Board game review: Ultimate Werewolf Legacy

reader comments

16 with 14 posters participating

Share this story

Welcome to Ars Cardboard, our weekend look at tabletop games! Check out our complete board gaming coverage at cardboard.arstechnica.com.

Our 16 games of *Werewolf* sprawled across 20 hours and two lengthy play sessions. They began well enough, with enthusiastic people enjoying each other's company, keen to backstab, betray, and devour their fellow participants. Villagers—and the occasional werewolf—were hanged, and each person's hands were bloodied. Yes, this was the decade-old social deduction game we all knew well—but now with sealed boxes, fistfuls of stickers, and a huge leather tome for the moderator to scribble in.

Ultimate Werewolf Legacy takes an old concept and pairs it with newfangled "legacy" game mechanisms. This means components are permanently altered—mostly the moderator's diary—and decisions are made that impact future plays. In other words, it's a campaign game with irreversible decisions, promising all the drama that premise entails.

Unfortunately, it doesn't deliver.

The furry devil is in the details

If you've played *Werewolf* before, you know what to expect from this social deduction game. You need a minimum of eight players in addition to a dedicated moderator who runs the game—but the experience is better with more, ideally with a group of 12-16.

Each session consists of a chapter of play, with each chapter being broken down into three individual games. In each game you're dealt new roles—some of you are voracious (but hidden) werewolves, while the rest are villagers full of suspicion. Play is broken into day and night phases, with the entire group attempting to suss out the villains by day and the werewolves devouring a villager at night.

The campaign consists of five such chapters, each offering a glimpse of story courtesy of the game's leather diary. The narrative itself follows a series of strange occurrences around a New England town in the late 1600s. This is the heart of the game, where players make group decisions that alter the village's path and provide consequences. Or, at least, that's the idea.

The implicit promise of a legacy game is that choices have permanence. We want our decisions to make waves that ripples

throughout the campaign. *Ultimate Werewolf Legacy* loses sight of this principle. It has moments where you think something special is really going to happen. Some of the time it will even trick you into thinking *this* particular outcome will shape things drastically. (The players don't know most of this, however, as all of the stickering, note-taking, and branching paths are recorded solely by the moderator.) Feedback to players is obscured.

The single benefit to this design philosophy is increased replayability. It would be entirely reasonable to play through a campaign and then grab another copy of the game (or the recharge pack from the publisher's website) and play it again. You would indeed see new roles and a couple of small changes.

But overall, the legacy elements here are a bit of a whiff due to their softness. Most player choices only have a very short reach. Some of the game's best moments have these wild unpredictable narrative decision points full of promise—but they usually offer some bright new ability or radical option only for the next game. Then they're gone like a horror in the dark.

For the sake of avoiding spoilers, let me fabricate an example. Before beginning play, the moderator (well, the diary) might offer a couple sentences of story concerning a villager accused of witchcraft. As a group, players must vote to cast the villager out or to declare her innocent. The outcome either adds or removes a new role to next game—although no one knows this when voting. Major decisions such as these seem like they should matter beyond modifying a single role for a single play, but that's not the case.

Even the excellent use of props is often undermined by limiting the fallout to one session. Perhaps worse, the vast majority of stickering and branching pathways alter the role selection for the next chapter only. This is almost always based on the number of team villager wins in the previous set

of three scenarios. In this way, the game seeks to balance play over the long haul through incremental shifts, adjusting as it goes. An experienced moderator hoping for some agency in crafting the role-set will find no love here, as their position is reduced to directing traffic and reading the small slices of story that are scattered throughout the game (and not particularly well written).

The single mechanism that *does* extend over the length of the campaign is the family system. In addition to a hidden role, each player possesses a family member card which groups the player with a neighbor or two sitting nearby. This has some interesting implications for early voting, but it more importantly facilitates the "Most Valuable Villager" (MVV) system.

- This game has plenty of role cards!
- The "family" cards.
- The game's narrative "diary."
- The introductory narrative…
- The game's rules booklets.
- Inside the diary.

The MVV is determined by different parameters in each game but generally is awarded to the player who makes the most significant contribution. Once a chapter concludes, the family with the most awarded members receives stickers. These offer special powers of various usefulness, with the most interesting arriving late in the game. Unfortunately, this system feels a little thin as it's detached from the narrative. In fact, most of play remains distinct from the story that's occurring since you can't actually impact it in a significant way.

While the game's full potential is not realized, it's not a total loss. This release would serve a group well that is perhaps inexperienced with *Werewolf*. It offers role-sets that

are roughly balanced, although not always terribly interesting, and it guides you through play with a few interesting mechanical twists and a memorable moment or two.

It also handles players dropping in and out between chapters quite astutely. (This is the positive aspect of decisions having a short yield and the narrative not boasting any extreme twists.)

Finally, the complexity of the game is introduced in small chunks. It’s not a heavy or opaque set of rules, but the unique structure and odd nature of the diary make for a bit of a hurdle. A preface chapter ensures that everyone is on the same page by offering solid explanations throughout play.

Assembling the pack

One of the fundamental qualities that separates this release from other legacy designs is the dedicated group size. While individual players can drift in and out, you will want a core fellowship of individuals committed to the journey. This is rough.

We all know that coordinating a dozen people’s schedules is one degree away from rocket science. The need to do this multiple times is a huge barrier, which will probably fall on the moderator—and there’s a high likelihood the group may tucker out and give up on the thing before completion.

In this regard, *Ultimate Werewolf Legacy* works best as a product aimed squarely at a younger audience. A group of college students in a dorm wanting to kill time is the ideal player base. The sheer length of the campaign precludes even convention play, as you could spend an entire day and still not finish.

The real point of failure here lies in the game’s unwillingness to commit. It wants to simultaneously support

short-term play with people dropping in and out, while also making promises of long-term decisions with serious consequence. Investing in one approach might have worked better.

This is an undeniably ambitious release. While I think it has a modest audience, the design ultimately struggles to deliver on its premise. Much like the single paragraph of story that wraps up the end of the campaign, nearly everything here feels both understated and hollow.

https://arstechnica.com/?p=1479069

Valve Software dreams of analyzing your brainwaves to tailor in-game rewards

Enlarge / Valve Software's Mike Ambinder offers a joking photo of what people think his job as Principal Experimental Psychologist looks like. Valve co-founder Gabe Newell was not on hand to confirm or deny Valve's use of power tools on his head.
Sam Machkovech

reader comments

74 with 66 posters participating

Share this story

GDC 2019

View more stories

SAN FRANCISCO—Valve Software's famously "flat" structure means most of its game-making staffers have vague titles. One of the few exceptions is its Principal Experimental Psychologist, who presented a futuristic gaming vision at this year's Game Developers Conference—in particular, he made a few peculiar admissions about how Valve might one day study your brain activity in the middle of a game and what the company might do with it.

Before speaking, Valve Software's Mike Ambinder laid out a very loud disclaimer about GDC's "vision" track of panels: "This is supposed to be speculative," he said. "This is one possible direction things could go." Even with that caveat in mind, Ambinder's choice of details is interesting to sink our teeth into, especially coming from a company that seems to offer more speculation about the future of gaming than it does actual applications of it (i.e. new games).

The slot machine of your mind?

The above and below images of Ambinder goofing off with Valve co-founder Gabe Newell weren't just for yuks: "Every talk I've given, this reliably gets a laugh. Think about that. What if we could elicit reliable reactions [from video games] and determine we were doing so?"

- Management, Valve style.
  Sam Machkovech
- A variety of existing headsets and mechanisms used to measure brainwave activity and synaptic responses.

- Ambinder shows one off.
- Newell tests another.
- Examples of existing, relatively affordable brainwave sensing systems.
- Valve’s dreams of how those signals could affect gameplay, in a confusing chart.

Ambinder compared a game player’s use of a controller to an average conversation. In his analogy, when you talk to a person, your words and gestures are equivalent to pressing controller buttons: loud and clear signals. But in terms of subtler signals like facial cues, “video games leave us out of the non-verbal part of a conversation,” he said.

As a result, games face interface limitations like memory, Ambinder argues. A player may memorize as many as 20-100 combinations of mouse-and-keyboard commands, he says. That may seem like plenty for some games, “but that is still a limitation.” And the exact amount of time that it takes for a synapse to reach a finger can average around 100 milliseconds, he estimates. “What if we could shave as much as 20-30 milliseconds off that time?” he asked.

Existing hardware, if money wasn’t a limit, could allow game makers to track everything from synaptic responses to “galvanic skin response,” from eyes’ gaze to muscle tension and posture. Many of Ambinder’s suggestions for what a game maker might do with this data sound like heart-rate sensor experiments, such as adjusting difficulty based on a player’s feelings at a given moment. We’ve seen Valve (and other studios) talk about the gameplay-tweaking potential of eye-tracking in previous years.

But one suggestion in particular raised our alarms: adjusting virtual goodies in a game on the fly. “We can figure out what kinds of rewards you like, and the kinds you don’t,” Ambinder suggested, potentially based on the physiological responses a player might have from getting loot. He didn’t talk to the

very severe privacy implications of this feedback loop, however, nor about the abuse potential for having a game pump players with loot-driven endorphins at the moment they might start getting bored. (Slot machine and loot box mechanics are already decried for artificially toying with player expectations to hook them longer.)

“Make you like the Predator”

Other ideas, on the other hand, sounded positive, if not absolutely insane. For instance, this level of connectivity could potentially help blind, deaf, and otherwise sense-deprived people perceive a video game in new ways. This would require transcranial magnetic stimulation (TMS), a particularly wild kind of treatment that is already being used to help sufferers of insomnia, PTSD, and other conditions.

“Can we make you see in infrared like the Predator? Give you access to spatial location data, or even echolocation? Could we add a sense? Improve your sense of touch? Help you notice new frequencies?” Ambinder continued. “Taste and smell things you’ve never tasted or smelled? Focus your attention? Stimulate certain areas of your brain to recruit neurons for other tasks? Help you hold more memory at once? improve memory retrieval?”

Remember that “speculative” caveat from earlier? There you go. That particular hypothetical scenario would require TMS systems, and those are admittedly much further out as a possibility than the head-mounted, synapse-monitoring systems that Newell was seen wearing in the above photos.

Of course, the near-future possibilities of these tracking systems are also limited by the amount of head- and skin-interaction required. Such connectivity would require both purchasing a bulky headset and players choosing to wear it. But Valve does have at least one path to getting players to adopt such a crazy scenario. After asking about its

feasibility, Ambinder gestured to a photo of himself sporting an HTC Vive VR headset. “Not everyone plays VR, but VR gives you semi-consistent contact with a source of brain activity,” he said. It’s an interesting thing for Ambinder to point out on Valve’s behalf as rumors of a Valve-produced VR headset mount.

Until this sort of system becomes more commonplace, Valve clearly has practical work to do to either create hardware or wait for more game makers to buy into this ultra-connected theory, that games would offer an “improved, qualitatively better player experience if developers had access to states, emotions, cognition, and decisions.” In the meantime, now—not later—might be the time to start talking about the privacy implications inherent in “games as a service” that read and respond to our brainwaves.

https://arstechnica.com/?p=1479021

Jordan Peele’s Us should cement his status as a master of modern horror

Enlarge / Lupita Nyong’o stars as Adelaide Wilson, whose family encounters their own evil *doppelgängers* in Jordan Peele’s new horror film, *Us*.

reader comments

47 with 35 posters participating

Share this story

A family is terrorized by their own doppelgängers while vacationing in Santa Cruz in Jordan Peele's new film, *Us*. With its spot-on writing and pacing and fantastic performances from its ensemble cast, the film should cement Peele's status as a master of modern horror.

(Mildest of spoilers below, because anything more would spoil the fun.)

Us is the much-anticipated follow-up to *Get Out*, Peele's surprise box office hit that earned more than $250 million and snagged Peele an Oscar for best original screenplay—the first time the award has gone to a black recipient. *Get Out* is a subtle exploration of racial tensions that quietly builds to reveal its horrifying premise and inevitably bloody conclusion. In *Us*, the theme isn't so much racial tension—it's exploring, in Peele's words, the myriad ways in which "we are our own worst enemies."

Adelaide (Lupita Nyong'o) and Gabe (Winston Duke) Wilson take their two kids (Shahadi Wright Joseph and Evan Alex) to the beach to relax and unwind with friends (one of whom is played by Elizabeth Moss). Their young son, Jason, encounters a strange figure on the beach, and that night four more mysterious figures show up at the Wilson's summer house. They break in, and that's when the real terror begins—those figures are creepy *doppelgängers* of each family member. (Peele calls them "the Tethered.")

- Elisabeth Moss plays Kitty Tyler. She and her husband join the Wilsons for a day at the beach in Santa Cruz.
- Did something sinister follow the Wilsons home from the beach?
- Strange figures dressed in red and holding hands.
- "They're us." Dark, twisted mirror images.
- Jason (Evan Alex) with Pluto (also Evan Alex), his

Tethered doppelgänger.
- Kill it with fire!
- Adelaide Wilson (Lupita Nyong’o) will stop at nothing to protect her family.

It’s a classic kill-or-be-killed horror scenario for much of the film’s second act, with plenty of jump scares and just enough snippets of humor to keep things from getting too heavy. (I loved the various family members arguing about who had the highest kill rate.) But Peele could never just leave it at that; he’s always got a few more tricks up his sleeve. This time the twist involves a network of abandoned underground tunnels, a secret laboratory, Jeremiah 11:11, and a demented take on the 1986 campaign to battle world hunger, Hands Across America. To say anything else would spoil the fun.

The hot takes and think pieces are already starting to emerge about what *Us* ultimately “means,” especially to analyze the final twist at the end (telegraphed early on for the eagle-eyed viewer). Is the film a metaphor for the darker, violent aspects of our individual natures? Or, does it say something about “the unseen underbelly lurking beneath every sunny American dream”? Maybe the film’s a story about ascending from your own personal hell or the stigma of mental illness? Can it be an elaborate commentary on capitalism, slavery, and rising up to overthrow one’s oppressors? If none of that rings true for you, Peele’s film can certainly be enjoyed as top-notch horror entertainment, with no need for a deeper meaning.

> *“The feeling that we all feel we are the good guy in our own story prevents us from facing our demons.”*

Take your pick. That’s the advantage of Peele’s decision to keep things just ambiguous enough to leave room for lots of personal interpretation and speculation. What is *Us* ultimately about? Well, what do you think it’s about? It acts as a mirror for the audience. That aspect may frustrate some viewers who

like a neatly packaged ending, but I found that the uncertainty and unanswered questions kept me reflecting on what I'd just seen for much longer than the average film. The questions enrich the overall experience.

That's not to say Peele doesn't have his own broad take on his movie's themes. "The feeling that we all feel we are the good guy in our own story prevents us from facing our demons," he recently told the *Los Angeles Times*. "I wanted to make a movie that allows everybody to face their demons, in whatever faction you want to put this movie through the prism of. But as a starting point … the United States and our xenophobia was the front and center idea to grapple with."

Like *Get Out*, this is horror that makes you think, and it's already doing brisk business at the box office. That said, *Us* doesn't quite scale the artistic heights of its Oscar-winning predecessor. In particular, its ambitious larger premise starts to strain one's willing suspension of disbelief to the breaking point in the third act. But ultimately, this is still a fine, satisfying, highly creative piece of filmmaking that makes us all the more keen to see what Peele does with his forthcoming *Twilight Zone* reboot on CBS All Access.

[embedded content]
Trailer for Jordan Peele's *Us*.
https://arstechnica.com/?p=1478397

A dev trained robots to generate "garbage" slot

machine games—and made $50K

Enlarge / Two indie devs explain how they used automation, a single Google Play account, and a single slot-machine template to create and distribute over 1,000 slot machine apps.
Alex Schwarz

reader comments

62 with 53 posters participating, including story author

Share this story

GDC 2019

View more stories

SAN FRANCISCO—This year's Game Developers Conference saw two game makers emerge with a possible chapter in a future dystopian sci-fi novel: the story of making money by letting robots do the work. In their case, that work was the procedural generation of smartphone games.

A single "game jam" event led to a data machine that ultimately pumped out a decent amount of cash: $50,000 over a couple of years. Years later, with that data (and money) in hand, the makers of this game-making machine, which focused entirely on "garbage" free-to-play slot machines, used GDC as a wake-up call to an industry where the "right" messages often revolve around listening to players, sidling up to publishers, and racking up critical acclaim. In their case, eschewing all of that worked a little too well for their comfort level.

Winning the “race to the bottom”

In 2013, two video game makers had been trying for years to make it in the burgeoning mobile games space. One of them, Alex Schwartz, had helped get the solid mobile swiping-action game *Jack Lumber* off the ground. (In a past life, I gave that game a good review at the now defunct tablet-only magazine *The Daily*.) The other, Ziba Scott, had put together a fine mobile-friendly puzzle game, *Girls Like Robots*.

Both games operated in a pay-once, play-forever model without microtransactions. Both attracted awards, recognition, good expo showings, and publishers. Both failed to take off.

They looked at the meager income they were making doing it the “right” way, as had been established by the old publishing guard. They then looked at iOS and Google Play marketplaces and saw that “freewares, clones, and junk” dominate the general selection, let alone the actual money-making charts.

In one casual chat about the sheer weight of that business reality, the duo came to a conclusion: “We could do better… at doing worse!”

- One example of how the auto-generated slot machines looked on Android devices. Missing here is the text-to-speech singing about curling that would play upon any jackpot.
- By using Selenium, the duo could automate the clicking, typing, and file selection needed to upload each of its games automatically.
- Early stats.
- Surprise: these things are making money.
- Some of the takeaways that came at the experiment’s end.

They teamed up during the 2013 Global Game Jam to push something out that resembled the “race to the bottom” they saw on mobile platforms. Thanks to the time-restricted nature of a

game jam, they opted to buy a 3D slot machine asset off of the Unity Store (a marketplace that lets game makers pay modelers and animators for unrestricted use of various 2D and 3D assets) for $15. They then spent the rest of the jam creating a system that would automatically generate the rest of the skinning needed to make this basic virtual slot machine *just* unique enough to be published as its own smartphone app.

“Let’s customize these like other slot machine companies do,” Schwartz said. “They make themed slots. What’s the minimum set of things to change to make a different slot machine? Let’s change the title. Change the one image on the reel that might be relevant to your topic. So, like, a dolphin slot: put a dolphin in there [as the jackpot slot logo] with a special icon. Then the background is a scrolling dolphin image.”

The original word list was hand-curated based on what the team thought was interesting but also generic and safe-for-work. The most scintillating name they went for at that point was “3D Sexy Librarian Slots.”

With the visuals knocked out, the duo went one step further: creating custom music.

“A crappy song would play, then use text-to-speech to sing the word ‘dolphin’ in the Google Translate monotone voice. It’d play that every time you won. It’d say the name of your game in the music,” Schwartz said.

As a result, with the press of a single button, a Unity script could put those steps together and essentially auto-generate hundreds of “custom” slot machines. Schwartz and Scott confirmed that their automated system’s scraping of public images exposed one issue: Google Image Search would throw up errors for exceeding the rate limit. “We found a use for Bing,” Schwartz said in a phone interview with Ars. “Its image search had a number of things that were looser. I’m not trying to knock them, but they have a reputation for being second

class. That felt like a kindred spirit for what we were trying to achieve here.”

“A portal to a better world”

With that slot-creation template set, the team automated the process of feeding information to Google Play (a much easier marketplace to exploit than iOS at the time) and creating publicly available freeware slot machine apps with ads. One simple Selenium script later, and that process was done.

The duo could feed a single slot-machine keyword into their combined scripts, which took “a few hours” in all to build, then watch on a “ghost monitor” as its system faked like a human, clicked every appropriate checkbox, picked every country, agreed to every terms-of-service agreement, and filled in every appropriate text box—then took the auto-generated slot machine and uploaded it for anyone to play.

> *“The slots didn’t have much to do with bowling.”*

They attached mobile ad network Playhaven to the whole thing because the duo’s philosophy was that they never wanted to take actual money from users who would download their bizarrely named apps. They then “walked away” for two months. After that period of dealing with real-life work, they peeked at their income and advertising statement and were stunned: people were downloading their apps, and 27 percent of those people were clicking on their ads, driving roughly $211 of ad revenue per day.

The team came up with a theory: “All of our advertising keywords were related to casino related content,” Schwartz said to Ars. “We had an epiphany: our game looks so fucking terrible, but people downloaded it for some reason. When they see an ad for a much better slot machine or casino, they click it because… of course you do! That’s a greener pasture! A way better future you could be having! We think the quality was so

low in our shit that the ads were a portal to a better world."

Yet the duo incredulously admits that its average rating for many of the apps was in the four-star range and that reviews were quite kind. One review stood out to Scott, for the auto-generated "3D Bowling Slots" app: "Someone wrote that they were disappointed that the slots didn't have much to do with bowling."

Supervillain origin story?

- Quite a bit of money by the end.
- One of the duo's favorite auto-generated slot machines.
- The automation system minded Google Play's upload limit of 15 apps a day. Their system ran automatically for *quite* some time to get up to roughly 1,500 uploaded apps.
- A slot-machine template that the team ultimately discarded.
- This is as "popular" as the duo's wares became.

The engineering half of their brains wanted to see how far this enterprise could take them. So they began tinkering with the existing template with things like the automation of slot-machine descriptions.

Google Trends seemed like a good idea, but that usually led to trademarked or public-figure names, and the team wasn't interested in getting out of the "automation" loop by having to deal with a high volume of takedown requests. So they opted for a slightly updated app-name template: the word "3D," plus an adjective, plus either an animal, location, or country, plus the word "free."

Examples included 3D Tremendous Face Pain Slots, 3D Rough Elbow Slots, 3D Mild Dogwood Slots, 3D Viceroy Butterfly Slots, and 3D Inexperienced Great Horned Owl Slots. (They

eventually made T-shirts to commemorate the latter.)

Schwartz and Scott also paid a small Romanian studio a pittance to build a higher fidelity slot machine, which they eventually discarded. That happened in part because the duo's mix of newer full-time work and ethical concerns crowded out their excitement and availability.

> *"You have the worst users."*

"Someone said, you could raise money on this idea, or sell this data to someone else, or sell your company," Scott said to Ars. "We were at a crossroads where the joke was similar to the origin story of a supervillain. Do we abandon all creative pursuits to make the most intense money-making slot-creating enterprise? Or does this continue being a tiny background of 1/20 of our day?"

Eventually, the headaches of keeping up with Google Play caught up to the team. Apps were removed for violating an updated terms of service that gave Google more leeway to cut out apparent crapware. Google also updated the Web interface on a somewhat regular basis. Moving a single box a few pixels could throw a wrench into the Selenium robo-clicking works, which the team had previously designed to auto-upload 15 apps a day (Google Play's upload limit for a single developer account at the time).

At one point, the app network Playhaven called the duo with a flat declaration. "We're seeing erratic data on your account," Schwartz said to paraphrase. "We're not sure what's up. We're not interested in continuing to serve ads to your slot machines. But we want to be clear: you didn't break our ToS. You're just, I don't know, inconvenient."

Playhaven then added, "You have the worst users. People who come from your apps don't spend money." The team switched to another ad provider immediately, Chartboost. "They knew roughly what we were doing," Schwartz said. "They've been

great.”

“Optimized to remove our content”

This many years later, all of the roughly 1,500 apps generated by this experiment are dead. And Schwartz and Scott think of the whole affair as a mix of a joke and a productive wake-up call.

“Our half-joking argument: by offering the largest target of low-quality garbage apps, these marketplaces became optimized to remove our content,” Schwartz said to Ars. “You could almost say that our company trained their algorithm so that what we were doing could eventually not be possible. But we were the first to bring it to that level.”

In their GDC presentation, Schwartz mostly left the data up for people to peruse in a laughing manner, but also with a human angle about the costs that might arise if you chase such a robo-generated app dream with profit, not laughs, as the goal.

“This whole project was an itch that felt so good to scratch,” Schwartz said at GDC. “We thumbed our noses at the mobile market that had broken our hearts in 2013. We laughed the whole way… The truth is I don’t give a fuck about slot machines. We’ve moved on. If there’s a business lesson in that, I’d say: scratch your crazy itches, give that insane idea a try. But even if it *does* work out… be prepared to walk away because it may only have been the attempt that made you happy, not the result.”

https://arstechnica.com/?p=1477805

Latest trailer for John Wick 3: Parabellum is sheer guns-and-glory mayhem

reader comments

44 with 35 posters participating

Share this story

[embedded content]
Keanu Reeves gives us a *Matrix* callback in latest *John Wick: Chapter 3—Parabellum* trailer.
Fresh on the heels of the announcement that *Bill and Ted 3: Face the Music* will start shooting this summer, we get a new trailer for another Keanu Reeves-starring vehicle: *John Wick: Chapter 3—Parabellum*.

(Spoilers for first two films below.)

For those who missed the first two movies, John Wick (Reeves) is a legendary hitman (known as "Baba Yaga") who tried to retire when he fell in love and got married. Unfortunately, he's drawn back into the dark underground world by an act of senseless violence after his wife's death. Nothing will stop John Wick from seeking retribution. The first *John Wick* grossed more than $88 million worldwide for a film that cost around $30 million to make, and it was praised for its brisk pace, heart-stopping action sequences, and stylish *noir* feel.

- John Wick, excommunicado.

YouTube/Lionsgate

- An over-confident assassin.
  YouTube/Lionsgate
- “If you think you can take John Wick, you’ve got a nasty surprise coming.”
  YouTube/Lionsgate
- Anjelica Huston plays the director of the High Table. They want John dead.
  YouTube/Lionsgate
- Channeling Neo: “I need guns. Lots of guns.”
  YouTube/Lionsgate
- Winston definitely has the guns.
  YouTube/Lionsgate
- Halle Berry plays a reluctant ally.
  YouTube/Lionsgate
- Wick in action.
  YouTube/Lionsgate
- Take that guy out, Halle.
  YouTube/Lionsgate
- Wick has a new nemesis who’s been longing to meet him.
  YouTube/Lionsgate
- A fight to the death, no doubt.
  YouTube/Lionsgate
- They have no idea what’s coming.

So naturally, there was a sequel. In *John Wick: Chapter Two*, Wick takes on a whole slew of Russian baddies but finds himself honor-bound (per the professional hitman’s code) to go on one last assignment. Of course he succeeds—he’s John Wick—but the client, Italian crime lord Santino D’Antonio (Riccardo Scamarcio), tries to cover his tracks by posting a bounty on Wick’s head. Wick ends up killing Santino inside the Intercontinental Hotel, which is supposed to be a safe space for assassins, and is declared “excommunicato,” cutting him off from all the resources the hitman underground provides.

And that’s where *John Wick: Chapter 3—Parabellum* picks up,

with Wick scrambling to find allies (and ammo) to fight off pretty much the entire underworld of assassins, all after a $13 million bounty is put on his head. The first trailer dropped in January, a juggernaut of action set against an ominous countdown to replicate the one-hour head start Wick has been given before his excommunication goes into effect.

> *“You have no idea what’s coming.”*

In the new trailer, we get even more action and tantalizing glimpses of the fascinating underground culture of assassins, with their gold coins, markers, and almost medieval code of honor—an element that has contributed a great deal to the franchise’s lasting success. There’s even a call-out to *The Matrix* (yet another of Reeves’ iconic roles) when Intercontinental owner Winston (Ian MacShane) asks what he needs. Wick responds, “Guns. Lots of guns.”

New cast members include Anjelica Huston as The Director of the crime lord High Table—”There’s no escape for you”—and Halle Berry as fellow assassin (and dog lover) Sofia, a reluctant ally who owes Wick a debt. “After this, we are less than even,” she says in the new trailer. And Wick faces a formidable Chinese assassin, who admits he’s been looking forward to meeting him for ages. “So far you haven’t disappointed.” He rarely does. As Winston says in a portentous voiceover, “You have no idea what’s coming.”

John Wick: Chapter 3—Parabellum opens May 17, 2019.

Listing image by YouTube/Lionsgate

https://arstechnica.com/?p=1478025

It’s been a long wait, but Bill and Ted 3: Face the Music is happening

reader comments

61 with 56 posters participating

Share this story

[embedded content]
“Be excellent!” Alex Winter and Keanu Reeves reprise their iconic roles next summer in *Bill and Ted 3: Face the Music*.
John Wick 3 isn’t even out yet, but Keanu Reeves is already onto his next project, reprising another of his iconic roles: a third film in the Bill and Ted franchise that has long been desired by fans. And it’s been announced via Twitter in a quintessentially Bill and Ted way: with stars Alex Winter and Reeves—collectively Wyld Stallyns (the name of their band)—standing in front of the bandshell at the Hollywood Bowl (“where we will never play”) and thanking fans for their support.

The third film will be called *Bill and Ted 3: Face the Music*, and it will hit theaters August 21, 2020, if the *Bill and Ted 3* Twitter account is accurate. Naturally, time travel will be involved. Per an accompanying press release:

> *Following 1989’s Bill & Ted’s Excellent Adventure and 1991’s Bill & Ted’s Bogus Journey, the stakes are higher than ever for William “Bill” S. Preston Esq. (Winter) and Theodore “Ted” Logan (Reeves). Yet to fulfill their rock-and-roll destiny, the now middle-aged best friends set out on a new*

> *adventure, when a visitor from the future warns them that only their song can save life as we know it and bring harmony to the universe. Along the way, they will be helped by their families, old friends, and a few music legends.*

The film will be directed by Dean Parisot of *Galaxy Quest* and will reunite the screenwriters of the two previous films, Ed Solomon and Chris Matheson (who is incidentally the son of *I Am Legend* scribe Richard Matheson).

Listing image by YouTube/Orion Pictures

https://arstechnica.com/?p=1477595

“Energizing Times”: Microsoft to “go big” at E3 in response to Google Stadia

Enlarge / This controller attachment was shown in Microsoft’s Xcloud promo video, and Bluetooth wireless controller support is also planned.

reader comments

61 with 43 posters participating

Share this story

Microsoft announced its Xcloud game-streaming service last August, with the ambition of streaming console-quality games

to gamers wherever they are—on their tablets, smartphones, PCs or even consoles. Yesterday, Google joined the streaming gaming fray with its announcement of Google Stadia, one-upping Redmond by offering the assembled press limited hands-on access to Stadia games.

Google promises that Stadia will be "coming 2019," potentially stealing a march on Xcloud, which is due only to enter public trials this year. But in an internal email sent to rally the troops, Phil Spencer, Microsoft's gaming chief, seemed unsurprised and apparently unconcerned.

Spencer wrote that Google "went big" with its Stadia announcement, but Microsoft will have its chance to do that, too: he promised that the company will "go big" with its E3 presentation and raft of announcements. He also said that Google's launch endorsed Microsoft's decision to launch its streaming service and said that Microsoft offered all the key elements Google identified—"Content, Community, and Cloud"—but that ultimately, "it's all about execution."

While the broad strokes of Google's announcement didn't surprise Spencer, he did note that he was impressed by certain elements: the "leveraging of YouTube, the use of Google Assistant and the new WiFi controller." Google has been building up YouTube as a live streaming, gaming-friendly platform, though last year it decided to kill off the standalone YouTube Gaming app and roll gaming features into the regular YouTube apps and site.

Microsoft's gaming community is built around its Mixer streaming platform and Xbox Live. Mixer retains its focus on gaming and includes unique ways for streamers to interact with their audiences: viewers can vote on in-game actions to help or hinder the streamer; pick and choose which stats, leaderboards, and overlays they can see; and even be given control of the gamepad. As of the end of last year, Xbox Live had 64 million monthly active users, all with their

achievements and friends lists, and Microsoft is about to vastly expand the reach of Xbox Live with its new support for Nintendo’s Switch, iOS, and Android. Thanks to the friends lists and other features, Xbox Live arguably already represents a stronger community offering than Google has with YouTube, and if Xbox Live’s cross-platform expansion is successful, this advantage can only grow further.

Enlarge / Microsoft included this map of Azure data centers in its announcement of Xcloud. These are the locations to which the service could be deployed, but don’t expect it all to happen at once.

Microsoft’s cloud might also have an important advantage over Google’s. Perhaps the most critical single element of these streaming services is their latency; if latency is too high and the game is too slow to respond to controller input, certain classes of games, such as first person shooters, become unusable. Technology is one element to achieving low latency, with both companies investing in development of low-latency video compression, but another element is simple distance between a player and the data center the game is being streamed from. As parts of its Azure build-out, Microsoft boasts more cloud data centers in more parts of the world than either Google or Amazon. This should mean it can

offer low-latency streaming to a wider audience than Google.

https://arstechnica.com/?p=1477125

An astronaut with PTSD loses her cool in first trailer for Lucy in the Sky

reader comments

58 with 33 posters participating

Share this story

[embedded content]
Natalie Portman stars as an astronaut who starts to unravel after returning from space in *Lucy in the Sky*.
A female astronaut returns to Earth and starts a downward psychological spiral in the first trailer for *Lucy in the Sky*, a forthcoming film from Fox Searchlight Pictures.

Lucy in the Sky tells the story of married astronaut Lucy Cola (Portman), who has an affair with fellow astronaut Mark Goodwin (Jon Hamm). When he dumps her for another woman in the program, she begins to lose her grip on reality. Originally titled *Pale Blue Dot*, the movie is the feature-film debut for director Noah Hawley, whose TV credits include *Fargo* and *Legion*. Reese Witherspoon was initially considered for the lead role of Lucy but dropped out to shoot the second season

of *Big Little Lies*. Portman came aboard instead, making this at least her second song-titled film alongside *Jane Got a Gun*.

The film is loosely based on the real-world case of NASA astronaut and US Naval officer Lisa Nowak, who became involved with fellow astronaut William Oefelein in 2004 after his divorce. The affair lasted a couple of years, until Nowak discovered her lover had taken up with an Air Force engineer named Colleen Shipman. In February 2007, Nowak drove from Houston to Orlando International Airport with a car full of kidnapping gear (including an 8-inch folding knife) and confronted Shipman in her car in the airport parking lot. Nowak copped a guilty plea in 2009 and received two years' probation; she also received an "other than honorable" discharge from the Navy.

The case notably raised questions about the psychological effects of long-term space missions, in addition to how well NASA screens its astronauts for mental stability, themes that clearly inform *Lucy in the Sky*. That's earned the film some advance criticism for its stated premise from retired astronaut Marsha Ivins about how Hollywood misrepresents female astronauts: "Hollywood may consider spaceflight traumatic, and certainly our mothers do, but astronauts do not," she wrote for *Time* in 2017.

Granted, the film was still being cast when Ivins wrote her article, so the criticism was premature. But I'm not sure there's anything in this trailer that will put her concerns to rest. We first see Lucy during a spacewalk, mesmerized by the sight of Earth at night, until her fellow astronaut (Goodwin) tells her, "Time to wrap it up. We're going home." Then we see her back on Earth and get a series of flashbacks—her husband and kids, the beginnings of her affair with Goodwin—interspersed with a voiceover counting down.

"I just feel a little off," Lucy admits. "You go up there and see the whole universe, and everything here just seems so

small.” Goodwin commiserates, noting that everything she saw in space “blew your mind, and now nothing else makes sense.” But hey, fleeting glimpse of a pistol aside, she’s probably fine. Right?

Lucy in the Sky hits theaters later this year.

Listing image by YouTube/Fox Searchlight Pictures

https://arstechnica.com/?p=1476523

Hands on with Google Stadia: It works, but is that enough?

reader comments

83 with 65 posters participating

Share this story

GDC 2019

View more stories

SAN FRANCISCO—Shortly after Google announced its upcoming Stadia streaming platform this morning at the Game Developers Conference, the company opened up a few kiosks showing off the technology in a corner of a Moscone Center West hallway. Unfortunately, these extremely limited demos didn’t answer

most of the burning questions that Google has still left unanswered about it much-hyped platform.

The most popular demo, judging by the crowds gathered around the screen, was an opportunity to play *Assassin's Creed Odyssey* on a standard Chromebook via Stadia streaming. The game—running at apparently native resolution and 60 frames per second on a 1080p display—was for all intents and purposes indistinguishable from a local copy running on a high-end gaming rig. Playing and watching the games for a few minutes, I didn't notice any of the input delay, dropped frames, or stuttering that sometimes characterizes the current state of game streaming.

- By far the most important part of this image is the Chromebook in the foreground powering the display.
- A quick gallery of the various, trippy machine-learning skins being shown in a Stadia tech demo. Here's the standard skin…
- Rainbow powered.
- Cherry blossoms.
- Old school.
- Sketchpad.
- Out-of-focus.
- The set-up.

The important caveat here, of course, is that the demo was running on a wired Ethernet connection hooked to the Moscone Center's industrial-strength Internet hookup. The demo team couldn't confirm the location for the Google data center where the game was actually running, but we can't imagine it would be very far from the heart of San Francisco, where the demo was being played.

In other words, all the demo showed is that Stadia's streaming technology works very well in pretty-much-ideal technical circumstances. How it will work on your own home Wi-Fi

connection—complete with more limited bandwidth, a potentially low-grade router, and general local Internet congestion—was probably better demonstrated by Google's previous Project Stream beta test late last year.

There was only one other playable demo in Google's GDC Stadia display, focused on the machine-learning-based environmental re-skinning technology shown in the announcement this morning. As promised, the tap of a button could change an entire 3D platform game so that every surface broadly resembled the color and patterns of a new source image, shown in the bottom left of the screen.

It's a cool trick and a neat demonstration of the current power of machine-learning algorithms, which staff said can generate a new environmental skin from an image in just a couple of hours. That said, the extremely simple, mostly barren environments weren't really a great demonstration of Stadia's purported 10.7 teraflops of power. As with *Assassin's Creed*, it's fine as a streaming proof of concept but not indicative of the all-important real-world use cases.

Known unknowns

The other Stadia kiosks at GDC used pre-recorded videos to highlight other platform features Google stressed in its keynote announcement. One showed how Stadia's "Stream Connect" can easily send a copy of one player's perspective to another player. This would enable a sort of "command center" gameplay or couch-co-op split screen wherein everyone has their own server instance providing full graphical fidelity. Other videos highlighted YouTube gaming integration efforts like crowd voting and easy, single-instance multiplayer scaling.

But while being able to touch some early version of Stadia's technology at GDC was nice, the demos didn't really answer the many unanswered questions left by Google's scattershot announcement. Chief among them—and glaring by its complete

absence in the presentation—was any mention of pricing. High-end servers and streaming-video bandwidth aren't cheap, even with Google's scale advantages, and we have no idea how Google intends to pay for all that infrastructure. Subscriptions, a la carte purchases, hourly meters, or even ad-based revenue models could all work for a service like this. But Google isn't even hinting at which ones will be on offer here.

The new Stadia controller shown off on stage this morning was also conspicuous by its absence in the demo kiosks. Instead, we used a relatively cheap third-party USB controller to play in the cloud. More than that, everything from Stadia's interface to its platform features to the minimum required bandwidth and router hardware to any indications of a potential game library (aside from *Assassin's Creed* and *Doom Eternal*) are still completely unknown.

That lack of information isn't really out of the ordinary at this point; there are still nine months left for Stadia to launch in 2019, after all. Nintendo didn't reveal any key information about the Switch hardware until just two months before it hit stores, for comparison.

But without important details like these, it's hard to know if Stadia will be an OnLive-style flop, a hardware-agnostic revolution in the way we play games going forward, or something in between. Showing that Stadia can work in ideal conditions is just the first step. We're looking forward to Google taking further steps to prove out the concept going forward.

Listing image by Kyle Orland

https://arstechnica.com/?p=1476705

D3D raytracing no longer exclusive to 2080, as Nvidia brings it to GeForce 10, 16

Enlarge / A screenshot of *Metro Exodus* with raytracing enabled.

reader comments

36 with 30 posters participating

Share this story

Microsoft announced DirectX raytracing a year ago, promising to bring hardware-accelerated raytraced graphics to PC gaming. In August, Nvidia announced its RTX 2080 and 2080Ti, a pair of new video cards with the company's new Turing RTX processors. In addition to the regular graphics-processing hardware, these new chips included two extra sets of additional cores, one set designed for running machine-learning algorithms and the other for computing raytraced graphics. These cards were the first, and currently only, cards to support DirectX Raytracing (DXR).

That's going to change in April, as Nvidia has announced that 10-series and 16-series cards will be getting some amount of raytracing support with next month's driver update. Specifically, we're talking about 10-series cards built with Pascal chips (that's the 1060 6GB or higher), Titan-branded cards with Pascal or Volta chips (the Titan X, XP, and V), and 16-series cards with Turing chips (Turing, in contrast to the Turing RTX, lacks the extra cores for raytracing and machine learning).

Enlarge / The GTX 1060 6GB and above should start supporting DXR with next month’s Nvidia driver update.

Unsurprisingly, the performance of these cards will not match that of the RTX chips. RTX chips use both their raytracing cores and their machine-learning cores for DXR graphics. To achieve a suitable level of performance, the raytracing simulates relatively few light rays and uses machine-learning-based antialiasing to flesh out the raytraced images. Absent the dedicated hardware, DXR on the GTX chips will use 32-bit integer operations on the CUDA cores already used for computation and shader workloads.

Nvidia says that Turing and Pascal cards will take two to three times longer, respectively, to render each frame than a Turing RTX card. This difference is particularly acute on Pascal cards. In Turing, the 32-bit integer workload used for raytracing can run concurrently with the 32-bit floating-point workload used for other graphical tasks. That’s not the case on Pascal, where the workloads will have to be run consecutively.

This weaker performance means that Nvidia recommends that developers use only simpler raytracing effects on the older chips. On the RTX parts, the raytracing performance can be

good enough to enable global illumination—a form of raytracing that enables indirect lighting from reflections, in addition to the usual direct lighting from light sources—but on the GTX parts, Nvidia recommends the use of simpler tasks such as material-specific reflections.

Enlarge / With the RT cores for raytracing and the tensor cores for machine-learning algorithms, the Turing RTX (bottom graph) can process a frame with raytracing relatively quickly. The Turing (middle graph) doesn't have the dedicated cores, but it can still run the integer workload simultaneously with its floating-point workload for a total frame time of about double the RTX. The Pascal (top graph) has to run the integer and floating-point tasks back to back, so it takes much longer than the Turing, let alone the Turing RTX.

That raytracing can be performed on these older chips isn't a big surprise. During DXR's development, Microsoft used a compute-shader-based raytracing algorithm, so clearly the dedicated hardware isn't necessary. However, the substantial performance difference shows that the dedicated hardware is going to be important, at least for now.

Existing games that use DXR should automatically start using raytracing on the 10-series and 16-series cards once the drivers are updated, without requiring any game updates.

However, given the performance differential, we would imagine that developers will want to tailor their raytracing to the older hardware. The burden imposed by raytracing on the RTX chips is substantial, and it’s simply not going to be practical to try to use the same level of raytracing quality on the older cards.

https://arstechnica.com/?p=1476009