IndexNow Speed Test: How Fast Do Bing, (2026)

June 22, 2026
https://tygartmedia.com/indexnow-speed-test-bing-gpt-google-crawl-comparison/

IndexNow promises instant content discovery. But how fast is it really? We ran a controlled speed test — 40 articles published simultaneously to tygartmedia.com with IndexNow pings fired on every one — then measured exactly how long it took Bing, GPTBot, Google, and every other crawler to show up. The timestamps tell a story that IndexNow's marketing materials do not.

This is the second article in Tygart Media's AI Search Intelligence series, based on proprietary server log data from our 40-article Microsoft Copilot content experiment conducted on June 22, 2026. Every timestamp and crawl interval cited here comes directly from our server access logs.

What Is IndexNow and Why Speed Matters

What IndexNow is and why speed matters.

IndexNow is an open-source protocol that lets websites notify participating search engines the moment content is published or updated. Instead of waiting for a crawler to discover your new page organically — which can take days or weeks — IndexNow sends a direct ping saying "this URL has new content, come get it."

Microsoft developed IndexNow and Bing is its primary participant. Yandex, Naver, Seznam, and several other engines also participate. Google does not. As of early 2026, over 60 million websites use IndexNow, and 22% of clicked Bing URLs come from IndexNow submissions, according to Bing's published data.

For publishers, the speed question is not academic. If you are publishing time-sensitive content — news, product launches, competitive analysis — the difference between a 3-hour crawl delay and a 3-day crawl delay determines whether your content gets indexed before or after your competitors. And in the AI era, the question extends beyond traditional indexing: how fast do AI crawlers like GPTBot find your new content?

Our Test Setup: 40 Articles, One Timestamp

On June 22, 2026, we published 40 original articles about Microsoft Copilot to tygartmedia.com. The site runs WordPress with RankMath SEO on a Google Cloud Platform Compute Engine instance. RankMath handles IndexNow submissions automatically on publish.

Every article was published within a short window, and IndexNow pings were fired for each URL. We then monitored our raw server access logs for every subsequent crawler visit, recording the user-agent string, timestamp, and requested URL for each hit.

This gave us a clean dataset: 40 identical test cases (same site, same publish time, same IndexNow submission) with crawler-by-crawler arrival times we could compare head-to-head.

Head-to-Head Results: Who Arrived First?

Head-to-head: who arrived first.

Bing: 3 to 6 Hours via IndexNow

Bingbot was the first traditional search engine crawler to reach our content, arriving within **3 to 6 hours of IndexNow submission**. The pattern was remarkably consistent across all 40 articles — most fell within a tight 4-hour window from publication to first crawl.

This is fast by search engine standards but not instant. IndexNow does not trigger immediate crawling. It places your URL into Bing's priority crawl queue, and Bing processes that queue on its own schedule. For our batch of 40 articles, that schedule produced a 3-to-6-hour window with high consistency.

For context, without IndexNow, new content on a site with our domain authority profile might wait 24 to 72 hours for Bing to discover it through sitemap parsing or link following. IndexNow compressed that to under 6 hours — a meaningful improvement for any publishing operation.

GPTBot: Faster Than Bing

Here is the result that surprised us most: **GPTBot arrived at our content faster than Bingbot in many cases**, despite GPTBot not being an official IndexNow participant.

GPTBot is OpenAI's crawler. It does not receive IndexNow pings directly. Yet it consistently reached our newly published articles before Bing's own crawler had finished processing the IndexNow queue. At 11:00 UTC on June 22, GPTBot executed a **1,123-request structural crawl in a single hour**, hitting not just article URLs but every tag, feed, and REST API endpoint on the site (Tygart Media server log analysis, June 2026).

How does GPTBot discover content faster than IndexNow delivers it to Bing? The most likely explanation is that GPTBot monitors RSS feeds, sitemaps, or other real-time content signals independently. WordPress sites broadcast new content through multiple channels — RSS feeds update instantly, XML sitemaps regenerate on publish, and REST API endpoints reflect new posts immediately. GPTBot appears to be monitoring one or more of these channels with higher polling frequency than Bing's IndexNow processing queue.

The implication for publishers is significant: even if you do not use IndexNow, GPTBot is likely to find your new content quickly through other discovery mechanisms. But IndexNow remains essential for Bing-ecosystem discovery, which feeds Microsoft Copilot's citation pipeline.

YandexBot: 30 Seconds Behind Bing

YandexBot arrived at each article approximately **30 seconds after Bingbot**, with remarkable consistency across the full batch. Yandex participates in the IndexNow protocol, and this timing suggests Yandex processes IndexNow submissions from the same shared queue but with a slight processing delay relative to Bing (Tygart Media server log analysis, June 2026).

The 30-second shadow is too consistent to be coincidental. It points to either a shared IndexNow notification infrastructure where Yandex processes submissions fractionally behind Bing, or to Yandex monitoring Bing's crawl activity directly. Either way, publishers who submit to IndexNow get both Bing and Yandex coverage from a single ping.

Googlebot: Effectively Absent

Googlebot recorded only 1 hit on our Copilot content in the initial crawl window (Tygart Media server log analysis, June 2026). One hit. Across 40 articles. While Bing had crawled every article within 6 hours and GPTBot had mapped the entire site architecture.

Google does not participate in IndexNow. Google has stated publicly that it relies on its own crawl scheduling, which considers factors like site crawl budget, historical update frequency, and sitemap change signals. For a batch of 40 new articles on a topic the site had not previously covered, Google's algorithms apparently did not prioritize rapid discovery.

This is not a criticism of Google's approach — its crawl scheduling optimizes for different goals than real-time discovery. But for publishers who need content indexed quickly, the data is unambiguous:
IndexNow-participating engines discover content in hours. Google discovers it on its own timeline.

The IndexNow Technical Gotcha We Discovered

During our experiment, we identified a technical issue that could affect other publishers: **the IndexNow key file was returning a 404 at the standard verification paths** where search engines expect to find it.

IndexNow requires a verification key file at your site root (e.g., yourdomain.com/{key}.txt). Search engines check this file to confirm you authorized the IndexNow submission. In our case, the key file was not accessible at the expected root-level path, which should have caused verification failures.

RankMath SEO's fallback mechanism saved us — it handles IndexNow key verification through an alternative method that does not require the physical key file to exist at the root URL. But publishers using manual IndexNow implementations, or other SEO plugins without this fallback, should verify their key file is accessible by navigating directly to the expected URL.

If your IndexNow submissions seem to be ignored by Bing, check the key file first. A 404 on the verification file silently kills the entire pipeline — Bing will not crawl the submitted URLs without successful verification.

What the Speed Test Means for Your Publishing Strategy

What the speed test means for publishing strategy.

For Bing and Copilot Visibility

IndexNow is the fastest path to Bing's index, and Bing's index feeds Microsoft Copilot's citation system. Our 40-article experiment earned 3 confirmed Copilot citation referrals within 48 hours, and that pipeline started with IndexNow getting our content into Bing's index within hours of publication.

If you are publishing content that you want Copilot to cite, IndexNow is not optional — it is the first link in the citation chain.

For AI Crawler Discovery

GPTBot does not use IndexNow, but it finds new content fast anyway — faster than Bing in our test. This means your site's real-time content signals (RSS feeds, sitemaps, REST API endpoints) are the discovery mechanism for OpenAI's crawler ecosystem. Keep these endpoints clean, accessible, and unblocked in your robots.txt if you want AI systems to discover your content quickly.

For Google

Google's crawl scheduling operates independently of IndexNow. If rapid Google indexing is important to you, continue submitting sitemaps through Google Search Console and requesting indexing for priority pages through the URL Inspection tool. Do not rely on IndexNow for Google discovery — the protocol has no effect

on Google's crawl behavior based on our data.

For Multi-Engine Strategy

The practical recommendation is to run both systems in parallel: IndexNow for Bing, Yandex, and the downstream AI systems that rely on Bing's index, plus Google Search Console for Google's independent crawl pipeline. Most WordPress SEO plugins handle IndexNow automatically, so the incremental effort is near zero.

The Speed Hierarchy: From Fastest to Slowest

Based on our server log data from the 40-article experiment, here is the definitive crawl speed ranking for newly published, IndexNow-submitted content (Tygart Media server log analysis, June 2026):

- **GPTBot** — fastest overall; arrived before IndexNow results in many cases; 1,123-request structural crawl in one hour
- **ChatGPT-User** — 3,404 hits over 48 hours; activates when real users query ChatGPT about relevant topics
- **Bingbot** — 3 to 6 hours via IndexNow; consistent, predictable timing
- **YandexBot** — ~30 seconds behind Bingbot; piggybacks on IndexNow shared infrastructure
- **OAI-SearchBot** — 3 hits total; minimal presence; appears highly selective
- **AzureAI-SearchBot** — 3 hits total; minimal presence
- **Googlebot** — 1 hit in initial window; operates on its own schedule independent of IndexNow

The gap between the top of this list and the bottom is not hours — it is the difference between same-day discovery and multi-day (or longer) discovery. For publishers who need content discovered quickly, the AI crawlers and IndexNow-participating engines are delivering results that Google's independent crawl schedule simply does not match.

A Note on Methodology and Reproducibility

Every crawl timestamp and interval cited in this article comes from raw server access logs on Tygart Media's Google Cloud Platform Compute Engine instance, analyzed in June 2026. Crawler identification was performed by user-agent string matching, with IP range verification against OpenAI's and Microsoft's published crawler IP ranges for additional confirmation.

The 40-article batch was published simultaneously to control for timing variables. All articles were submitted via IndexNow through RankMath SEO's automatic submission feature. No manual crawl requests were submitted through Google Search Console, Bing Webmaster Tools, or any other interface — we wanted to measure organic and IndexNow-driven discovery only.

This experiment is reproducible. Any publisher running a WordPress site with IndexNow enabled can monitor their server access logs after a batch publish and observe the same crawler patterns. The specific timing intervals may vary based on domain authority, server location, and crawl budget allocation, but the relative ordering — GPTBot fastest, Bing via IndexNow in hours, Google on its own schedule — should hold across most publishing environments.

For the complete dataset including all crawler hit counts and the full methodology, see our anchor article: We Published 40 Articles and Watched Every AI Crawler in Real Time.

Frequently Asked Questions

How fast does IndexNow actually get content crawled by Bing?

In our controlled test of 40 simultaneously published articles, IndexNow submissions resulted in first Bingbot crawls within 3 to 6 hours, with most articles falling in a consistent 4-hour window. This is significantly faster than the 24-to-72-hour organic discovery timeline for sites without IndexNow, but it is not instant — Bing queues IndexNow submissions and processes them on its own crawl schedule (Tygart Media server log analysis, June 2026).

Does GPTBot use IndexNow to discover content?

No. GPTBot is not an IndexNow participant, yet it arrived at our content faster than Bingbot in many cases. GPTBot appears to monitor RSS feeds, XML sitemaps, or REST API endpoints independently, giving it a faster discovery pipeline than Bing's IndexNow processing queue. In our experiment, GPTBot executed a 1,123-request structural crawl at 11:00 UTC, mapping the entire site architecture within a single hour (Tygart Media server log analysis, June 2026).

Does Google support IndexNow?

No. Google does not participate in the IndexNow protocol as of June 2026. In our experiment, Googlebot recorded only 1 hit on our 40-article batch while Bingbot and GPTBot had fully crawled the content. Google relies on its own crawl scheduling algorithms and recommends using Google Search Console's sitemap submission and URL Inspection tool for prioritized crawling (Tygart Media server log analysis, June 2026).

Why was YandexBot always 30 seconds behind Bingbot?

YandexBot, as an IndexNow participant, appears to process submissions from a shared notification infrastructure with a slight delay relative to Bing. The consistent 30-second gap across all 40 articles suggests either a shared queue processed fractionally behind Bing or direct monitoring of Bing's crawl activity. The practical result is that a single IndexNow ping delivers both Bing and Yandex crawls almost simultaneously (Tygart Media server log analysis, June 2026).

What should publishers do if IndexNow submissions are being ignored by Bing?

Check your IndexNow key file first. The key file must be accessible at your domain root (e.g., yourdomain.com/{key}.txt). In our experiment, the key file was returning a 404 at standard paths, which would have silently killed the pipeline. Our RankMath SEO plugin's fallback mechanism handled verification, but publishers using manual implementations should navigate directly to their key file URL to confirm it returns a 200 response (Tygart Media server log analysis, June 2026).