Millstone AI Solutions

On-Prem AI Solutions

BENCHMARK REPORT

Beam-1.0-501B-A23B

Performance Analysis on 4x RTX Pro 6000 Blackwell

MODEL

Organization	**Reflection**
Parameters	**501B**
Precision	**NVFP4**

TEST HARDWARE

GPU	**4x RTX Pro 6000 Blackwell**
VRAM	**384GB**
Engine	**vLLM**

HIGHLIGHTS

392.5
Tok/s Peak Throughput
@ 10 Concurrent Requests

100.0%
Success Rate
Across All Scenarios

~16
Concurrent Users
@ 32K Context

October 17, 2026

MillstoneAI.com

CONTENTS

Table of Contents

Executive Summary 3

Use Case Guidance 4

Performance Analysis 5

- System Throughput 5
- Per-User Generation Speed 6
- Time to First Token 7

Capacity Analysis 8

- Code Completion (1K Context) 8
- Short-form Chatbot (8K Context) 9
- General Chatbot (32K Context) 10
- Long Document Processing (64K Context) 11
- Automated Coding Assistant (96K Context) 12

Technical Deep Dive 13

- Queue Wait Times 13
- Per-User Prefill Speed 14
- Inter-Token Latency 15
- Scaling Efficiency 15

Power & Efficiency 16

> **Interactive Data Available Online**
> This report provides a static snapshot of benchmark results. For interactive charts with hover tooltips, exact data point values, and interpolated metrics, visit the full benchmark page:
> MillstoneAI.com/inference-benchmark/beam-1-0-501b-a23b-nvfp4-4x-rtx-pro-6000-blackwell

OVERVIEW

Executive Summary

This benchmark evaluates **Beam-1.0-501B-A23B** (Reflection, 501B parameters, Mixture-of-Experts) running in NVFP4 precision on 4x RTX Pro 6000 Blackwell (384GB VRAM).

This report measures usable capacity: how many concurrent requests a configuration can serve at a given context length before performance falls below accepted thresholds. The KV cache holds 3,067,683 tokens across all requests at once. Testing stopped at the highest concurrency that remained usable.

Test parameters: Context lengths from 1K - 1M tokens. Concurrency from 1 - 10 requests. 1024 output tokens per request. No prompt caching. FP8 KV cache. Speculative decoding enabled, tuned for multiple users. A single-user configuration reached ~150 tok/s but significantly slowed concurrent requests.

Benchmark methodology →

Key Findings

Peak System Throughput	392.5 tok/s @ 10 concurrent requests, 1K context
TTFT Single Request	140ms (1K context) → 201.4s (1024K context)
Generation Speed Single Request	119.4 tok/s (1K context) → 99.6 tok/s (1024K context)
Chatbot Capacity	~16 concurrent requests @ 32K context
Throughput Scaling	4.4× from 1 to 10 concurrent requests
Success Rate	100.0% across 2.6K requests

> Throughout this report, "**concurrent requests**" refers to simultaneous active requests. For applications with natural user pauses (chat interfaces, coding assistants), each request slot typically serves 4–5 active users.

RECOMMENDATIONS

Use Case Guidance

The table below maps this configuration's performance to common deployment scenarios. Capacity limits are where TTFT or generation speed falls below accepted thresholds for a comfortable user experience. Detailed charts and analysis for each use case are available in the Capacity Analysis section.

USE CASE	TTFT THRESHOLD	SPEED THRESHOLD	ANALYSIS
Code Completion	2s e2e	N/A	Supports **7** concurrent requests within accepted thresholds.
Short-form Chatbot	10s	10 tok/s	Supports **36** concurrent requests within accepted thresholds.
General Chatbot	8s	15 tok/s	Supports **~16** concurrent requests within accepted thresholds.
Long Document Processing	12s	15 tok/s	Supports **10** concurrent requests within accepted thresholds.
Automated Coding Assistant	12s	20 tok/s	Supports **~4** concurrent requests within accepted thresholds.

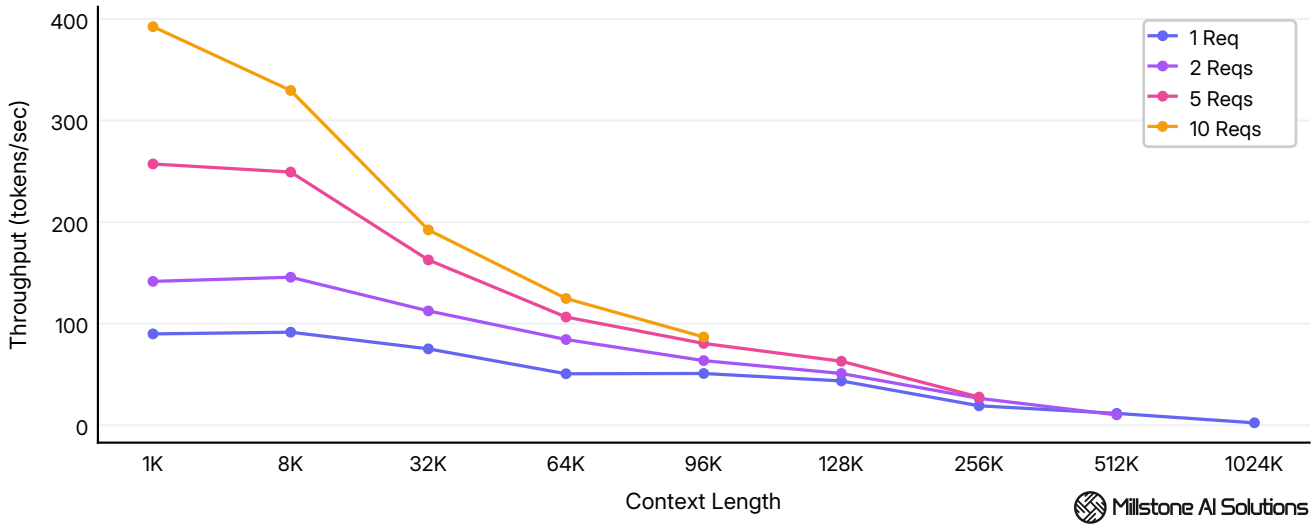
The limits shown are conservative. Beyond these points, the system continues functioning with slower response times that may still be acceptable for your specific use case.

> Want to validate your specific configuration? Request a Custom Benchmark →

PERFORMANCE

System Throughput

Aggregate token generation across all concurrent requests. Measures output tokens only. Prompt tokens processed during prefill are excluded.



Average system throughput across 1K - 1024K tokens context lengths at 1 - 10 concurrency levels.

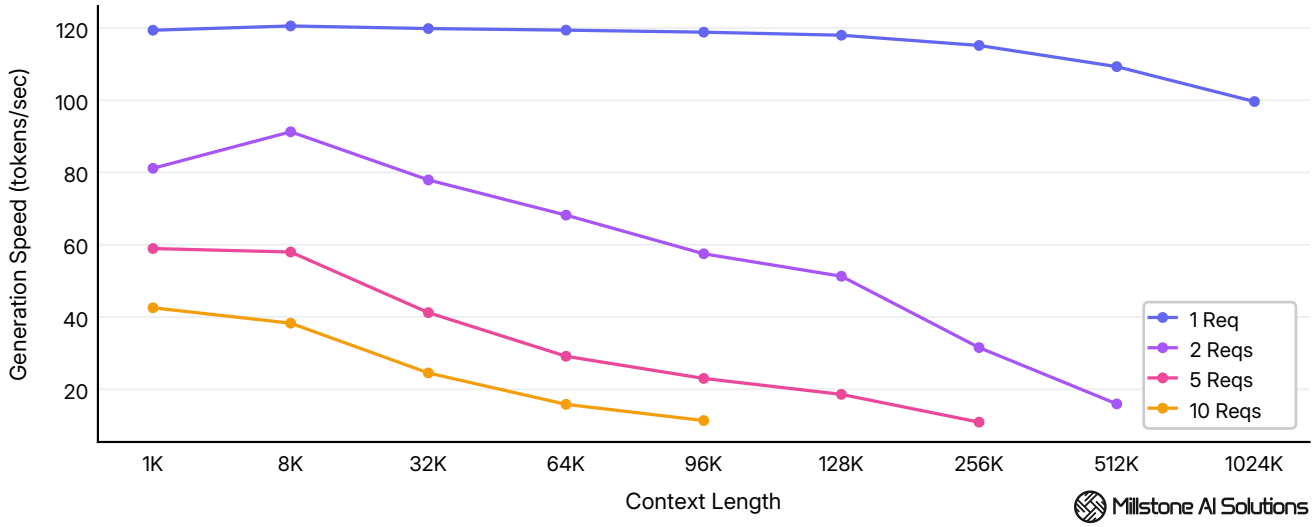
CONDITION	THROUGHPUT
Peak (1K context, 10 requests)	**392.5** tok/s
32K context, 10 requests	**192.3** tok/s
96K context, 10 requests	**86.9** tok/s

At peak throughput, this configuration produces approximately **1.4 million** tokens per hour. This is relevant for batch workloads like document processing, synthetic data generation, or offline analysis. Higher concurrency or shorter contexts can increase this further.

USER EXPERIENCE

Per-User Generation Speed

Token generation rate experienced by each individual user. This is the speed at which text streams into their response, also referred to as "decode speed" or "decode throughput." As concurrency increases, per-user speed decreases since GPU resources are shared across requests.



Average per-user generation speed across 1K - 1024K tokens context lengths at 1 - 10 concurrency levels.

How Fast is This?

SPEED	USER EXPERIENCE
< 15 tok/s	Slow; may be slower than reading speed
15–25 tok/s	Acceptable; keeps pace with reading
25–50 tok/s	Fast; exceeds reading speed
> 50 tok/s	Very fast; text appears nearly instantly

At **10.9** tok/s (the lowest measured point: 256K context, 5 concurrent requests), this configuration slows below fast reading speed in the most demanding scenarios. Single-user performance at 1K context reaches **119.4** tok/s.

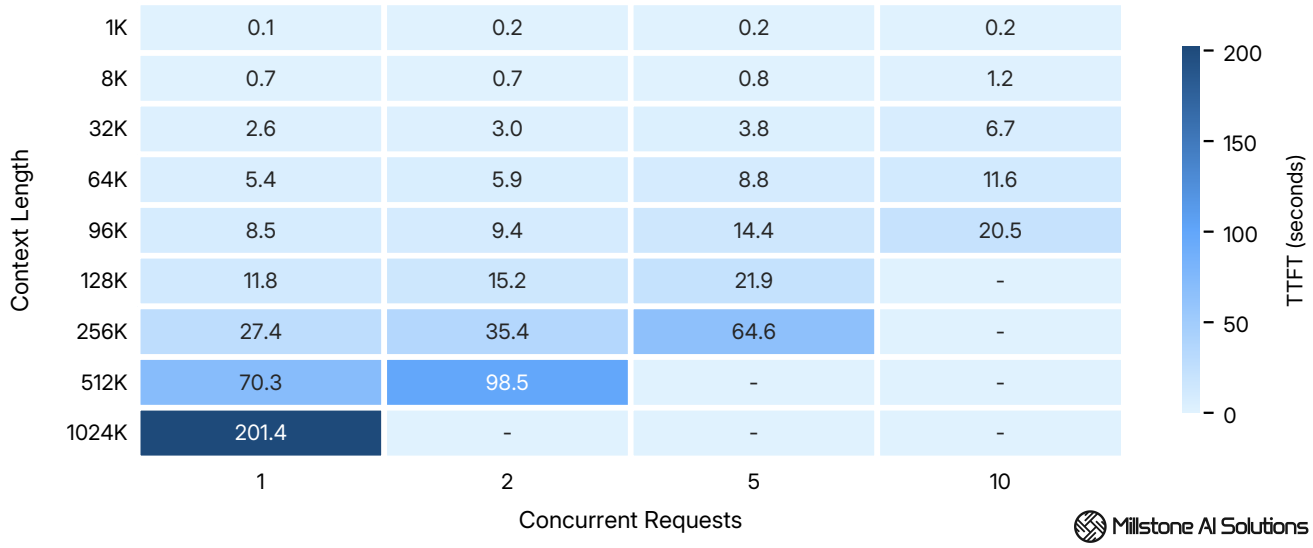
LATENCY

Time to First Token

Time from request submission to first response token. The primary metric for perceived responsiveness. TTFT has two components:

- **Queue wait:** Time waiting for GPU availability (increases with concurrency)
- **Prefill:** Time to process input context (increases with context length)

At low concurrency, prefill dominates. Under load, queue wait takes over. See Technical Analysis for more.



Average time to first token across 1K - 1024K tokens context lengths at 1 - 10 concurrency levels.

How Responsive is This?

TTFT	USER EXPERIENCE
< 500ms	Feels instant
500ms–2s	Feels responsive
2–5s	Noticeable but still acceptable
5–10s	Feels slow; generally acceptable at higher context lengths
> 10s	Can be frustrating; users may retry or abandon

Important note about caching. These benchmarks use fresh context with no caching enabled, representing worst-case TTFT. In production with caching enabled, only new tokens require processing. For example, a 64K conversation where you add 1K of new context would have a TTFT similar to the 1K results above, not the 64K results. **For most real-world use cases where context is built incrementally (chatbots, coding assistants, multi-turn agents), TTFT with caching enabled would be significantly faster than these results.**

CAPACITY PLANNING

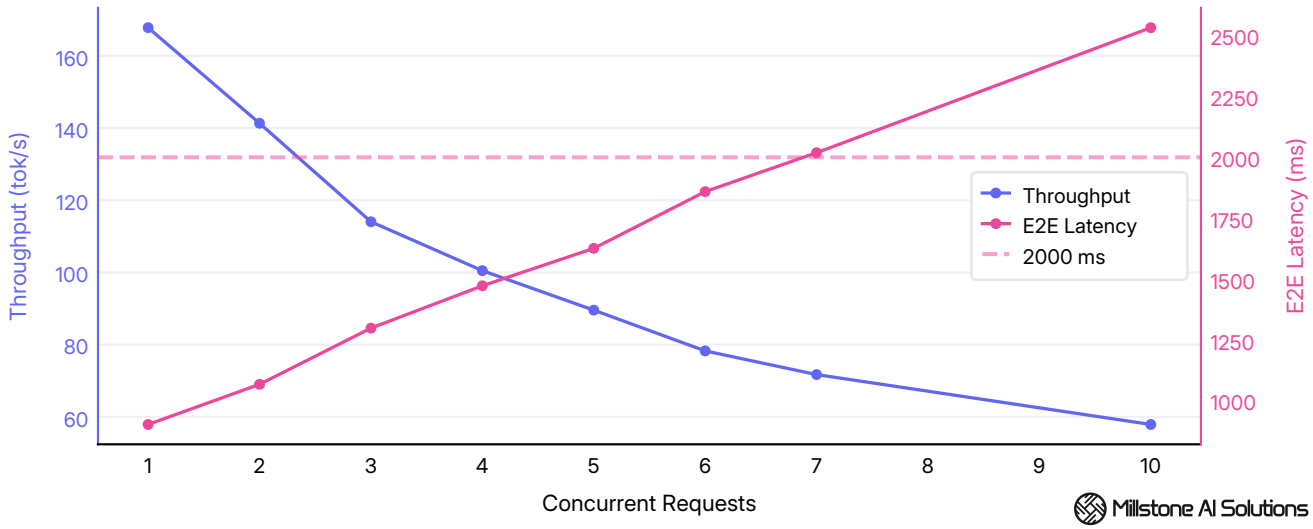
Capacity Analysis

How many concurrent requests can this configuration handle for different workloads? Each section below shows performance metrics as concurrency increases at a specific context length. Dashed lines indicate quality thresholds, the point where user experience degrades below acceptable levels. The "capacity limit" is the tested or estimated point where the first threshold is reached.

Code Completion (1K Context)

Inline code suggestions in IDEs, like autocomplete. Responsiveness is critical. This test generates 128 output tokens per request (vs. 1024 elsewhere) to match typical autocomplete length. The key metric is end-to-end latency, not TTFT.

Threshold: End-to-end latency < 2,000ms



Average end-to-end latency and throughput at 1K context. Dashed line indicates quality threshold.

METRIC	@ 1 request	@ 7 requests	@ 10 requests
End-to-end latency	902ms	2,019ms (threshold exceeded)	2,533ms (threshold exceeded)
Throughput	168 tok/s	72 tok/s	58 tok/s

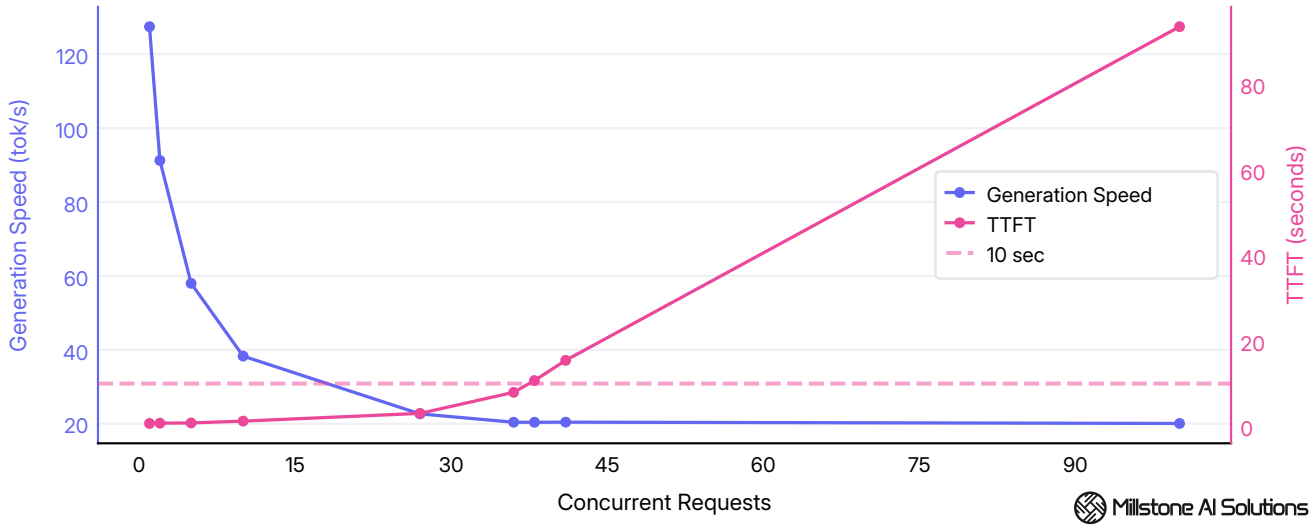
Capacity limit: 7 concurrent requests

At 7 concurrent requests, end-to-end latency reaches 2019ms, just above the 2,000ms threshold.

Short-form Chatbot (8K Context)

Quick conversational exchanges: customer support queries, simple Q&A, single-turn requests. 8K context accommodates a few back-and-forth messages plus system prompt. User expectations are more forgiving for these scenarios. 10+ tok/s is acceptable for reading streamed responses from a support chatbot.

Thresholds: TTFT < 10s, generation speed > 10 tok/s



Average per-user generation speed and TTFT at 8K context. Dashed lines indicate quality thresholds.

METRIC	@ 1 request	@ 36 requests	@ 100 requests
TTFT	0.7s	7.9s	93.5s (threshold exceeded)
Generation speed	127 tok/s	20 tok/s	20 tok/s

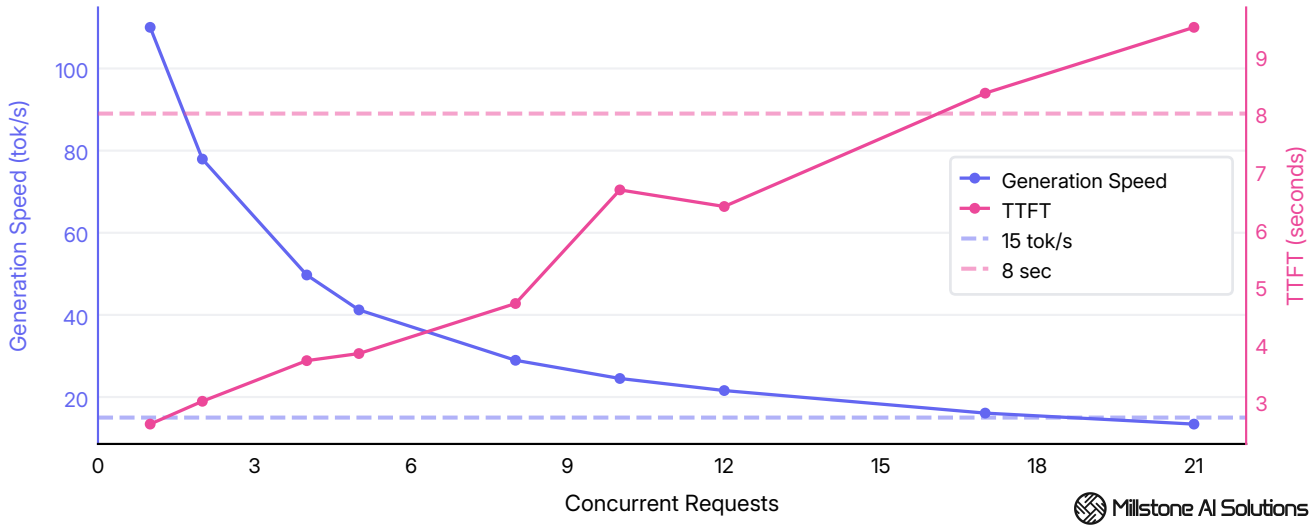
Capacity limit: 36 concurrent requests

At 36 concurrent requests, TTFT reaches 7.9 seconds, just under the 10-second threshold. Generation speed at this concurrency is 20 tok/s, above the 10 tok/s minimum.

General Chatbot (32K Context)

ChatGPT-style chatbot. If you're deploying a multi-turn conversational chatbot, this benchmark shows how many concurrent requests you can support while matching acceptable responsiveness.

Thresholds: TTFT < 8s, generation speed > 15 tok/s



Average per-user generation speed and TTFT at 32K context. Dashed lines indicate quality thresholds.

METRIC	@ 1 request	@ 16 requests	@ 21 requests
TTFT	2.6s	~8.0s	9.5s (threshold exceeded)
Generation speed	110 tok/s	~17 tok/s	13 tok/s (below threshold)

Capacity limit: ~16 concurrent requests

At 16 concurrent requests, TTFT reaches ~8.0 seconds, right at the threshold. Generation speed at this concurrency is ~17 tok/s, above the 15 tok/s minimum.

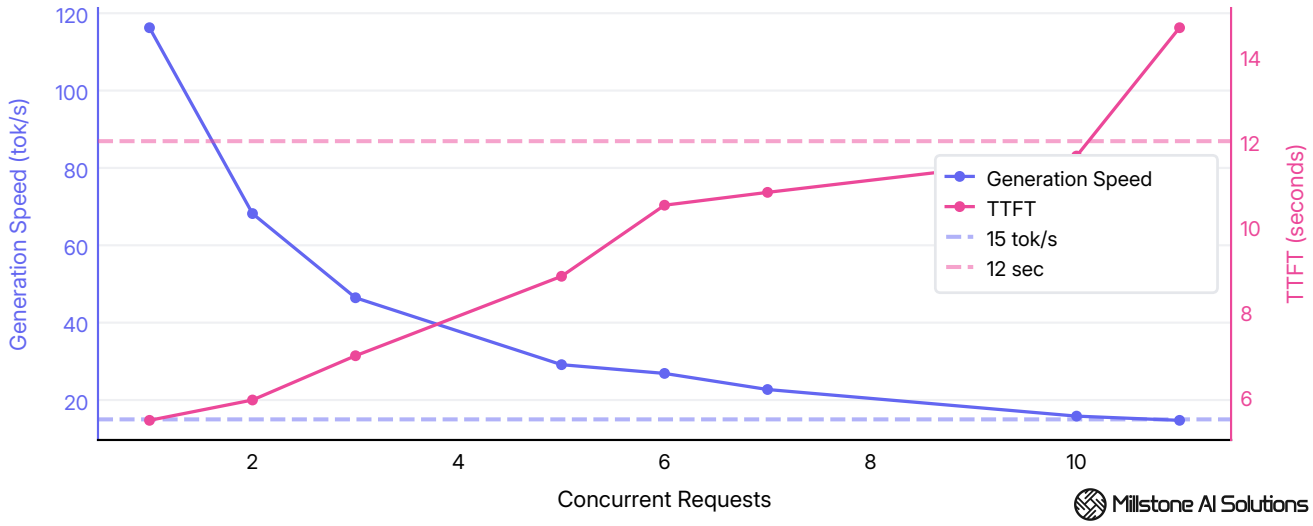
> **Note about caching:** Most chatbot users build context incrementally over a conversation. With caching properly configured, TTFT is dramatically reduced since only new tokens require processing. These results represent worst-case TTFT where all context is processed at once.

Long Document Processing (64K Context)

Summarizing reports, extracting data from contracts, analyzing lengthy documents. 64K tokens handles documents up to roughly 125-160 pages depending on formatting and density.

Users typically tolerate higher latency for document processing since they understand large inputs require more processing time. However, generation speed still needs to stay at or above reading speed.

Thresholds: TTFT < 12s, generation speed > 15 tok/s



Average per-user generation speed and TTFT at 64K context. Dashed lines indicate quality thresholds.

METRIC	@ 1 request	@ 10 requests	@ 11 requests
TTFT	5.4s	11.7s	14.7s (threshold exceeded)
Generation speed	116 tok/s	16 tok/s	15 tok/s (below threshold)

Capacity limit: 10 concurrent requests

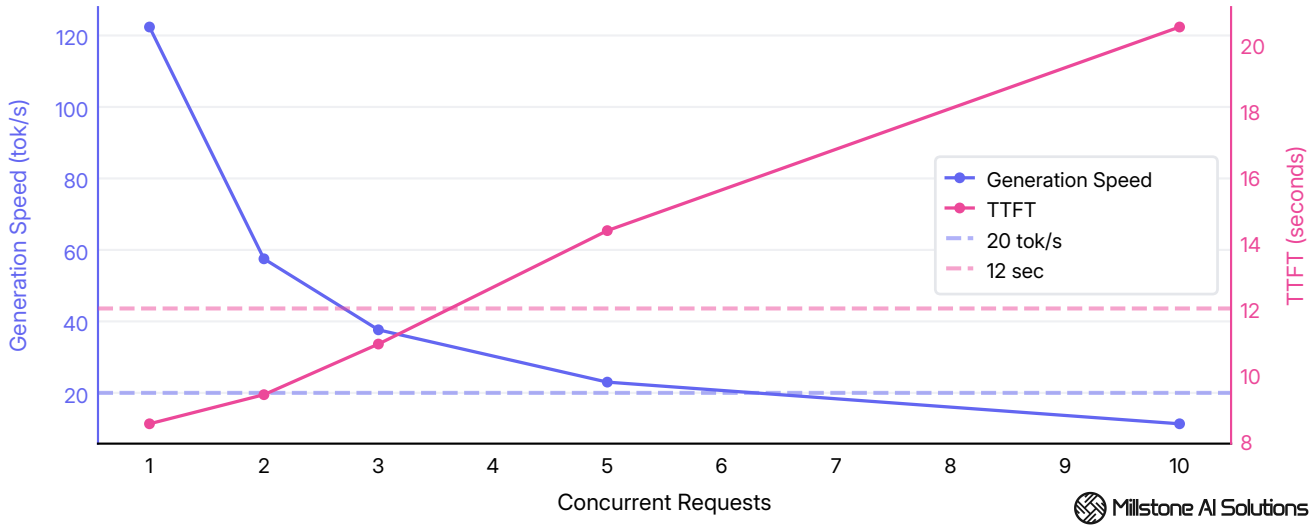
At 10 concurrent requests, TTFT reaches 11.7 seconds, just under the 12-second threshold. Generation speed at this concurrency is 16 tok/s, above the 15 tok/s minimum.

Automated Coding Assistant (96K Context)

Agentic coding workloads: AI assistants that read large portions of a codebase to answer questions, refactor code, or implement features. 96K tokens handles roughly 8,000-9,000 lines of code, enough for significant repository context.

Agentic workflows chain multiple LLM calls (tool use, retrieval, iterative refinement). With caching properly configured, context persists between requests and only new tokens require processing, dramatically reducing TTFT for each step. These results represent worst-case TTFT where all context is processed at once.

Thresholds: TTFT < 12s, generation speed > 20 tok/s



Average per-user generation speed and TTFT at 96K context. Dashed lines indicate quality thresholds.

METRIC	@ 1 request	@ 4 requests	@ 10 requests
TTFT	8.5s	~12.6s (threshold exceeded)	20.5s (threshold exceeded)
Generation speed	122 tok/s	~30 tok/s	11 tok/s (below threshold)

Capacity limit: ~4 concurrent requests

At 4 concurrent requests, TTFT reaches ~12.6 seconds, just above the 12-second threshold. Generation speed at this concurrency is ~30 tok/s, above the 20 tok/s minimum.

DEEP DIVE

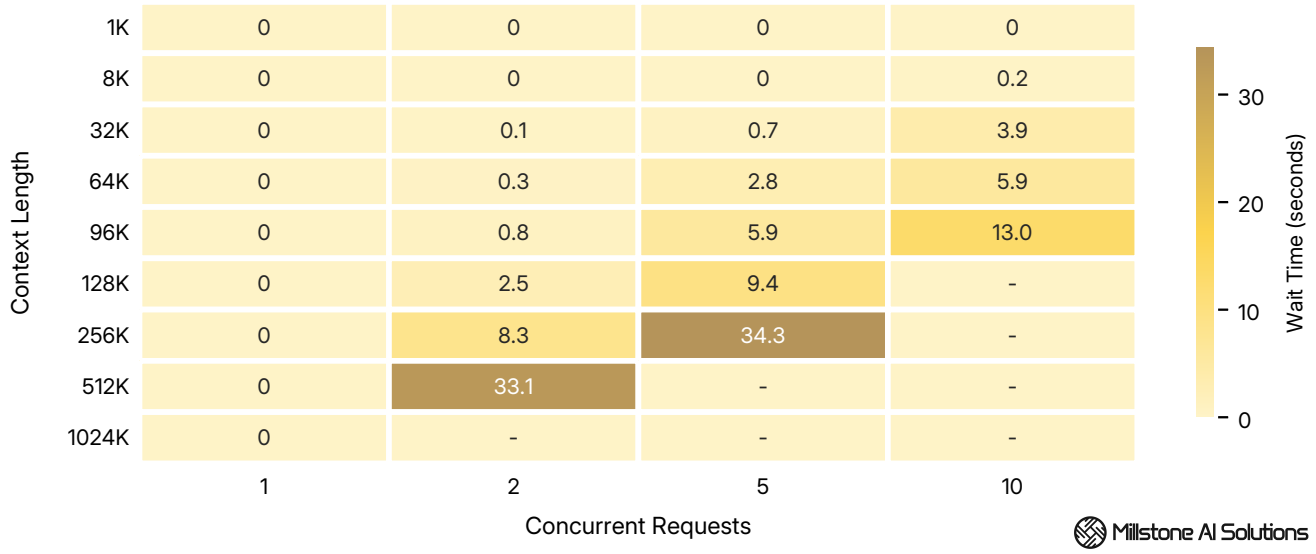
Technical Analysis

Infrastructure-level metrics that explain user-facing performance. Queue depth, prefill throughput, token generation latency, and scaling efficiency across load conditions. These help diagnose bottlenecks and validate infrastructure decisions.

Queue Wait Times

Time a request waits for GPU availability before processing begins. At low concurrency, queue wait is near zero. As load increases, requests queue and wait times grow.

Queue wait is included in TTFT. Breaking it out separately helps identify whether latency is caused by GPU saturation (high queue wait) or context processing (high prefill time).



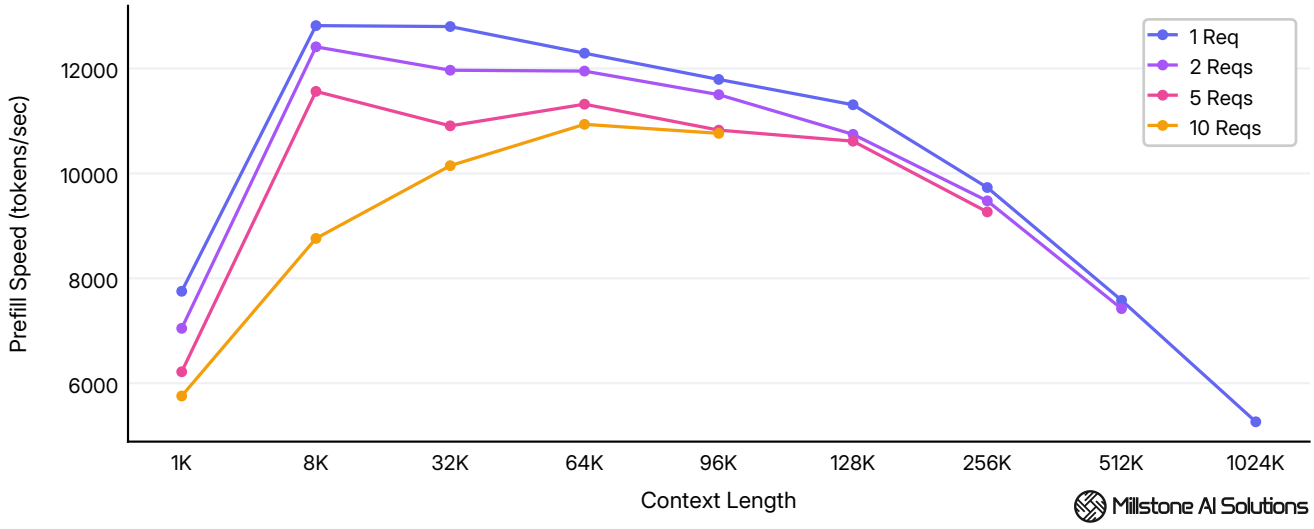
Average queue wait time across 1K - 1024K tokens context at 1 - 10 concurrent requests.

At single concurrency, queue wait is effectively zero regardless of context length. At **5 concurrent requests** with **256K context**, queue wait reaches **34.3 seconds**. Rising queue times signal GPU saturation, meaning requests are waiting for resources rather than being processed immediately.

> **Interpretation:** Queue wait time and prefill time are measured independently and may not sum exactly to TTFT. Under heavy load, chunked prefill and preemptions can cause these metrics to overlap, sometimes resulting in queue wait + prefill exceeding TTFT. Use queue wait for capacity planning and identifying bottlenecks. Use TTFT for actual user wait time before streaming begins.

Per-User Prefill Speed

Rate at which the model processes input context before generating output. Prefill speed determines the non-queue portion of TTFT. Higher prefill speeds mean faster time-to-first-token at a given context length.



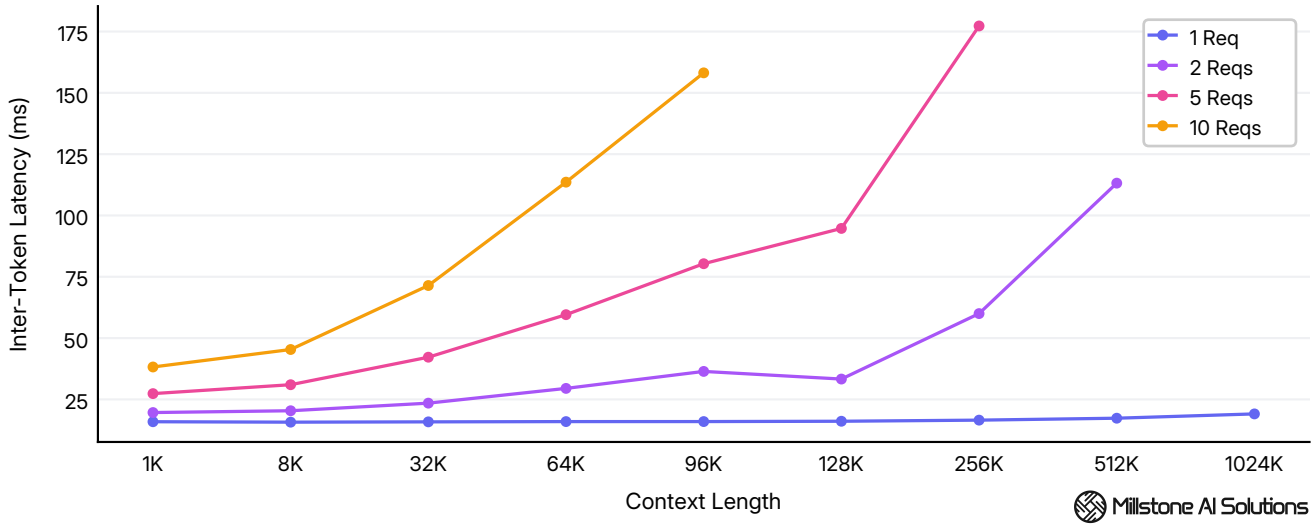
Average per-user prefill speed across 1K - 1024K tokens context at 1 - 10 concurrent requests.

CONCURRENT REQUESTS	PEAKS AT	PEAK SPEED
1	8K context	**12,818** tok/s
2	8K context	**12,414** tok/s
5	8K context	**11,564** tok/s
10	64K context	**10,936** tok/s

Prefill speed peaks at a certain context length and then declines as additional context increases computational overhead. This peak can reflect GPU saturation (compute or memory bandwidth fully utilized) or engine configuration such as chunked prefill limits, which cap tokens processed per forward pass to maintain responsiveness under load. On the chart, this appears as lines that peak before reaching the longest context.

Inter-Token Latency

Time between consecutive tokens during generation. Determines the smoothness of responses. Lower latency produces more fluid output. ITL helps diagnose the underlying token-level behavior.

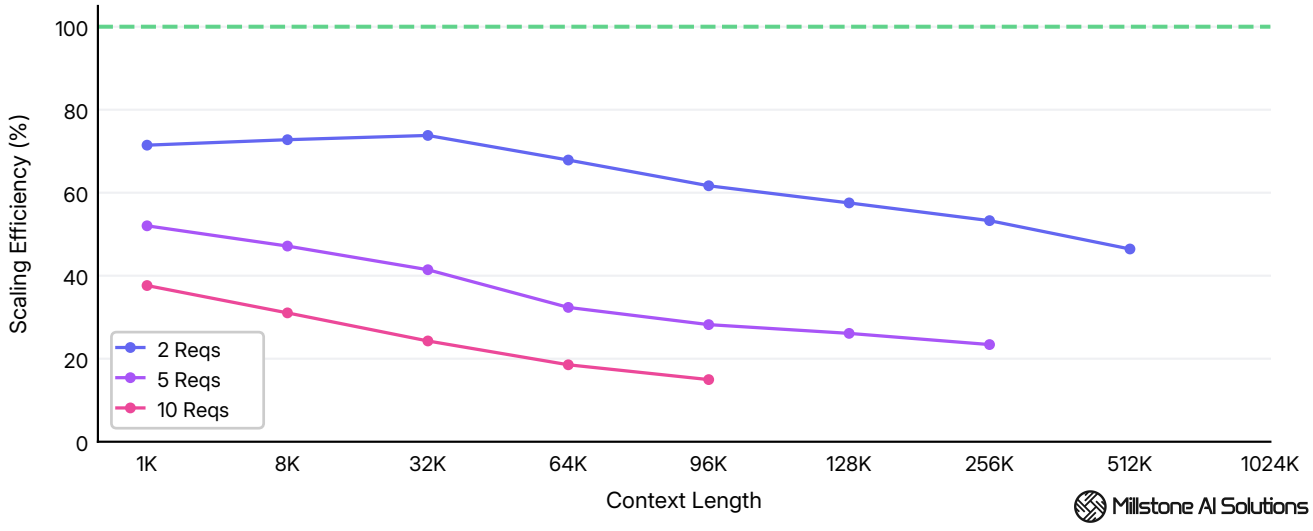


Average inter-token latency across 1K - 1024K tokens context at 1 - 10 concurrent requests.

At single-user short context, inter-token latency is imperceptible (**16ms**). The highest latency observed was **177ms** at **256K** context with **5 concurrent requests**, where individual tokens become visible as they stream.

Scaling Efficiency

Percentage of ideal linear scaling achieved as concurrency increases. 100% efficiency means doubling concurrent requests doubles total throughput with no per-user degradation. Real-world efficiency is always lower due to shared GPU resources.



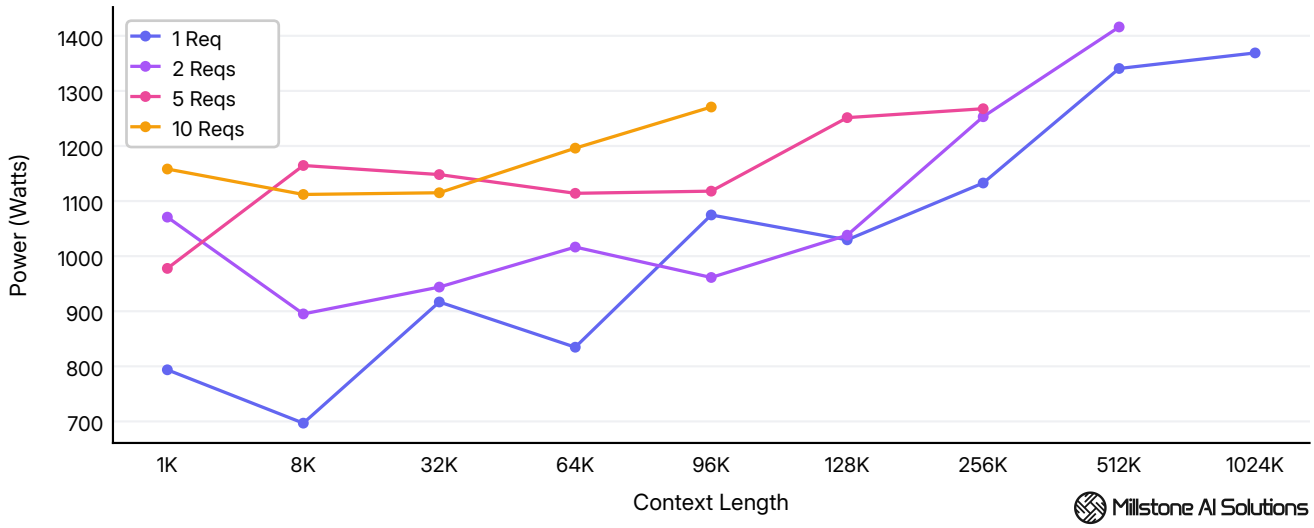
Scaling efficiency across 1K - 512K tokens context at 1 - 10 concurrent requests.

Efficiency remains high at low concurrency where GPU resources can serve requests without contention. At higher concurrency, efficiency drops as requests compete for shared resources. High efficiency at your target concurrency indicates headroom for growth. Sharply dropping efficiency signals diminishing returns.

EFFICIENCY

Power Consumption

GPU power draw under varying load conditions. Relevant for operational cost estimation, cooling requirements, and data center power budgeting.



Average GPU power draw across 1K - 1024K tokens context at 1 - 10 concurrent requests.

Power consumption scales with both context length and concurrency. The highest power draw observed was **1416W** at **512K** context with **2 concurrent requests**, costing approximately **$0.14/hour** at $0.10/kWh. Higher concurrency or sustained load beyond tested conditions may increase power consumption further. For infrastructure planning, budget for peak power draw.

Need Help Deciding?

Not sure what configuration you need? Our team can help you identify the right model, hardware, and deployment strategy for your specific use case.

Schedule a Conversation →

Additional data available on request: full percentile breakdowns (P50–P99) and GPU metrics.