttle, Founder & CEO, kris@research2zero.com

Stephen Waite, Partner and Director of Strategy, steve@research2zero.com

Leigh Deacon, Administrative Manager, leigh@research2zero.com

Benjamin Stein, Associate Analyst, ben@research2zero.com

Heidi Manwaring, Research Assistant, heidi@research2zero.com

COPYRIGHT

It is a violation of federal copyright law to reproduce or distribute any part of this publication to anyone by any means without permission. Institutional and corporate clients have internal redistribution rights and in some cases may distribute this publication externally to investors, suppliers and customers with written authorization.

DISCLAIMER

The content of this publication is intended for informational purposes only. Research 2.0 is not a broker-dealer, investment bank or registered investment advisor. We do not make recommendations as to whether to buy or sell securities nor do we solicit orders in publicly traded securities. Proper investment decisions can only be made with a complete understanding of return expectations, risk tolerance and duration. Research 2.0 does undertake to run a "Research 2.0 Model Portfolio" which contains stocks that we publish research on. Research 2.0 employees are allowed to buy, sell and own securities in companies covered subject to certain compliance rules.