Zaha Hadid Architects generative AI design workflow

IMPLEMENTED 2022–2025 · CONSTRUCTION AND PROPERTY · UK · BAND C · CLAIMED ONLY
VERSION 1.0 · PUBLISHED 15 NOVEMBER 2026 · UPDATED 15 NOVEMBER 2026 · FACTS CHECKED 15 NOVEMBER 2026

Zaha Hadid Architects has used generative AI in design since early 2022 and says it is now a standard part of project initiation. By the practice's account, renderings take 20% of the time they did and mid-stage preparation is up to 50% more productive. No figure is independently checked.

ABOUT THIS SYNOPSIS

This is MarLuc's summary and grading of publicly available sources, for information only. It is not investment, legal, financial or professional advice, and should not be relied on for any decision. Readers must validate it against the original sources and, where appropriate, challenge its conclusions.

Facts and figures are as reported by the sources listed below, which may be incomplete, self-reported or out of date. Check the original source before relying on any figure.

The impact band, proof marker and SEAM view are MarLuc's opinion. The organisation is not a MarLuc client and has not reviewed or endorsed this synopsis. Names and trademarks belong to their owners.

NOTICE WORDING LAST UPDATED 16 NOVEMBER 2026

AT A GLANCE

AI used: Generative AI image models (Stable Diffusion, DALL-E 2, Midjourney, Gendo, Rhino.ai) and custom fine-tuned models, used for concept images and visualisation.

Implemented: 2022–2025. Grade: band C, Process change, Claimed only.

Headline result, as reported: By the practice's account: renderings in 20% of the time and up to 50% productivity gain in mid-stage design preparation

1 · The case

The facts, as reported by the 4 sources listed at the end. Where a source is silent, the field says "Not reported".

Organisation

Zaha Hadid Architects. London-based architecture practice. Patrik Schumacher is described by Building Design as chief executive. 60 projects in development and projected revenue of £77 million for the financial year (January 2025). Canvas block: 4. Operations and process. Also 3. Customers and proposition.

Business challenge

The challenge. By NVIDIA's account, producing rendered visuals for clients was slow, with weeklong cycles and scene datasets of up to 60 million polygons, and the practice's 40 years of 3D models and imagery had to stay confidential. Schumacher told The Times, as relayed by

Who implemented it

NVIDIA (Software vendor): Supplied the GPUs, Omniverse and NIM used in the practice's AI pipeline, by its own customer story..
HP Inc. (Equipment or robot maker): Named by NVIDIA as a partner; HP technologies are credited with faster computation..
In-house (In-house): Built the internal tools and custom models..
The model makers behind Stable Diffusion, DALL-E 2 and Midjourney are used as tools and are not named as partners in the sources opened.

Technology (SEAM pillar E)

SEAM types: T5 Custom / fine-tuned models. Deployment: Own infrastructure. Vendors and products: NVIDIA, HP Inc., Stable Diffusion models and LoRAs fine-tuned on the practice's data, DALL-E 2 (earlier), Midjourney and Gendo (earlier), Rhino.ai, NVIDIA Omniverse, NVIDIA NIM, ZSPACE Create (internal).

Where it was built in: Concept and visualisation work in the design pipeline, linked to the practice's image database and to its Computation & Design group's tools through Omniverse and OpenUSD, by NVIDIA's account. In the earlier talk the models ran in the cloud, with a plan to move to an in-house data centre.

What changed

Before: rendering was slower, with weeklong client feedback cycles. After, by the practice's account: AI systems are a standard part of project initiation, renderings take 20% of the time, and feedback that took a week takes under 24 hours. [1] [2] [3] [4]

What it cost

Not reported.

What it saved

By the practice's account, renderings in 20% of the previous time and productivity gains of up to 50% in mid-stage design preparation (Building Design). Feedback loops from one week to under 24 hours (NVIDIA). Money or hours saved: Not reported. [1] [2] [3] [4]

Governance and control

By NVIDIA's account, AI workflows run on local compute to meet IP protection and client confidentiality needs. Schumacher urges firms to disclose AI use. In the earlier talk, 3D artists built the models by traditional methods after concepts were chosen. A written AI policy, accountable owner and monitoring: Not reported.

S2 · Accountable owner · Claimed
Schumacher, described as chief executive, leads public statements on AI use, and the Computation & Design group builds the tools; a named owner for AI risk is not reported. [1] [2]

A1 · Safe use · Claimed
By NVIDIA's account, AI workflows stay on local compute to meet IP protection and client confidentiality needs. [2]

A2 · Human in the loop · Claimed
In the earlier talk, the practice used 10 to 15 per cent of model output and 3D artists built the models by traditional methods. [3]

A3 · Monitoring and assurance · Not reported
Not reported.

A4 · Vendor AI · Claimed
By Building Design, the practice moved from off-the-shelf tools to proprietary software, and by NVIDIA's account it fine-tunes models on its own data. [1] [2]

A5 · Agent controls · Not reported
No agents described.

A6 · Standards and regulation · Not reported
Not reported.

What worked, what did not, what is unproven

WORKED

– The practice says AI is a standard part of project initiation and generates a wider range of options in less time (via Building Design).

– Schumacher says AI is a very strong boost in winning work (Building Design). No win-rate figure is given.

– The practice built its own toolchain so that confidential data stays on local compute (NVIDIA's account).

DID NOT

– The efficiency figures are the practice's own. Building Design says the firm claims them, and no baseline or method is given.

– In the earlier talk the practice used only 10 to 15 per cent of model output, and 3D artists then built models by traditional methods, so image generation did not replace modelling.

– Sector context, not this practice: in RIBA's 2025 member survey, 67% worried that AI will increase the risk of work being imitated and 35% were concerned that it will threaten the profession.

UNPROVEN

– Whether the 20% and 50% figures hold across projects, and how they were measured: Not reported.

– Effect on fees, win rate or headcount: Not reported.

– The date and scope of the NVIDIA customer story's figures: Not reported on the page.

Timeline

2022 Practice begins integrating AI systems into its workflows, by its account [1]

Jan 2025 Building Design reports the 20% rendering time and up to 50% productivity claims [1]

Jun 2025 RIBA reports that 59% of architect practices use AI, up from 41% in 2024 (sector figure) [4]

2 · Grade and SEAM assessment

MarLuc's opinion, made on the evidence above. How cases are graded: marluc.co.uk/case-studies/how-we-rate

Level of change and impact

Claimed: Not applicable. Grade: C, Process change. Proof: Claimed only. Level of change: Process change.

Process change (Level 2, band C), by the practice's account. AI sits in project initiation and the visualisation pipeline. The sources show no change to structure, roles or business model.

The practice describes AI in project initiation and a proprietary toolchain, with faster client feedback and less outsourcing, which is a change to how early design and visualisation are done. All results are the practice's own or its supplier's, with no baseline or method, so the marker is Claimed only. The sources show no change to structure, roles or fees.

The SEAM view and lessons learnt

The practice treats AI as part of the design process and has invested in its own tools to protect client data. Every result in the sources is the practice's own or its supplier's. The earlier talk shows people still choose concepts and rebuild them as models. For a mid-sized or larger organisation, the transferable point is to record a baseline before claiming a time saving and to keep selection and modelling with named people.

1. Record a baseline before claiming a time saving.
2. Keep client data on infrastructure you control.
3. Disclose AI use to clients and keep selection with named designers.

Evidence

Score 1 of 4. Own account: the practice's chief executive in trade press, a supplier's customer story and an earlier talk reported by DeepLearning.AI. No independent measurement.

Evaluation method. Company statements relayed by trade press, a vendor customer story and a conference talk report; no baseline or audit.

Status. Ongoing. Latest source for the practice's figures: Building Design, 6 January 2025. Later results: Not reported.

3 · Original sources

Read these in full. To challenge a fact or a grade, use the contact form at marluc.co.uk or the link on https://marluc.co.uk/case-studies/zaha-hadid-architects-generative-ai-design

[1] Building Design: Patrik Schumacher claims AI keeps ZHA 'a step ahead' (6 January 2025)
https://www.bdonline.co.uk/news/patrik-schumacher-claims-ai-keeps-zha-a-step-ahead/5133607.article

[2] NVIDIA: Zaha Hadid Architects Accelerate Design Innovation with NVIDIA-Accelerated AI Pipelines (customer story)
https://www.nvidia.com/en-us/customer-stories/zaha-hadid-architects-with-omniverse-and-usd/

[3] DeepLearning.AI, The Batch: How a Top Architecture Firm is Using Generative AI
https://www.deeplearning.ai/the-batch/how-a-top-architecture-firm-is-using-generative-ai

[4] RIBA: Second RIBA AI Report shows surge in usage among UK architects (27 June 2025)
https://www.riba.org/knowledge-and-resources/knowledge-landing-page/second-riba-ai-report-shows-surge-in-usage-among-uk-architects

Zaha Hadid Architects website: https://www.zaha-hadid.com

Version 1.0 · Published 15 November 2026 · Updated 15 November 2026 · Facts checked against sources 15 November 2026 · Downloaded 27 November 2026. This PDF may be out of date: the latest version, and a way to report an error or a better source, is at https://marluc.co.uk/case-studies/zaha-hadid-architects-generative-ai-design

VERSION HISTORY

v1.0 15 November 2026 First published