Episode Book

EPISODE #5911

What Revit 2027 Actually Costs to Run

PUBLISHED	12 October 2026 at 17:31 UTC
DURATION	27:23
LISTEN	myweirdprompts.com/episode/revit-2027-workstation-cost-06946/
HOSTS	Corn, Herman, Hilbert
SCRIPT MODEL	deepseek-flash
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.4
SOURCES	12 cited by the research (listed at the end)
BOOK GENERATED	19 October 2026 at 17:41 UTC

SYNOPSIS

A two-person architecture practice wants Revit 2027 running persistently so AI agents can query it through Autodesk's official MCP server. The hardware question looks simple until you price it out. We walk through Revit's actual system requirements, why single-core clock speed matters more than core count, what the read-only MCP tech preview can and can't do, and the three ways to arrange this setup — local workstation, modest box plus cloud rendering, or fully cloud-hosted. Along the way: Chaos's blunt warnings about vGPU and Enscape, why V-Ray's CPU rendering quietly changes the math, and where the budget actually lands.

TRANSCRIPT



Corn

Michael's back, and this time he wants a number.



Herman

Part one was the architecture. Part two is the invoice.



Corn

Right. Michael runs a two-person practice. He wants Revit sitting on a persistent Windows machine, humans dialing in over remote desktop, AI agents talking to it through Autodesk's official MCP server. Part one we worked out how that would hang together. Now he's asking what it costs, and specifically whether he needs to buy a big GPU or whether a modest box plus cloud rendering is the smarter play.



Herman

And the interesting thing is the second half of that question is where all the money is. The GPU is the thing everyone fixates on. It's not the thing that decides this.



Corn

So we start with what Revit 2027 actually demands, then which of those demands matter for his use case, then the three ways you can arrange this, then what cloud rendering really costs, and where the budget actually lands.



Herman

And I'll say now, before we get in, one of those three arrangements is going to look obviously right by the end and one of them is going to look like a trap.



Corn

Let's find out which.



Herman

So the framing I want to start with is that Michael's requirement is unusual, and the unusual part isn't rendering. A two-person practice that renders occasionally does not need anything strange. What's strange is persistence. Revit has to stay live for an agent even when neither architect is at the desk. That's the requirement doing the work here.



Corn

Live as in the process, not live as in the machine.



Herman

Exactly that, and it's a distinction worth being precise about, because it changes the hardware decision completely. The official Revit MCP server ships as a separate addon for Revit 2027, announced mid-June, and it's a tech preview. It talks over stdio to any MCP-compliant client. Claude Desktop, Cursor, whatever you're running.



Corn

And it's read-only.



Herman

It's read-only and that's deliberate. Autodesk said so explicitly. The current release focuses on read-only access to your model, quote, that's intentional. Seven tools. Get Running Revit Instances, Query Model, Get Element Data, Select Elements, Zoom to Elements, Open View, Export View. Write capability is deferred to a separate write server that doesn't exist yet.



Corn

Seven tools for a tech preview is honest, at least. They're not pretending.



Herman

The tool that matters for Michael's budget is the first one. Get Running Revit Instances returns the process ID and the open document name for each running instance of Revit. Which means the machine has to be running Revit. Not Windows. Revit. A powered-on VM with Revit closed satisfies nothing.



Corn

That's the whole architecture in one tool description.



Herman

It caps what the agent can do today too, and that matters for budgeting, because you're not budgeting for a machine that edits models. You're budgeting for a machine that reads them, answers questions about them, and exports views. That's a much lighter workload than people imagine when they hear AI agents and building models in the same sentence.



Corn

And the write server when it comes changes that arithmetic.



Herman

It changes it, but it doesn't change it in the direction people assume. A write-capable agent hammering a live model with cascading parameter changes is a different reliability problem, not really a different hardware problem.



Corn

Which is the conversation we had about BIM and SQL, where the analogy holds right up until the moment it doesn't.



Herman

And it doesn't exactly at the point where an agent is doing cascading edits. That's the thing that makes a write server risky. Not the compute.



Corn

So the punchline we're building toward here, and I want to say it early so we can spend the episode testing it, is that the GPU question is probably a red herring and the licensing question is probably the real cost.



Herman

That's the shape of it. Autodesk publishes Revit 2027 requirements in three tiers and the CPU guidance is identical across all three, which tells you something. Two gigahertz or better base clock, and Autodesk explicitly recommends the highest single-core base clock you can get. Revit will use multiple cores for many tasks, they say, but single-thread speed dominates responsiveness.



Corn

That's the sentence. That's the one Michael should tape to the wall.



Herman

The three tiers are really about memory and model size. Minimum tier is sixteen gigabytes of RAM for a single model up to around three hundred megabytes on disk. DirectX 11 GPU with Shader Model 5, four gigabytes of VRAM recommended, twelve eighty by ten twenty-four display, thirty gigabytes of install plus a hundred gigabytes free for temp files.



Corn

Temp files.



Herman

Revit is not shy about temp files. The balanced tier is thirty-two gigabytes of RAM for models around six hundred megabytes, six gigabytes of VRAM, and eight plus if you want accelerated graphics, sixteen eighty by ten fifty. The performance tier is sixty-four gigabytes for models around a gigabyte, eight gigabytes of VRAM, nineteen twenty by twelve hundred.



Corn

And the CPU line doesn't move across any of them.



Herman

Doesn't move at all. Which is Autodesk telling you plainly that this application is not going to reward you for buying more cores. It's going to reward you for buying one fast core.



Corn

Two constraints that break plans, though.



Herman

Revit 2027 is Windows 11 only. Windows 10 was dropped. And it requires .NET 10. So if you were planning to spin up a VM image you've had sitting around, check it, because that's a hard wall, not a preference.



Corn

Now the virtualization layer, because that's where this gets slippery.



Herman

Autodesk's own position is deliberately non-committal and I want to read it close because the wording is doing work. You may virtualize a product only if the applicable terms and conditions expressly permit virtualization. And they offer virtualization information as is, with no warranties for virtualized use.



Corn

That's not a yes.



Herman

That's not a yes and it's not a no. It's a we'll tell you how, and if it goes wrong that's yours.



Corn

Flag it as a risk to verify, not a settled fact.



Herman

Which we'll come back to when we do the budget, because it's the single biggest open number in this whole thing.



Corn

The virtualization guidance that's actually concrete comes from Chaos, on the rendering side.



Herman

And it's blunter than Autodesk by a wide margin. Dedicated GPU passthrough is what they call ideal for Enscape. Partitioned GPU, vGPU, is shared and less reliable for real-time rendering, and they say it often causes Enscape to fail to launch or perform poorly.



Corn

Often causes it to fail to launch. That's not a caveat, that's a warning.



Herman

Their recommendation for shared environments is sixteen or more vCPUs and thirty-two gigabytes of RAM guaranteed per user, not shared. Resource sharing between sessions causes instability and degraded performance. That's their line.



Corn

So the virtualization answer for a two-person practice, if you take Chaos at their word, is one VM per user with its own guaranteed resources and a passthrough GPU. Which is not really sharing anything.



Herman

It's not sharing. It's two machines in a trench coat.



Corn

The remote protocol matters more than people expect.



Herman

It does, and Chaos tested them, which I appreciate. Microsoft RDP and Parsec come out best. VMware Blast they rated laggy. ScreenConnect they rated unusable.



Corn

Unusable is a strong word for a product review.



Herman

They tested it with Enscape running, which is the difference. A protocol that's fine for spreadsheets can be unusable for a real-time 3D viewport. And the network floor is wired Ethernet, fifty megabits up and down minimum, forty milliseconds latency or better.



Corn

Wired.



Herman

The latency on a viewport is brutal. Their test matrix is worth saying out loud because it shows the failure case. Hyper-V with an RTX A4000 passthrough, very good. Parsec to an Azure VM with an NVIDIA T4, very good. Azure VM with an AMD MI25, failed outright, because GPU passthrough wasn't configured.



Corn

Failed because somebody forgot a checkbox.



Herman

That's the whole point of the example. The GPU was there. It just wasn't attached to the session.



Corn

So now the split that actually decides Michael's budget. Modeling versus rendering.



Herman

Revit modeling is CPU bound and single-thread sensitive. The GPU handles display, view navigation, shadows, textures, large three-dimensional scenes. Revit's own native rendering is modest. It's not a path tracer and it was never meant to be.



Corn

Enscape is the one that changes the answer.



Herman

Enscape is GPU mandatory. Almost all the calculations are handled on the graphics card, in their own words, and it needs dedicated VRAM, not shared. Intel integrated graphics don't qualify at all. But here's the thing people get wrong about Enscape. It's a rasterized, lightly ray-traced real-time preview. It is not a full path tracer. It's built for speed in a viewport, not for final frame quality.



Corn

And V-Ray is the sleeper fact of this episode.



Herman

V-Ray can render on CPU alone, or CPU plus GPU hybrid. A studio can render on CPU alone, on hardware it may already own. That's their phrasing and it's the most consequential sentence we'll read today.



Corn

So Michael does not have to buy a GPU to produce final quality renders.



Herman

He does not. He can render on the CPU of a machine he already has, or he intends to buy for other reasons, and burst to the cloud when he needs more.



Corn

Which answers the sub-question directly.



Herman

If the workstation's primary purpose is keeping Revit available for human interaction and MCP-driven automation, then yes, it can operate satisfactorily with relatively modest graphics hardware. The official MCP server is read-only anyway. It is not asking the GPU to do anything. The GPU only becomes mandatory if Enscape is the chosen visualization tool, and only then because Enscape refuses to work any other way.



Corn

So that's what Revit asks for. The interesting question is which of those resources matter for Michael's use case, and that depends on what the machine is actually doing.



Herman

Right, and once you know which tasks are CPU bound and which are GPU bound, the deployment question almost answers itself. But there are three ways to arrange it and they have very different cost profiles.



Corn

Model A. Powerful local workstation, hosting the Revit VM, enough CPU and GPU for both modeling and rendering.



Herman

Advantages are real. Everything local, no data transfer, no cloud dependency, works offline. The limitation is capital cost, and in a two-person practice the GPU sits idle most of the time. That's the honest objection. You're buying a seven hundred dollar card to use it forty hours a month at most.



Corn

Complexity is moderate. You're running a hypervisor on a workstation.



Herman

It's good if the practice renders frequently and wants zero cloud friction. If someone in that office is doing visualizations every day, this is the right build and you should stop reading spreadsheets and go buy it.



Corn

Model B. Modest local workstation, Revit and the AI integrations on it, demanding rendering delegated to cloud services or remote GPU nodes.



Herman

Lowest capital cost, GPU paid for only when used. The limitation is the one we'll spend the most time on, which is that cloud rendering is not renderer-agnostic. Complexity is moderate to high because you're managing asset sync and licensing across two environments.



Corn

And suitability. Likely the best fit for a two-person practice.



Herman

Likely. With one condition attached that we'll get to.



Corn

Model C. Fully cloud-hosted Windows GPU workstation running Revit, accessed remotely by both architects and agents.



Herman

No local hardware, scales up and down. And Chaos, who sell cloud services, warn about it. Full VM deployments are often seen as complex and cost-intensive, due to high GPU instance pricing in cloud environments. Evaluate the total cost of ownership carefully. That's a vendor telling you to do the math before you buy their product.



Corn

I appreciate that more than I probably should.



Herman

It's an unusually honest line. Complexity is high. Suitability for a two-person practice is questionable unless the always-on requirement gets relaxed, and the always-on requirement is the whole point of the project.



Corn

Which is the trap in Model C, and I want to sit on it for a second.



Herman

Go ahead.



Corn

A cloud workstation can run continuously, which satisfies the persistent AI endpoint. And it bills continuously. Or it can be stoppable, which is cheaper, and then the agent can't reach Revit when the machine is off, which breaks the requirement. There's no third setting.



Herman

That's exactly the tension. You're paying for twenty-four hour availability to serve an agent that might use it for twenty minutes at two in the morning.



Corn

And you can't schedule around it, because the agent's whole value is that it's there when nobody's looking.



Herman

Which is why Model B keeps winning. The modest local box is always on anyway. It costs you electricity and a machine you'd have bought regardless. The persistence is nearly free.



Corn

Now the place where most plans fall apart. The cloud rendering support question.



Herman

Chaos Cloud is native to V-Ray only. One click from the V-Ray toolbar. Smart Sync uploads only changed data, Smart Vault stores assets so you upload once, Live View for watching progress. It is not a generic renderer-agnostic service.



Corn

Enscape has no Chaos Cloud render-offload path.



Herman

Enscape's cloud features are for sharing and review, not rendering. Its rendering is local GPU bound by design, which follows from the fact that it's a real-time viewport tool. You can't render what you can't calculate locally.



Corn

And the third-party farms.



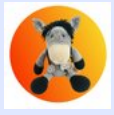
Herman

GarageFarm, RebusFarm, Ranch Computing, Fox Render Farm. They support V-Ray, Corona, Blender, 3ds Max, SketchUp. They do not support Enscape.



Corn

So this is the decisive constraint. If Michael wants cloud rendering, he has to choose V-Ray.



Herman

And if he chooses Enscape, he has to buy the GPU, because Enscape will never offload and never will, because offloading isn't what it's for.



Corn

So the "modest host plus cloud rendering" plan has a fork in it that people don't see when they're sketching it on a whiteboard. They write "render in the cloud" and they think it's renderer agnostic. It isn't. It's V-Ray in the cloud or it's nothing in the cloud.



Herman

Which is fine, because V-Ray is the better fit for this use case anyway. CPU rendering locally, cloud bursts when you need them, no GPU purchase.



Corn

Walk through what actually gets uploaded, because that's the hidden cost people miss.



Herman

Scene file, fifty to five hundred megabytes. Textures and materials, five hundred megabytes to five gigabytes. Proxies, two hundred megabytes to three gigabytes. HDRI, fifty to five hundred megabytes. Total on first upload, one to eight gigabytes.



Corn

Which at fifty megabits up is three to twenty-five minutes.



Herman

On the first upload. Subsequent renders re-sync only the changes through Smart Sync, so it drops to almost nothing if you're iterating on the same scene.



Corn

That's the design that makes cloud rendering actually usable day to day. The first upload is painful and every one after it is fine.



Herman

Now the pricing itself. Chaos Cloud sells credit packs. Thirty thousand credits, six hundred twenty-five dollars a year, which works out to about two cents a credit. Sixty thousand credits, one thousand fifty dollars a year, about one point seven five cents a credit.



Corn

Typical cost per image.



Herman

Five to thirty dollars. A thirty-second animation, two hundred to six hundred.



Corn

So a rendering-heavy month could run you a few hundred dollars without anyone noticing.

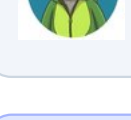


Herman

And the hidden costs are real. Storage for uploaded assets, data transfer, minimum billing increments, renderer licensing on the cloud side, and setup time. And the pricing itself is volatile. Chaos forum users have been vocal that the new credit system is a disproportionate price hike. One user reported 4K renders at medium plus quality consuming two to five credits per image under the old model, and called the new pricing completely unviable for daily professional work.

Corn

Which is a user complaint, not a price list. Take it as a signal that this is a moving target.



Herman

Take it as a signal that whatever number you budget for cloud rendering, your budget has a half-life.



Corn

Break-even.



Herman

Archgyan put a line on it. If you render more than twenty images per month, upgrading your hardware is almost certainly cheaper than cloud rendering over a year. Below that, cloud wins, especially for animations and occasional large batches. Local wins for single images with flexible deadlines and draft renders, where you can just leave it cooking overnight.



Corn

Which is a clean rule and I'd want to stress the assumptions. Twenty images a month, what quality, what resolution.



Herman

It's a ballpark. The shape is right even if the exact number isn't. If you're rendering every day, buy the card. If you're rendering twice a month, rent the cloud. If you're somewhere in the middle, do the arithmetic with your own render times.



Corn

So let's put real numbers on it. Not false precision, but realistic figures with transparent assumptions.



Herman

Revit subscription, US list. Three hundred eighty dollars a month, three thousand five dollars a year, or nine thousand twenty for a three-year term.



Corn

And the three-year term is five dollars more than three annual terms.



Herman

Five dollars more. Which is a price freeze, not a discount. You're paying five dollars to lock the rate for three years.



Corn

I'd take that deal in a heartbeat, actually.



Herman

Flex tokens are the other option. Ten tokens a day is about thirty dollars a day, and that breaks even against the annual at around a hundred days a year.



Corn

So it's for people who need Revit occasionally, not for a practice using it to earn a living.



Herman

Revit LT is five hundred sixty dollars a year, about eighty-one percent cheaper. And it's a false economy for Michael, for two reasons. It drops worksharing. And Enscape is not available as a plugin for Revit LT at all.



Corn

So if visualization matters, LT is out before you finish reading the spec sheet.



Herman

The AEC Collection is the one people skip past and shouldn't. Three thousand six hundred seventy-five dollars a year. Only six hundred seventy more than Revit alone. And it includes 3ds Max, AutoCAD, Civil 3D, Navisworks.



Corn

And 3ds Max alone is two thousand ten dollars a year.



Herman

So the six hundred seventy dollar uplift buys you a two thousand dollar application. If you need rendering at all, that's the real price of adding it to Revit.



Corn

Renderers, because there's a menu here.



Herman

Enscape Solo, five hundred seventy-four eighty a year. Premium, six hundred seventy eighty. Collection, seven hundred thirty eighty. V-Ray Solo, five hundred forty. Premium, seven hundred seventy-eight eighty. Collection, one thousand two hundred ten eighty. D5 Render Pro, three hundred sixty. Lumion View, two hundred twenty-nine. Twinmotion is included with Revit.



Corn

Twinmotion being included is worth flagging, because it changes the floor on this.



Herman

It does. If you're price sensitive and you're not chasing a specific look, you can spend zero on a renderer and get something usable out of the box.



Corn

Hardware reference points.



Herman

RTX 4070 Ti, around seven hundred dollars. Ryzen 9 7950X, around five hundred fifty. A rendering-capable workstation, two thousand to four thousand.



Corn

And the current leader as of May was an AMD Ryzen 9 9950X paired with an NVIDIA RTX 5070, which is a useful reference for what a single-user performance build looks like.



Herman

It's a good spec, and it's more than Michael needs for the read-only MCP workload.



Corn

Now the licensing question that has no answer in anything we found.



Herman

Whether a second Revit seat is required for a VM used by AI agents, versus the named user's existing seat. Autodesk's published terms do not resolve it.



Corn

And if it's required, that's three thousand five dollars a year added to the bill.



Herman

It could be. Or it could be nothing, if the named user's seat covers the virtual instance. Nobody outside Autodesk can tell you which.



Corn

Which makes it the single biggest unknown in the budget.



Herman

Verify with Autodesk directly before you commit to any of this. That's not a small caveat. That's the difference between a four thousand dollar a year system and a seven thousand dollar a year system.



Corn

Now two usage patterns, because the answer differs.



Herman

Occasional architectural rendering. A couple of images a month, maybe one animation a year. Modest local host, V-Ray CPU rendering, an occasional Chaos Cloud burst. No discrete GPU purchase at all. You're into this for a host and software, and the rendering bill is tens of dollars a month, not hundreds.



Corn

Regular demanding visualization.



Herman

Then the local GPU starts to pay for itself. Above twenty images a month you're into buy territory. Or you keep the cloud budget and accept that it's a subscription to someone else's GPU forever.



Corn

And the cloud-hosted workstation distinction, because I think this is where the plan actually breaks.



Herman

Continuous versus stoppable. A machine that runs continuously satisfies the persistent AI endpoint requirement and bills around the clock. A stoppable machine is cheaper and breaks persistence, because the agent can't reach Revit if the VM is off.



Corn

The always-available requirement and the cost-minimizing cloud configuration are in direct conflict.



Herman

They are. Which is the fundamental tension in Model C, and it's why I'd steer Michael away from it unless his practice is already bigger than two people.



Corn

The honest answer to "does Michael need a powerful GPU" is no. And the plain answer to "what's the most cost-effective infrastructure" is a modest local box, thirty-two gigabytes of RAM, a high single-clock CPU, modest graphics, running Revit 2027 and the official MCP server, with V-Ray doing CPU rendering and cloud bursts when he needs them.



Herman

With one asterisk. The second Revit seat is unverified.



Corn

Which is a weird place to land, honestly. The cheapest part of this whole plan is the part everyone worries about.



Hilbert

It's two oh two watts.



Hilbert

The three machines. Measured at the wall, all three running, two hundred and two watts. I had one of those plug-in meters. I wrote it on a sticky note.



Corn

What was the render farm for, Hilbert?



Hilbert

A garden shed. Six weeks. I'm not going to get into the details of the shed.



Herman

He's telling us he built a render farm for a shed.



Hilbert

It wasn't a farm by the end, it was three machines in the spare bedroom and the fan noise was so bad the neighbour asked if I was running a server. I told him I was. He said he'd guessed. That's when I knew I'd gone too far.



Corn

Six weeks for a shed.



Hilbert

It was a nice shed. It's not built. I still have it.



Herman

The shed.



Hilbert

The shed and the machines. They're unplugged now. The wife calls the spare room the museum.



Corn

Why the fridge, Hilbert?



Hilbert

The sticky note. It was the only flat surface left.



Herman

The wattage number ended up on the fridge.



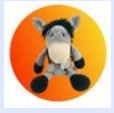
Hilbert

It did. She threw it away. I remember the number, though. Two hundred and two.



Corn

Let's land this. Where does the money actually go for a two-person practice that wants a persistent AI-accessible Revit?



Herman

The recommended configuration is a modest local host. Thirty-two gigabytes of RAM, a CPU with the highest single-core clock you can afford, modest graphics, Revit 2027 and the official MCP server running on it, V-Ray for CPU rendering, and occasional Chaos



Herman (continued)

Cloud bursts. No big GPU purchase unless Enscape becomes mandatory.



Corn

Initial budget is a host in the two thousand to four thousand dollar range if you're buying new, less if you're upgrading an existing desktop. Recurring is Revit at three thousand five a year, plus a renderer if you're not running Twinmotion, plus cloud rendering in the tens to low hundreds a month depending on how much you render.



Herman

The unresolved condition sits on top of all of it. If Autodesk requires a second seat for the AI-accessible VM, add three thousand five a year. Verify before committing.



Corn

Two ways this changes. If the practice renders more than about twenty images a month, or if Enscape becomes the primary visualization tool, a local GPU becomes worthwhile. And if the practice grows past two people, or if the always-available requirement can be relaxed to business hours only, a fully cloud-hosted workstation starts to make sense.



Herman

Past two people, the cloud VM argument gets much stronger. Because at that point you're paying for concurrency, and that's what cloud is actually good at.



Corn

The one thing I'd take from this whole episode, and it's not what Michael asked about, is that persistence is cheap and it's licensing that's expensive.



Herman

Which is the opposite of what everyone assumes going in. Everyone budgets for the hardware and gets surprised by the contract.



Corn

A man in the room said two hundred and two watts and walked away.



Herman

We'll be back soon with more. Our thanks to producer Hilbert Flumingtop, still at the desk, still unimpressed.



Corn

Send us your own prompt on Telegram at t dot me slash MWP listener bot. Everything else lives at my weird prompts dot com.



Herman

For more along these lines, there's episode forty-seven, Compressing Days into Minutes; episode twenty-five forty-four, How to Make AI Architectural Renders Photoreal Without Breaking Geometry; and episode twenty-four fifty-two, When BIM Breaks the SQL Analogy. This has been My Weird Prompts.



Corn

The human-AI collaboration podcast. See you tomorrow.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **Autodesk [primary]**
   Revit Public MCP Server announcement, 2026-06-17
   https://www.autodesk.com/blogs/aec/2026/06/17/revit-public-mcp-server/

2. **Autodesk [primary]**
   System requirements for Revit 2026 products (updated 2026-08-06)
   https://www.autodesk.com/support/technical/article/caas/sfdcarticles/sfdcarticles/System-requirements-for-Revit-2026-products.html

3. **Autodesk [primary]**
   Buy Revit (pricing/FAQ)
   https://www.autodesk.com/products/revit/buy

4. **Chaos [primary]**
Running Enscape in Virtual and Remote Environments
https://documentation.chaos.com/space/ENSCAPE/842006572/Running%20Enscape%20in%20Virtual%20and%20Remote%20Environments

5. **Chaos [primary]**
Cloud Rendering / Chaos Cloud pricing
https://www.chaos.com/cloud-rendering

6. **goto.archi**
Revit 2027 System Requirements, 2026-05-27
https://goto.archi/blog/post/revit-2027-system-requirements

7. **Archgyan**
Cloud Rendering for Architects: Cost, Services & Setup Guide (2026), 2026-04-04
https://archgyan.com/render-farm-guide-cloud-rendering-architectural-visuals/

8. **Render a House**
Revit Pricing 2026, updated 2026-09-05
https://www.renderahouse.com/blog/revit-pricing

9. **Render a House**
V-Ray Pricing 2026, updated 2026-09-07
https://www.renderahouse.com/blog/vray-pricing

10. **Render a House**
Enscape Pricing 2026, 2026-09-07
https://www.renderahouse.com/blog/enscape-pricing

11. **GitHub**
revit-mcp/revit-mcp (community MCP server, archived in favour of monorepo)
https://github.com/revit-mcp/revit-mcp

12. **Chaos Forum**
The new credit system is unsustainable
https://forums.chaos.com/t/the-new-credit-system-is-unsustainable-a-disproportionate-price-hike-that-drives-users-away/191734