

MIDLAND/ODESSA DIVISION

)

)

)

)

)

)

)

)

)

MEMORANDUM IN RESPONSE TO PETITIONER’S MOTION TO COMPEL COMPLIANCE WITH SUBPOENAS SERVED ON NON-PARTY GOOGLE LLC

TABLE OF CONTENTS

	Page
I. INTRODUCTION	1
II. FACTUAL BACKGROUND.....	3
III. ARGUMENT	4
A. The Motion Does Not Comply with Local Rule CV-7(g)	4
B. The Subpoenas Do Not Comply with Federal Rule 45.....	6
IV. CONCLUSION.....	10

I. INTRODUCTION

Neural AI, LLC (“Neural”) has filed at least **nine** separate actions to compel discovery.¹ The instant motion thus does not exist on an island: the subpoena to produce documents and to testify (“Subpoenas”) to which Google LLC (“Google”) objects is part of a larger litigation pattern by Neural seeking discovery that is overbroad, unduly burdensome, and not proportional to the needs of its litigation against NVIDIA. In each of these separate actions, Neural served nearly identical subpoenas on non-parties, and when each non-party fairly objected, Neural filed nearly identical motions to compel. It is unclear whether Neural has taken a moment to reflect and consider why each of the non-parties “objected to every request and deposition topic,” *see, e.g., Neural AI v. xAI*, 7:26-mc-00319-DC, Dkt. 1, at 1 (W.D. Tex. Aug. 17, 2026), but one thing remains constant across these actions: Neural’s unreasonable approach to non-party discovery.

In this action, Google timely objected to Neural’s boundless requests seeking information about roughly **240 devices** over an **8-year timeframe**. The parties then engaged in a meet-and-confer process, during which Neural never addressed Google’s objections nor narrowed a single document request or deposition topic. Instead, Neural focused on having Google sign a declaration “in lieu of further document production and/or deposition testimony.” Mot. at 4. The meet-and-confer process thus never addressed Google’s objections to the Subpoenas; instead, Neural filed a precipitate motion to compel just hours after providing Google a “revised declaration” with language Google had never seen. It follows that Neural’s rush to the

¹ See *Neural AI v. Meta Platforms*, No. 7:26-mc-00327-LS (W.D. Tex. Aug. 19, 2026); *Neural AI v. CoreWeave*, No. 7:26-mc-00323-LS (W.D. Tex. Aug. 18, 2026); *Neural AI v. Oracle*, No. 7:26-mc-00325-DC (W.D. Tex. Aug. 18, 2026); *Neural AI v. xAI*, No. 7:26-mc-00319-DC (W.D. Tex. Aug. 17, 2026); *Neural AI v. Tesla*, No. 7:26-mc-00318-LS (W.D. Tex. Aug. 17, 2026); *Neural AI v. OAI Int’l., Inc.*, 3:26-mc-80261-AGT (N.D. Cal., Aug. 17, 2026); *Neural AI v. Amazon.com*, No. 7:26-mc-00241-LS (W.D. Tex. June 24, 2026); *Neural AI v. Microsoft*, No. 7:26-mc-00242-LS (W.D. Tex. June 24, 2026).

courthouse was not driven by an impasse between the parties, but rather by the deadline in the underlying litigation against NVIDIA, which Neural openly admitted. Dkt. 1-21 at 2 (Aug. 18 email from T. Laiche) (“Given today’s deadline, we intend to file a motion solely to preserve our rights and avoid any potential waiver of discovery from Google.”). But filing a discovery motion to “preserve rights” does not comply with Local Rule CV-7(g), which requires that “counsel for the parties have conferred in a good-faith attempt to resolve the matter by agreement.” And, as Google explained in correspondence before Neural filed its motion, the parties did not confer on the “revised declaration” nor “any of the specific requests or topics” in the Subpoenas. Ex. A at 4.² The Court should deny the instant motion because Neural failed to comply with Local Rule CV-7(g).

The Court should also deny the motion because the Subpoenas are overly broad, unduly burdensome, and disproportionate to needs of its litigation against NVIDIA, especially in view of Neural making requests of a non-party. *See* Fed. R. Civ. P. 45. For starters, Neural has not demonstrated that the requested discovery is relevant. Neural simply *guesses* that Google possesses relevant information merely because Google buys NVIDIA GPUs. But this type of “fishing expedition” is inappropriate, especially when directed to non-parties. *Micro Motion, Inc. v. Kane Steel Co., Inc.*, 894 F.2d 1318, 1327-28 (Fed. Cir. 1990). And, while Neural says the Subpoenas seek “targeted” documents, Mot. 1, 5, the opposite is true. Indeed, the very first Definition in the Subpoenas—defining the term “NVIDIA GPUs”—demonstrates the unreasonableness of the various requests Neural makes. According to Neural, the term “NVIDIA GPUs” means nine different “architectures” of NVIDIA graphics processing units, which include over **240 devices**. Dkt. 1-6 at 6–8. Even more, Neural included this definition in

² Neural omitted this correspondence from its motion.

11 of 12 document requests, which would necessarily require Google to produce “voluminous” records. *See, e.g., Lopez v. State Farm Lloyds*, 348 F.R.D. 419, 429 (W.D. Tex. 2025). Google respectfully requests that the Court deny Neural’s motion.

II. FACTUAL BACKGROUND

Neural served Google with the Subpoenas on June 25, 2026, and Google timely served written objections on July 21, 2026. Contrary to Neural’s presentation of the facts, Google’s objections were not “boilerplate,” Mot. at 3, but instead explained in detail why each request or topic was improper. Dkt. 1-10 at 10–31, Dkt. 1-11 at 9–14. Google also offered to meet-and-confer to discuss each document request and deposition topic. *Id.*

The parties conferred on July 24, 2026. Rather than address Google’s objections to the Subpoenas, Neural sent over a list of questions and a template declaration for Google’s review, saying Neural would “consider accepting the declaration in lieu of further document production and/or deposition testimony, subject to resolving any material gaps.” Dkt. 1-21 at 7–8 (Aug. 6 email from R. Magni). Google responded that it might be able to “provide a declaration within a more reasonable scope.” *Id.* at 7 (Aug. 6 email from J. Jaffe). Neural agreed, and asked for “proposed edits to the declaration as soon as possible[.]” *Id.* at 4 (Aug. 11 email from T. Laiche). Google responded two days later that it had an update and requested a call with Neural, offering times that day and the following day. *Id.* at 3–4 (Aug. 13 email from J. Jaffe). Neural did not respond until that weekend. *Id.* at 3 (Aug. 16 email from T. Laiche). Google responded the next day, and the parties conferred that same day. *Id.* at 2 (Aug. 17 email from J. Jaffe).

After conferring on August 17, Neural sent Google two follow-up emails on August 18. In the first email, Neural wrote that it would file its motion to compel that day “to preserve our rights and avoid any potential waiver of discovery from Google.” *Id.* at 1–2 (Aug. 18 email from T. Laiche). In the second email, sent about five hours later, Neural attached a “revised

declaration,” which Neural said it believed “accurately reflects Google’s operations.” *Id.* at 1 (Aug. 18 email from T. Laiche). As Google pointed out in an email sent *before* Neural filed its motion, the parties never conferred regarding the “revised declaration” nor “any of the specific requests or topics” in the Subpoenas. Ex. A at 3–4. (Aug. 18 email from J. Jaffe) (“Needless to say, Google and Neural AI have not met and conferred regarding Neural AI’s latest positions and revised declaration.”). Because of this, Google explained that Neural’s decision to file a discovery motion was premature and in violation of this Court’s local rules. *Id.*

III. ARGUMENT

Neural’s motion to compel should be denied because the motion does not comply with this Court’s local rules and because the Subpoenas do not comply with Federal Rule 45.

A. The Motion Does Not Comply with Local Rule CV-7(g)

Local Rule CV-7(g) provides, in relevant part:

The court may refuse to hear or may deny a nondispositive motion unless the movant advises the court within the body of the motion that counsel for the parties have conferred in a good-faith attempt to resolve the matter by agreement and certifies the specific reason that no agreement could be made.

Here, Neural’s certificate of conference stated the following:

The undersigned certifies that counsel for Neural AI, LLC conferred in good faith with counsel for non-party Google, LLC regarding the issues raised in this Motion, including Zoom conferences on or about July 25, August 1, and August 17, 2026, and related email correspondence. Despite those efforts, the parties were unable to resolve the dispute.

Mot. at 12. The certificate is inaccurate. Counsel for Neural and Google did not confer “regarding the issues raised in this Motion.” *Id.* As the record shows, the parties spent their time discussing the scope of the declaration Neural wanted Google to sign; Google’s general and specific objections to the discovery sought in the Subpoenas were never addressed. Yet the

“issues raised in this Motion” are *entirely* about the discovery sought in the Subpoenas. Mot. at 6–10 (requesting that the Court compel Google to produce documents responsive to Requests 1–12 and to produce a witness on Deposition Topics 1–5).

There was ample opportunity for Neural to engage with Google about the objections Google raised. Google provided Neural with both general and specific objections to the Subpoenas, including, *inter alia*, that they seek information “not proportionate to the needs of the case,” Dkt. 1-10 at 3, use vague and overbroad definitions like the one for “NVIDIA GPUs,” *id.* at 4, seek “source code or the equivalent of a full source code review,” *id.* at 6–7, and impose an improper “temporal scope of September 13, 2018 to the present,” *id.* at 8. *See also id.* at 10–31 (specific objections to document requests); *id.* Dkt. 1-11 at 9–14 (specific objections to deposition topics). The parties’ discussions nevertheless focused entirely on the ancillary declaration that Neural wanted Google to sign. But the declaration is not what “the issues raised in this Motion” are about, Mot. at 12; and Neural’s failure to confer with Google on either the “revised declaration” or Google’s objections is an independent ground to deny the motion. *See Gonzalez v. Int’l Med. Devices, Inc.*, No. 1:24-cv-00982, 2025 WL 3453956, at *2 (W.D. Tex. Nov. 14, 2025) (“Gonzalez has not satisfied the conference requirement of Local Rule CV-7(g) because he has not actually talked with Defendants about the alleged deficiencies.”).

Whether driven by expediency or some other litigation strategy,³ Neural ignored the local rules, which required the parties to discuss “each item” underlying the motion to compel before running to the courthouse. *Anzures v. Prologis Texas I LLC*, 300 F.R.D. 314, 315–16 (W.D.

³ Neural sought a similar “declaration” from the other non-parties it is suing in federal court for discovery. *See, e.g., Neural AI v. Meta Platforms*, No. 7:26-mc-00327-LS, Dkt. 1 at 1 (explaining that Neural supplied “a draft declaration” and “invited Meta to propose revisions or provide equivalent information in lieu of broader discovery”).

Tex. 2012). That did not happen here. As a consequence, Neural’s premature motion wastes both party and court resources. *Id.* at 316 (“It appears that it may require several hours of court time to resolve the numerous issues raised; it seems logical that the parties will have spent an equal or greater amount of time attempting to resolve the issues.”); *In re Presto*, 358 B.R. 290, 293 (S.D. Tex. 2006) (“It is vitally important that counsel confer with one another in good faith, and so represent to the Court, before taking up court time. . .”). At bottom, Neural’s certificate that the parties “conferred in a good-faith attempt to resolve the matter” cannot be squared with the record, which shows that the parties never conferred on “the matter” now before the Court, *i.e.*, the twelve document requests and five deposition topics. Ex. A at 2; *id.* at 4 (“We have not discussed any of the specific requests or topics at any point. Instead, Neural AI has chosen to focus on a form declaration, of which it only provided a revised version earlier today.”). Neural failed to comply with Local Rule CV-7(g).

B. The Subpoenas Do Not Comply with Federal Rule 45

If the Court does not deny the instant motion based on Neural’s failure to appropriately meet and confer, Google respectfully requests that it deny the motion on substantive grounds.

The Subpoenas fail to comply with Federal Rule 45. To begin, Neural has not demonstrated the requested discovery is relevant. Neural says it seeks evidence to support indirect infringement allegations against NVIDIA, but Neural has not shown an evidentiary basis that Google uses NVIDIA GPUs or software in an allegedly infringing configuration. Neural is simply **guessing** that Google possesses relevant information merely because Google buys NVIDIA GPUs. This is laid bare by Neural’s indiscriminate “shotgun” approach to non-party discovery, in which it has served carbon-copy subpoenas and filed undifferentiated motions against a wide swath of third parties. Indeed, several of the requests ask **whether** Google uses NVIDIA products in a particular way. This type of speculation is improper, particularly against

a non-party. “A litigant may not engage in merely speculative inquiries in the guise of relevant discovery.” *Micro Motion*, 894 F.2d at 1327-28 (rejecting third party subpoena that was a “fishing expedition.”); *see also Strong v. Paradise*, No. 3:23-cv-2847-K, 2025 WL 1811766, at *3 (N.D. Tex. July 1, 2025) (“[U]nder Rule 45, non-parties have greater protections from discovery than parties do.”) (internal quotations omitted).⁴

The Subpoenas are also unduly burdensome, for at least the following reasons:

Temporal Scope: Neural argues that the Subpoenas are “limited” to September 13, 2018 to the present.” Mot. at 9. An eight-year period is not “limited,” and enhances the burden on non-party Google. *See, e.g., Treadway v. Otero*, No. 2:19-cv-244, 2020 WL 602225, at *3 (S.D. Tex. Feb. 7, 2020) (“Plaintiff’s request for production of these documents for a period of eight years, from 2010-2018, is unduly burdensome.”).

NVIDIA GPUs: Eleven of the twelve document requests rely on Neural’s definition of “NVIDIA GPUs,” which in effect turns each individual document request into “over 240 requests.” *See, e.g.,* Dkt. 1-10 at 31. Take for example Request No. 1, which seeks documents “sufficient to identify all software, frameworks, libraries, APIs, scripts, Source Code, configuration files, and custom code You use to perform computations on NVIDIA GPUs.” Dkt. 1-6 at 13. To properly respond to this request, Google would need to identify *all* [1] software, [2] frameworks, [3] libraries, [4] APIs, [5] scripts, [6] Source Code, [7] configuration files, *and* [8] custom code that Google uses to perform computations on *over 240 products*. Responding to this request (and the ten others using the “NVIDIA GPUs” definition) would require Google to produce “voluminous” records that are spread out across multiple departments at the company.

⁴ Google further disagrees with Neural’s view that this district is the appropriate place for compliance, Fed. R. Civ. P. 45(c), but given the other manifest deficiencies the subpoenas, is not relying on that issue here. *See* Dkts. 1-10, 1-11 at 2-3.

See State Farm Lloyds, 348 F.R.D. at 429. Jaffe Decl. ¶ 8. Requiring Google to search for information associated with “NVIDIA GPUs” over an eight-year period would be akin to asking Ford to catalog how it uses every steering wheel, and every associated blueprint for each car it makes that includes a steering wheel, over almost a decade.

Document Request No. 1: Request No. 1 is patently overbroad and unduly burdensome. It requests: “Documents sufficient to identify *all* software, frameworks, libraries, APIs, scripts, Source Code, configuration files, and custom code You use to perform computations on NVIDIA GPUs.” NVIDIA GPUS refers to the three-page definition discussed above, and “Source Code” is defined to include “*all* associated files necessary to understand, compile, and execute the code, such as scripts, header files, makefiles, configuration files, and documentation.” Dkt. 1-6 at 10. “Source Code” is further defined to include “*all versions* and revisions relevant to the time periods and subject matter.” This request thus seeks an exhaustive catalog of every piece of software, framework, library, API, script, source code, configuration file, and custom code that Google purportedly uses with at least ~240 NVIDIA GPUs and “every version” of any software that goes along with it for the past eight years. To even begin to assess how to comply with this request across the entirety of Google would be nearly impossible, and certainly could not be accomplished within any realistic timeframe. Jaffe Decl. ¶ 8. And this is for documentation which has not even been shown to be relevant or at a minimum proportional to the needs of the case under Rule 26.

The request further seeks Google’s highly proprietary software and workflow for its AI software and services, comprising Google trade secrets. Google should not be required to disclose some of its most highly confidential and competitive technical information without a sufficient demonstration of relevance and proportionality. *See Leonardo Worldwide Corp. v.*

Pegasus Sols., Inc., 2015 WL 13469920, at *4 (N.D. Tex. Apr. 16, 2015) (“Although the request may be tangentially relevant to a defense in the case, the request seems likely to procure far more irrelevant, burdensome, and potentially confidential information than it would procure information relevant to this action.”).

Document Requests Nos. 2–4: These requests are objectionable for the same reasons as Document Request No. 1. These requests are aimed at a laundry list of NVIDIA software or undefined sample code from NVIDIA. Request No. 4 goes so far as to seek information about “*any* software” Google uses “to perform computations on NVIDIA GPUs.” Dkt. 1-6 at 13. As with Request No. 1, these requests are boundless dragnets that lack any defining parameters or logical boundaries. Compliance would require a sprawling, unfocused investigation across the entirety of Google. Jaffe Decl. ¶ 8. The rules do not license such an indiscriminate fishing expedition, particularly against a non-party. *Micro Motion*, 894 F.2d at 1327-28; *Paradise*, 2025 WL 1811766, at *3.

Document Request No. 5: This request is even broader than Document Request No. 1. This request seeks documents to describe the architecture of “*any* system” at Google that uses NVIDIA GPUs to perform computations, including every version of source code over the past eight-years. For a company at the forefront of AI like Google, this would be a Herculean task. Neural has not, and cannot, justify the relevance of this request when compared to the disproportionate burden imposed on Google.

Document Requests No. 6–12: These requests are Neural’s attempt to shift the burden to Google to substantiate Neural’s infringement claims. *Compare* Dkt. 1-6 at 14 (“Documents sufficient to show whether You use a pointer to data stored in memory... swapping an input pointer with the pointer to data output from a GPU computation...”) *with* U.S. Patent No.

8,648,867 at col. 14:40–52 (claiming an accelerator controller . . . to swap the first pointer and the second pointer”). This is an unabashed fishing expedition with no demonstrated relevance to Google, *Micro Motion*, 894 F.2d at 1327–28, which also imposes undue burden on Google because of the highly technical (and legal) nature of the requests.

Deposition Topics Nos. 1–5: As Neural readily admits, the deposition topics seek information “on the same subjects as the document requests,” Mot. at 13, so necessarily suffer from the same flaws as the document requests, and in some instances even more so. For example, Topic No. 1 asks Google to provide a deponent to testify about **28 different** “NVIDIA software and libraries” used across Google’s entire computing infrastructure, Dkt. 1-6 at 24; Topic No. 2 adopts the unreasonably broad definition of “Source Code,” *id.*; Topic No. 3 asks for a witness to testify about every “customization” or “data input” Google has made to any NVIDIA software (whatever that means), and a description of this so-called “data input,” *id.*; Topic No. 4 requests that a witness testify to the identity of **all software** at Google “that uses NVIDIA GPUs to perform computation,” which as discussed above, would cover over 240 products, *id.*, and Topic No. 5 is analogous to Document Requests No. 6–12, seeking highly technical information and, arguably, a deponent with experience in patent law who can testify to the claim limitations at issue in Neural’s underlying litigation against NVIDIA. The deposition topics are improper.

IV. CONCLUSION

For the foregoing reasons, Google respectfully requests that the Court deny Neural’s motion to compel.⁵

⁵ Google reserves the right to seek its fees under Fed. R. Civ. P. 45(d)(1) based on Neural’s failure to take reasonable steps to avoid imposing undue burden or expense on a person subject to the subpoena.

Dated: August 25, 2026

/s/ Jason M. Storck

Jordan R. Jaffe (*pro hac vice forthcoming*)

jjaffe@wsgr.com

David Fox (*pro hac vice forthcoming*)

dfox@wsgr.com

WILSON SONSINI GOODRICH & ROSATI

One Market Street, Spear Tower, Suite 3300

San Francisco, CA 94105

Telephone: (415) 947-2000

Jason M. Storck, Texas Bar No. 24037559

jstorck@wsgr.com

WILSON SONSINI GOODRICH & ROSATI

900 South Capital of Texas Highway

Las Cimas IV, Fifth Floor

Austin, TX 78746

Attorneys for Respondent Google LLC

CERTIFICATE OF SERVICE

The undersigned certifies that, on August 25, 2026, all counsel of record are being served with a copy of this document via CM/ECF to the following:

Susman Godfrey

Brian D. Melton

1000 Louisiana St.

Suite 5100

Houston, TX 77002

bmelton@susmangodfrey.com

Max L. Tribble , Jr.

1000 Louisiana

Suite 5100

Houston, TX 77002-5096

mtribble@susmangodfrey.com

713-653-7820

Rocco Magni

1000 Louisiana, Suite 5100

Houston, TX 77002-5096

rmagni@susmangodfrey.com

Tanner H. Laiche

401 Union St., Suite 3000

Seattle, WA 98101

tlaiche@susmangodfrey.com

206-516-3880

Attorneys for Petitioner Neural AI, LLC

/s/ Jason M. Storck

Jason M. Storck