

S62797 - LLM Inference Sizing: Benchmarking End-to-End Inference Systems

Dmitry Mironov

Solutions Architect, NVIDIA

Sergio Perez

Solutions Architect, NVIDIA

About Us

Let's coordinate inference conversation



Dmitry Mironov, EMEA
dmitrym@nvidia.com

- Senior Deep Learning Solutions Architect @ NVIDIA - Supporting deployment of AI / Deep Learning solutions
- Focusing on large scale efficient deployment and inference
- Co-author of NeMo Inference Sizing Calculator



Sergio Perez, EMEA
sergiop@nvidia.com

- Senior Deep Learning Solutions Architect @ NVIDIA - Supporting delivery of AI / Deep Learning solutions
- Focusing on quantization in training and inference
- Co-author of NeMo Inference Sizing Calculator



Agenda

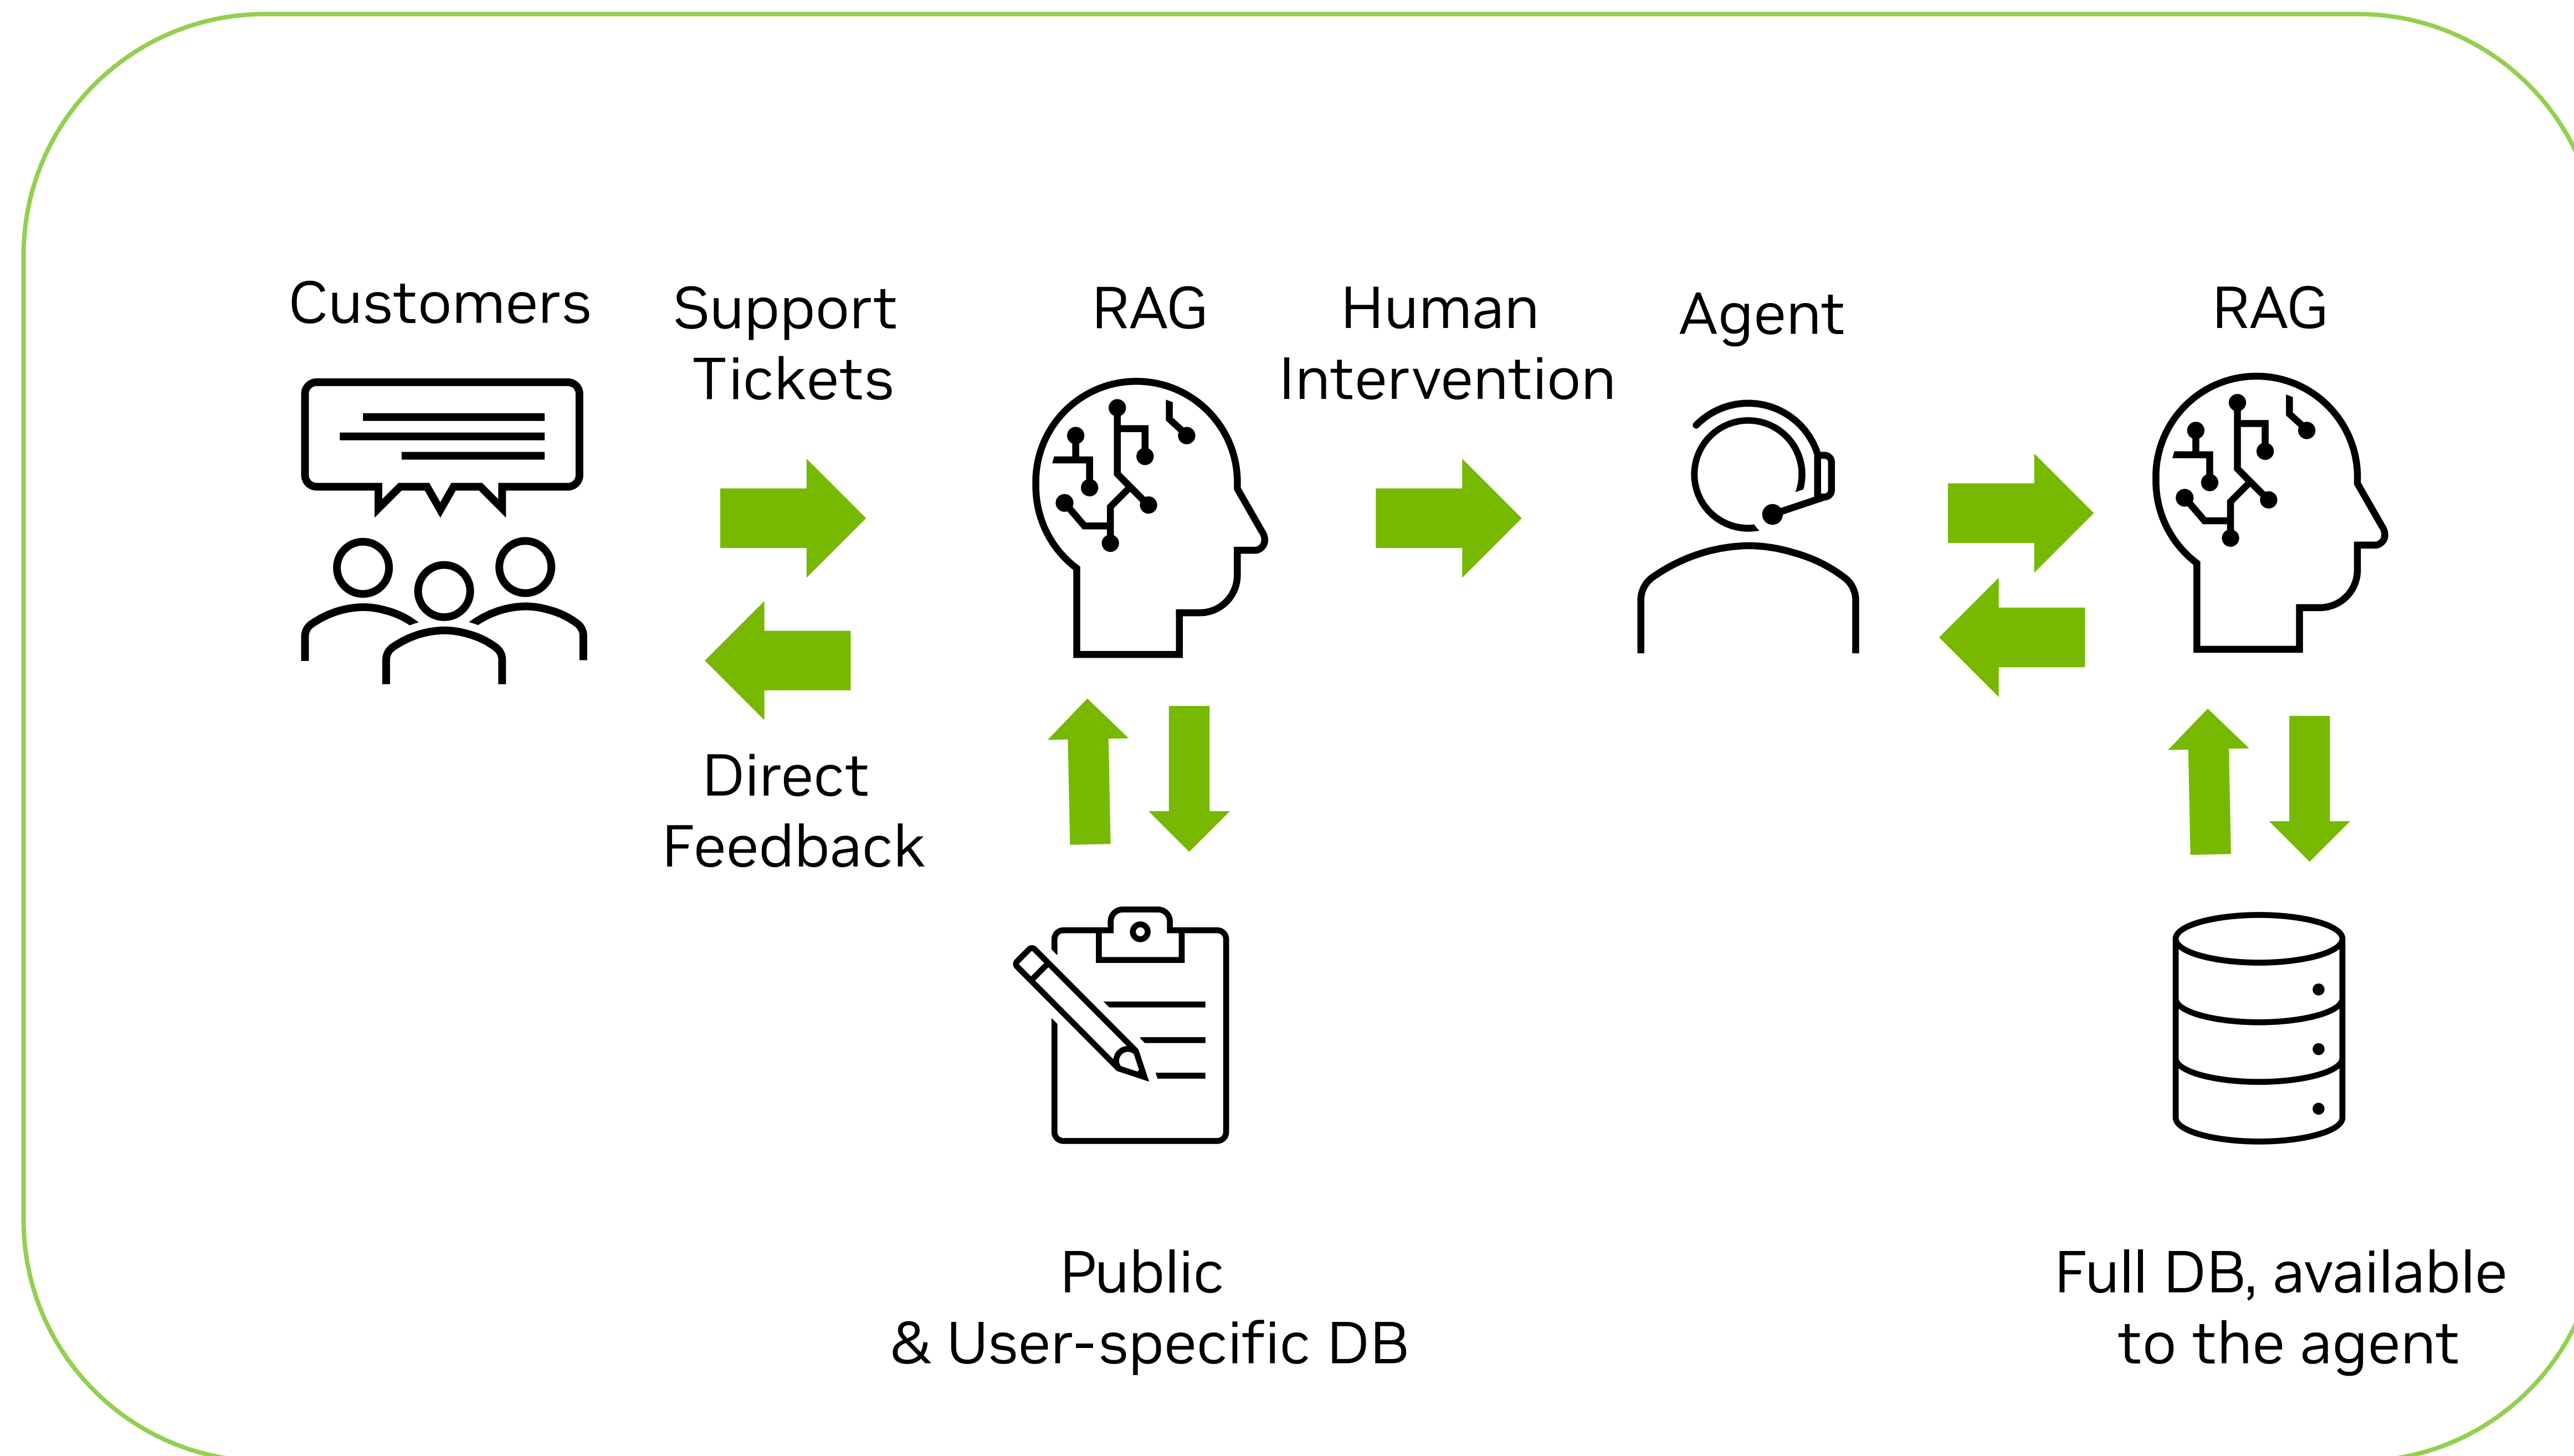
Sizing for Inference can get a bit complicated

Inference performance tools to help you are emerging

Short summary of how to think about a problem and live demo

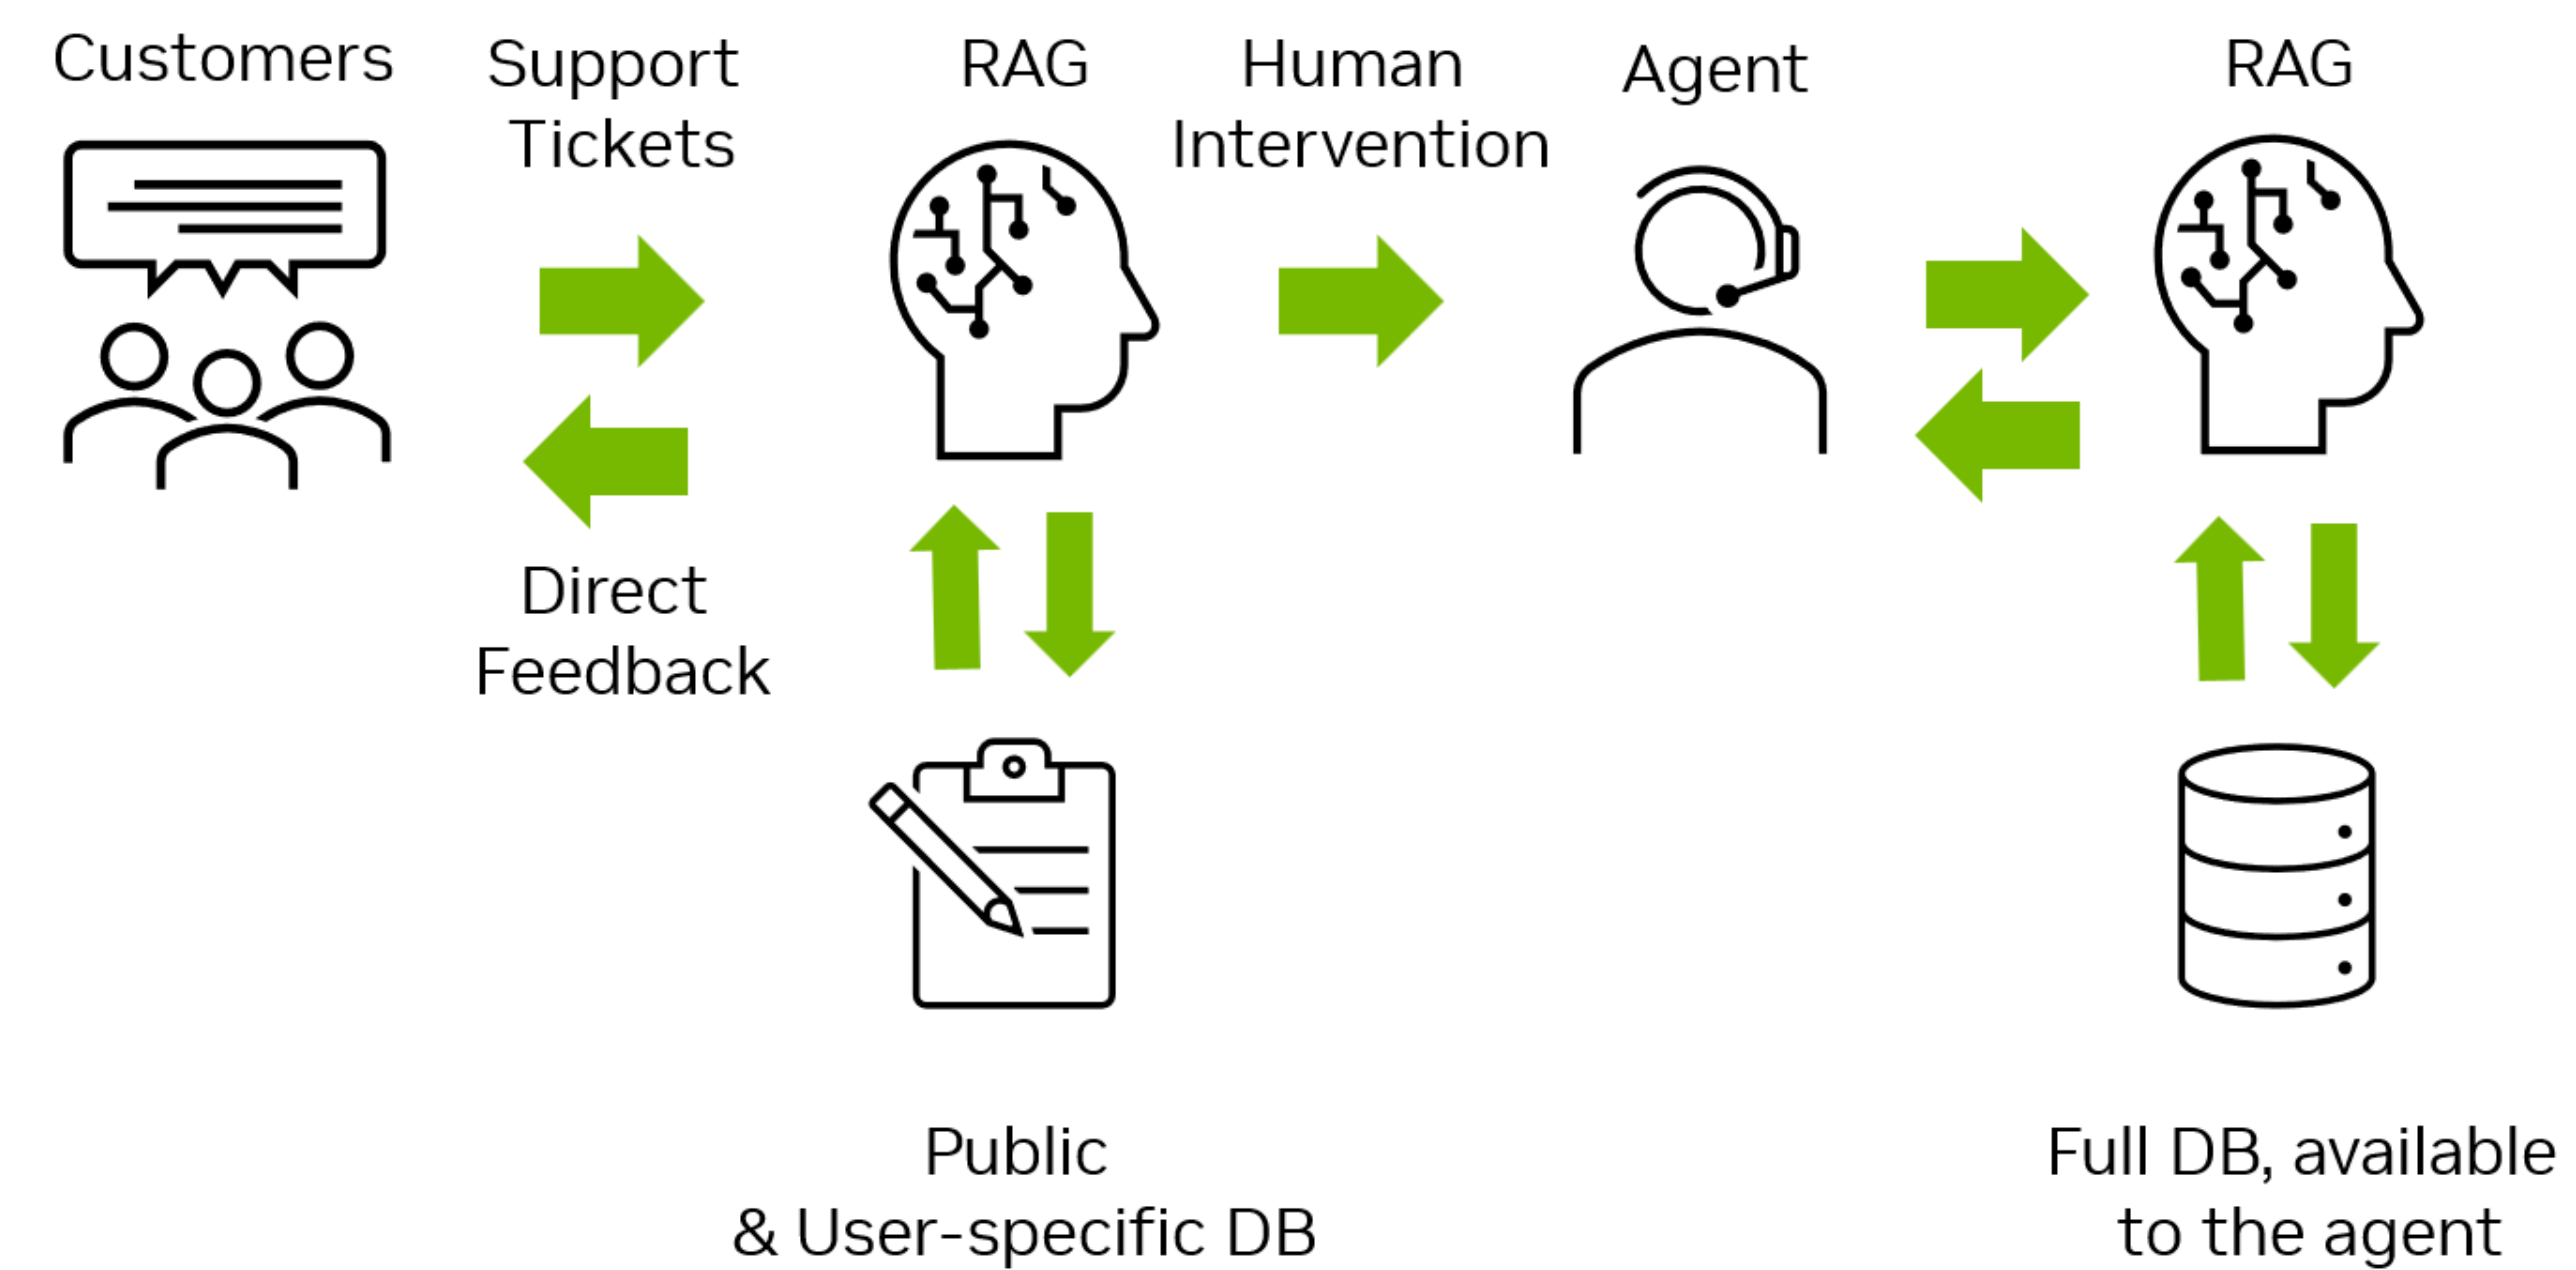
Customer Use Case Example

Challenges of sizing




Customer Use Case Example

Challenges of sizing

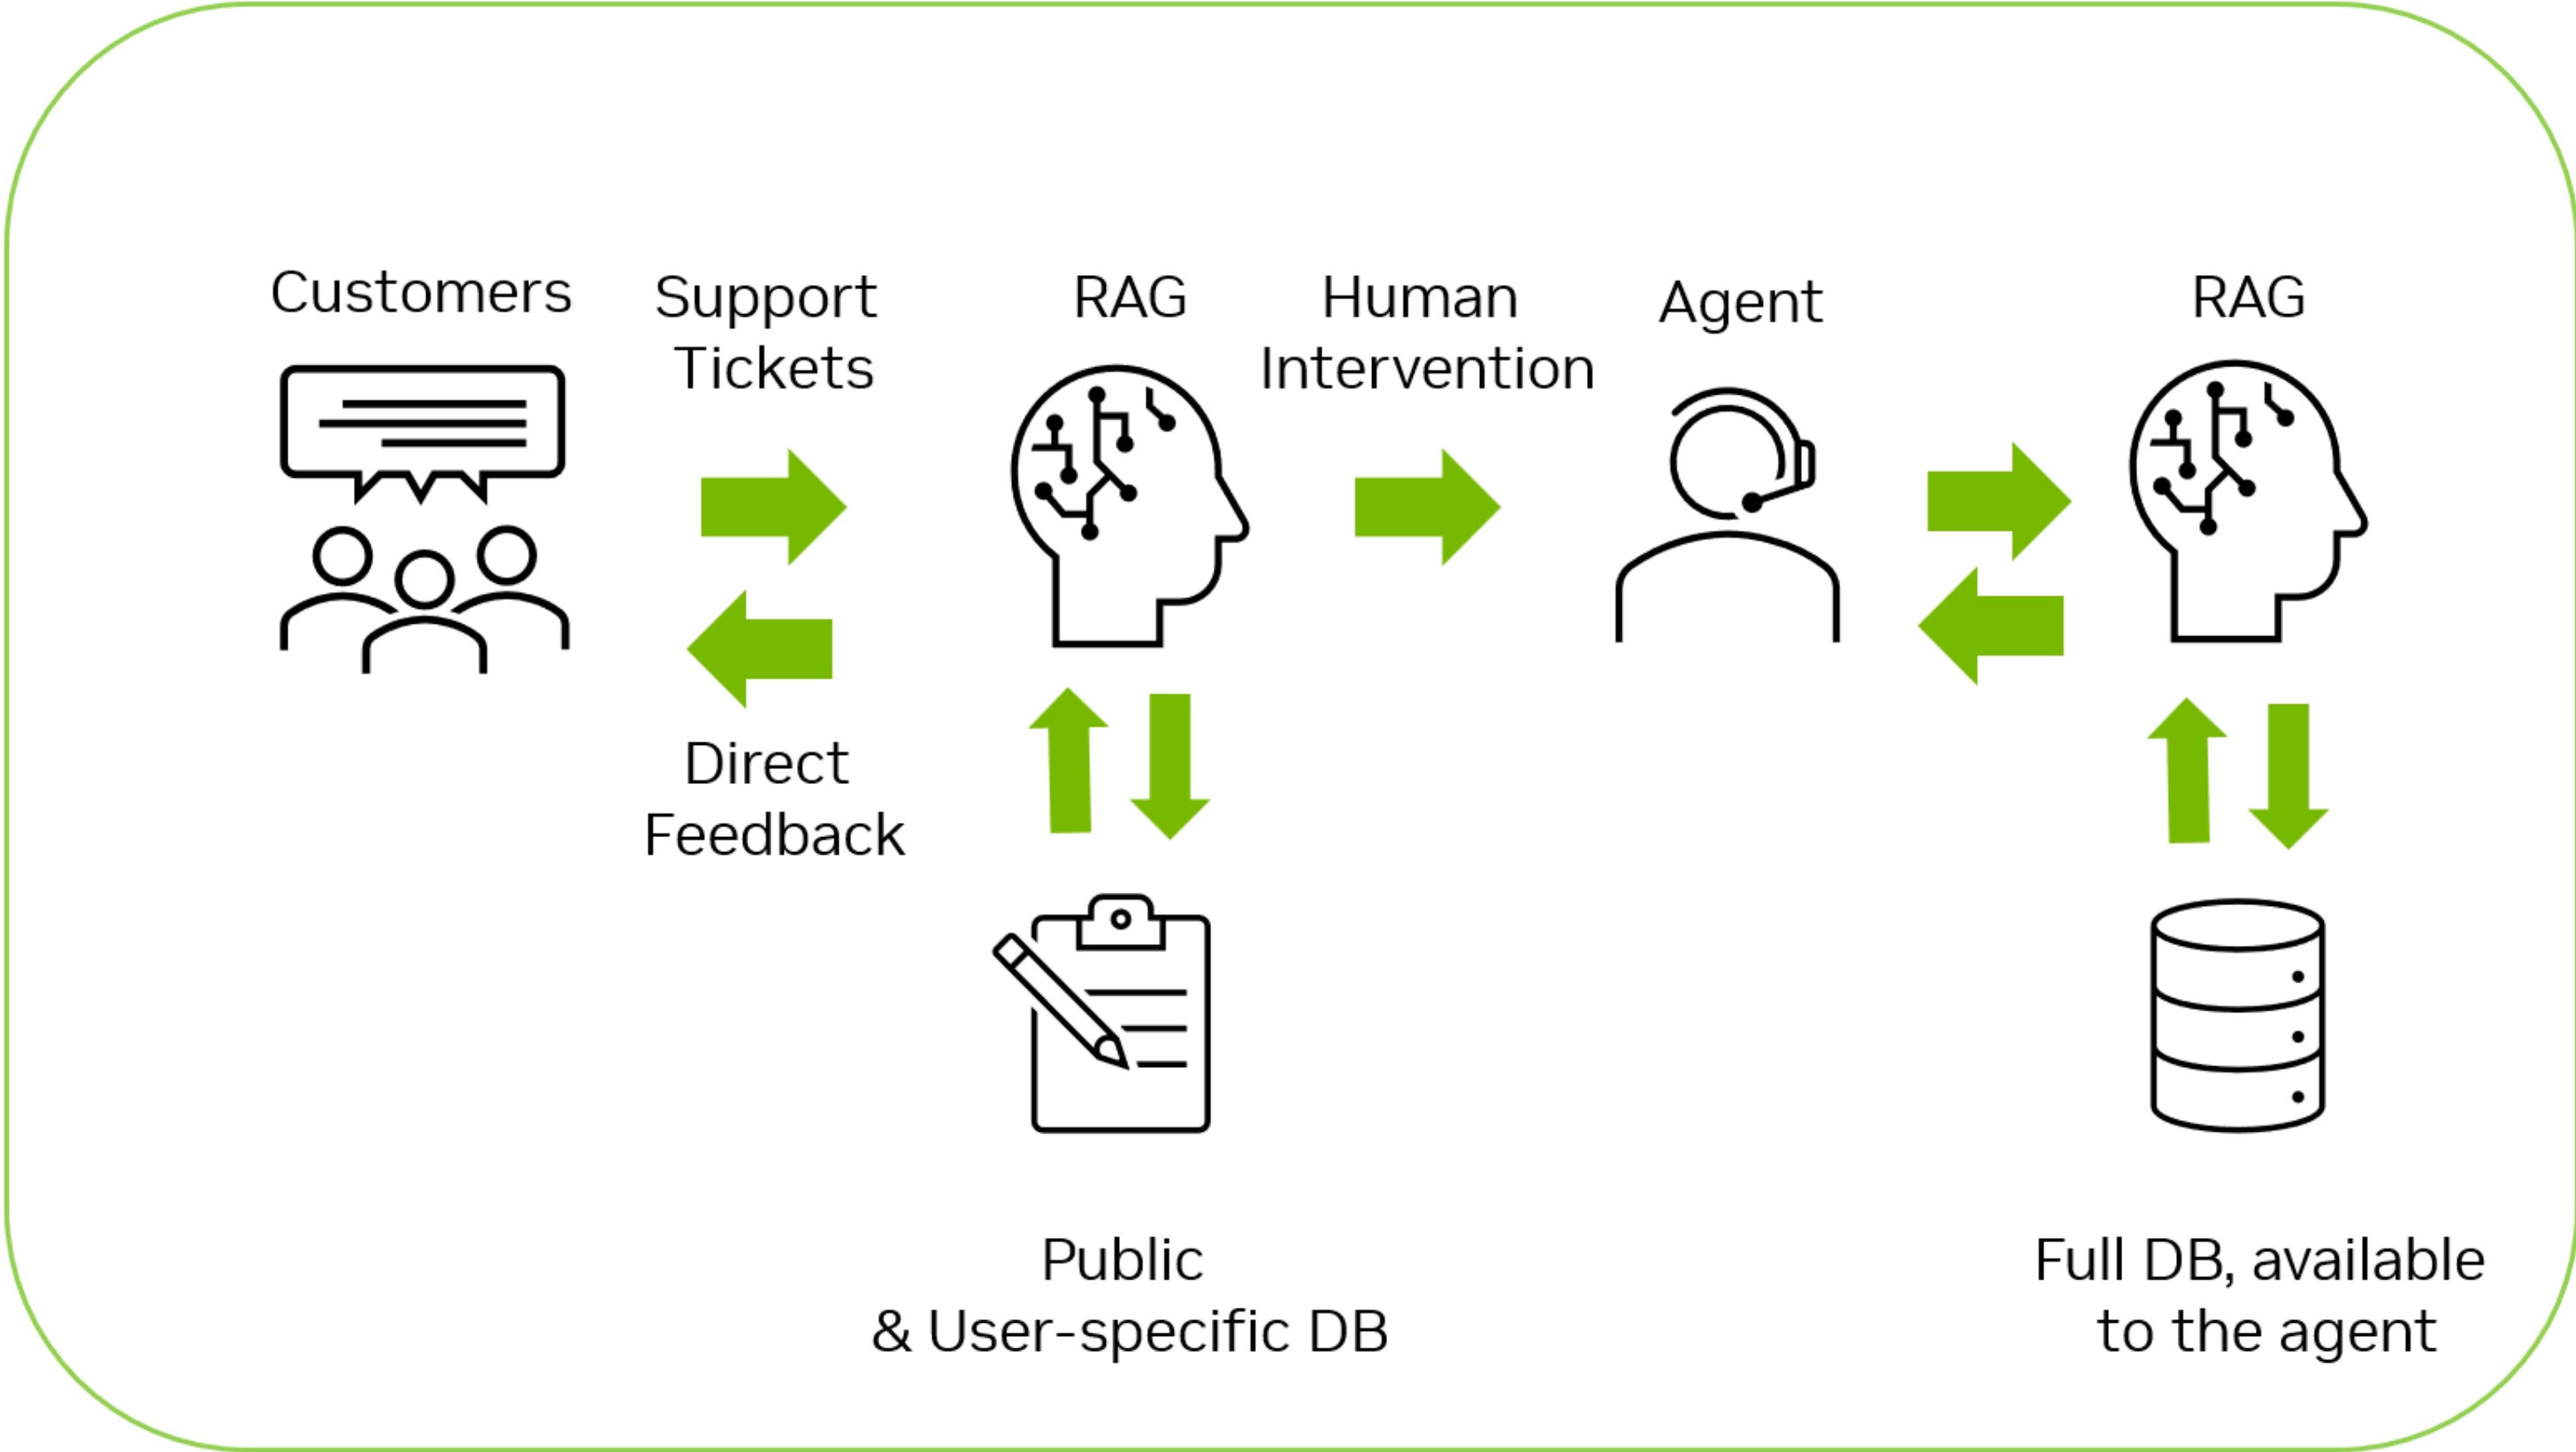


How many systems do we need to buy for this?




Customer Use Case Example


Challenges of sizing



How many systems do we need to buy for this?

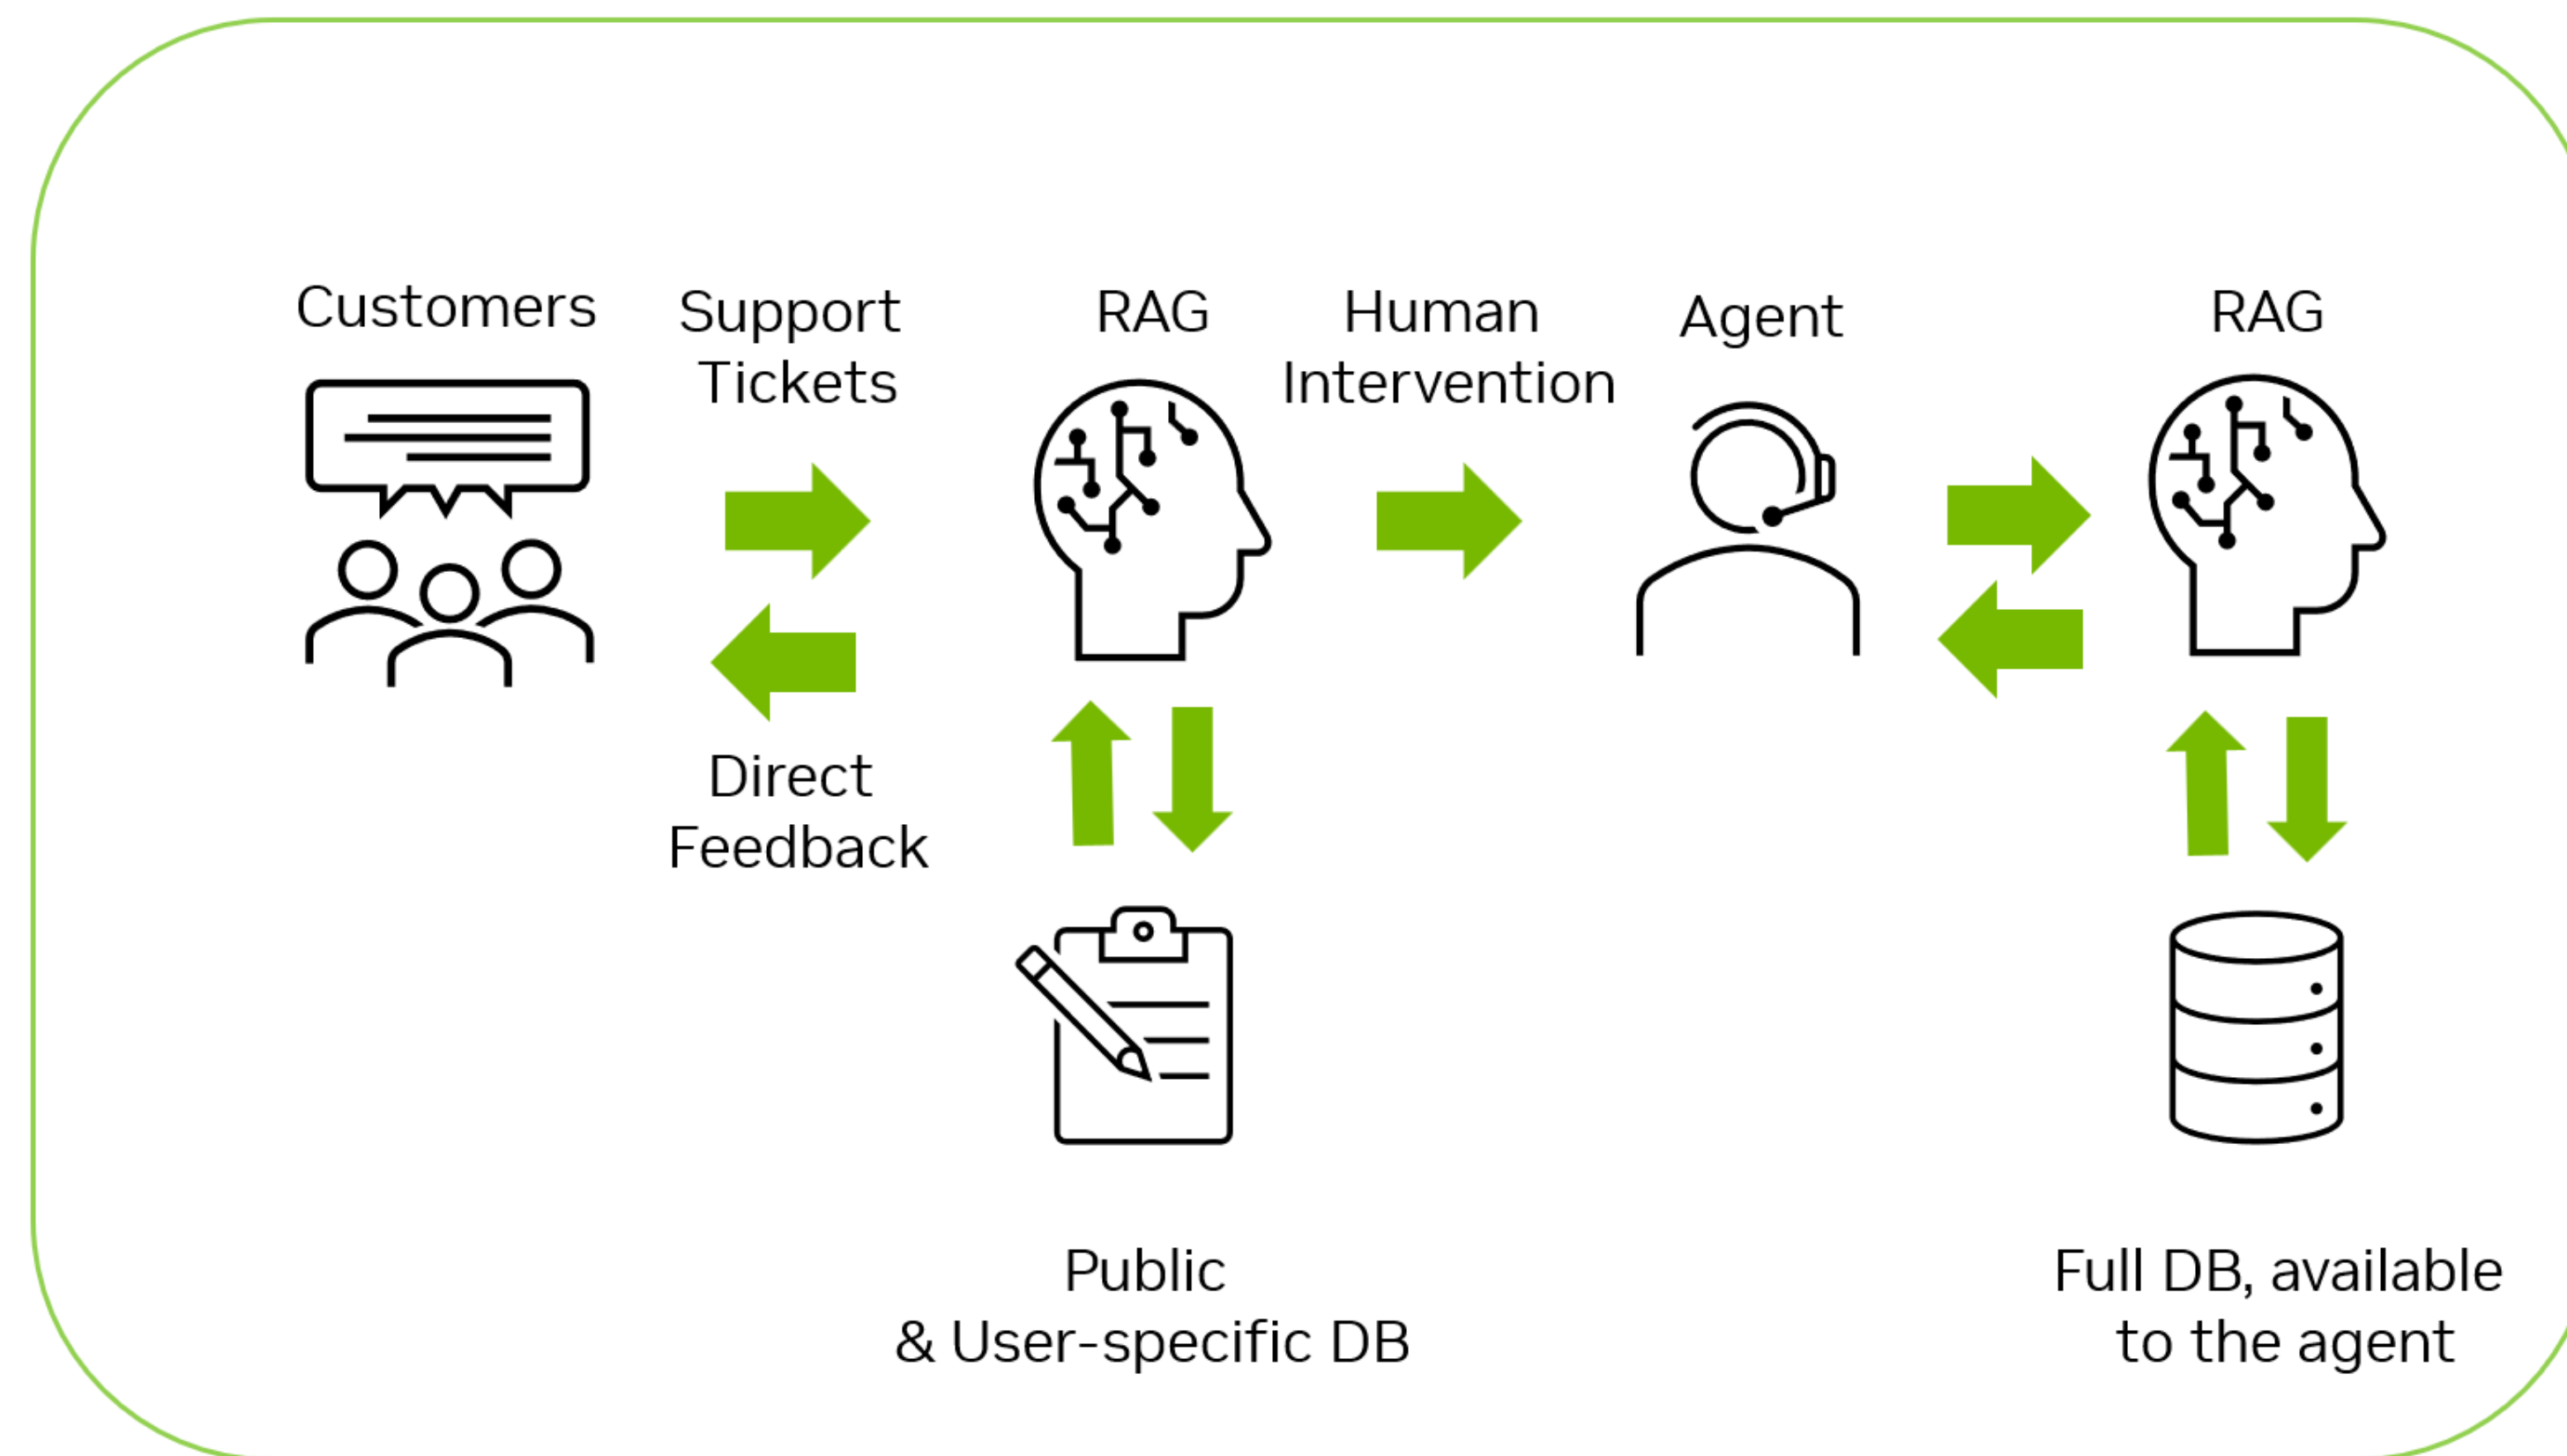


How long is a piece of string?




Customer Use Case Example



Challenges of sizing



How many systems do we need to buy for this?



How long is a piece of string?



- **The customer needs 13 DGX H100 systems**
 - DGX H100 can serve 2.4 requests per second
 - First token latency 2606 ms (prefill) is within the limit specified
 - Inter-token latency 21.4 ms/generated token
 - Generation latency of 500 tokens = $21.4 * 500 = 10\,700\text{ ms} = 10.7\text{ s}$

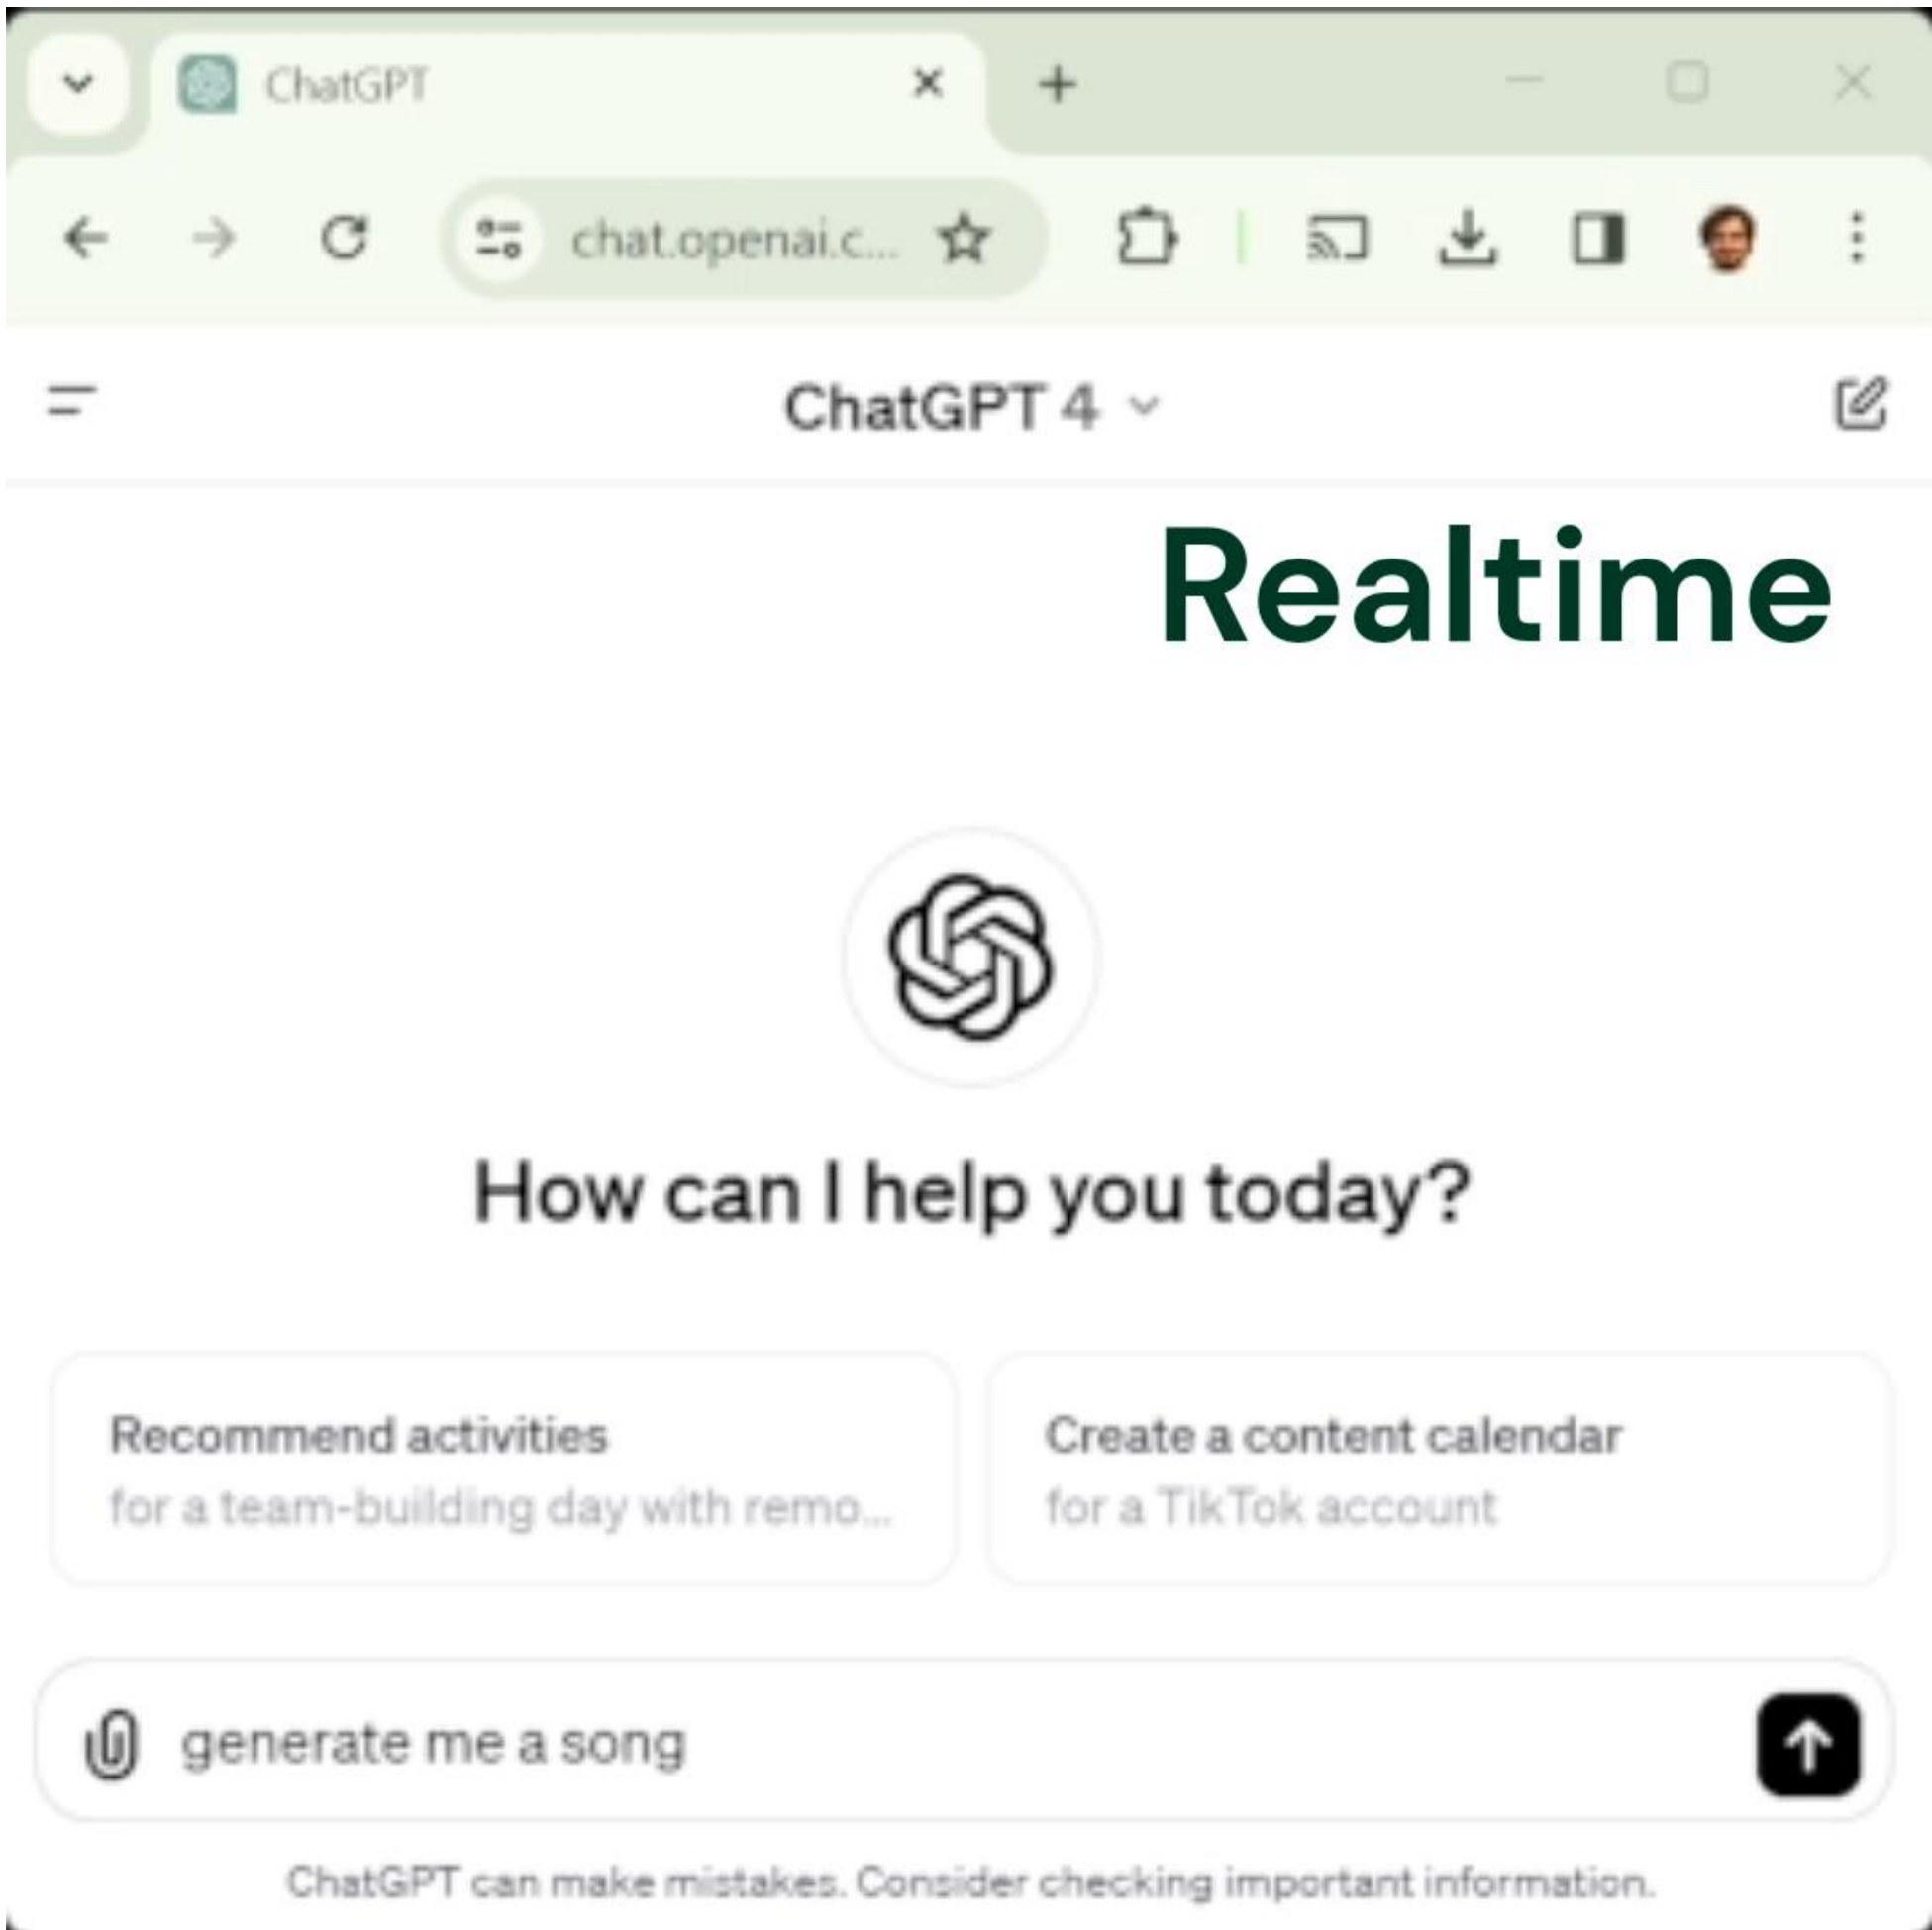


- 3500 words in, 500 words out
- NeMo 43B GPT
- First token latency limit 3s
- Max 31 requests (=prompts) per second

Two Stages of LLM Execution

Prefill vs Decoding

- **Prefill** = time to first token (~word)
 - Loading the user prompt into the system
 - From the request reception to the first token
 - Depends only on the number of input tokens
 - Populate KV-cache for all the tokens from the prompt.
 - Compute-bound for most of the reasonable prompt lengths
- **Decoding** = inter-token latency
 - Generating the response token by token, word by word
 - Inter-token latency depends on the total token number, both input and output tokens.
 - Usually memory-bound




Prefill: 1.14s, 5 input tokens, 1 output token

generate me a song

Decoding: 1.62s, 33 output tokens

Could you please provide me with some specifics


Prefill, first-response	Decoding							
User: generate me a song. AI: Could	you	please	me	with	some	specifics	for	your




The Two Things To Care About

Where and how do we execute inference?

Where?



How?



Where?

Significant impact of deployment location


On-prem

- Fixed Capacity: you need to understand the size for the maximum simultaneous load
- Pricing model: per peak capacity

Cloud

- Variable Capacity: APIs hide capacity concerns – in reality, similar limitations apply (GPU shortage)
- Pricing model: per token

- Minimal Capacity + Autoscaling for bursts



How?

Significant impact of inference strategy

Online

- Complexity: it matters to people how quickly they will get their response
- Imposing latency requirement significantly decreases available throughput. Need to balance between throughput and latency

Offline

- Simplest execution model
- Throughput, throughput, throughput: maximum GPU utilization, maximum batch size

Fun fact: Fast human reading speed is 90 ms/token (=500 words/minute at 0.75 tokens/word) (avg is 200 ms/token)

Online Streaming vs Sequential

Two facets of latency

- **Streaming**, when it is OK to give people answer one token at a time.
 - In this situation only the **TIME-TO-FIRST-TOKEN** matters (as we generate text faster than people can read).
 - One needs to develop the app streaming capabilities.
 - Simpler to satisfy real-time latency requirements
 - Can be implemented only in the last step of the pipeline
- **Sequential**, when one waits for the full response
 - Say you want to check whether the user question is not toxic BEFORE you start answering.
 - In this case **END-TO-END** latency/time to last token matters
 - Legacy apps can be simply updated with sequential mode
 - Latency requirements are too restricting for throughput

```
graph LR; A[Optional sequential request to LLM:  
"What should a database query look like for this prompt"] --> B[DB Query Run,  
Building the final prompt]; B --> C[Streaming request to the LLM, if the app allows that.  
Sequential otherwise.];
```

Optional sequential request to LLM:
"What should a database query look like for this prompt"


DB Query Run,
Building the final prompt

Streaming request to the LLM, if the app allows that.
Sequential otherwise.


In This Presentation

We focus on the most complicated part of the problem

Where?



How?






Sizing One Use Case

Questions for a Sizing Use Case


1. ☒ What model are you planning to use?
2. ☒ What is the average number of tokens in the prompt to your LLM (Length of input)?
 - For English one token is approximately 0.75 of a word.
 - Make sure to include system prompt.
3. ☒ What is the average number of tokens in your LLM output?
4. ☒ How many requests per second should your system process at its peak?
5. ☒ What is your latency limit? First-token? Last-token?



Which Model?

The most popular requests

- Typically, we get asked about LLaMa 2 family of the models:
 - Free for research and commercial use
 - Supported by NVIDIA SW stack, including NeMo
- The bigger the model, the more resources it needs for inference
 - The bigger the model the better the accuracy
 - Very roughly the resource amount scales with the model size
- If considering LLaMa 7B and 13B parameters, see also NVIDIA Nemotron-3 8B Family of models: [blog](#)



Input Length

There's a maximum budget of tokens to pass into the model

- Most of the models support up to 4096 tokens.
 - Context window = input tokens + output tokens
 - Llama2 supports 4096 context window
 - New models support even larger context windows
- Everything counts so be careful:
 - **System prompt** (a.k.a custom instructions): instructions you give to the model for every “dialogue”. Make sure to include them into the input token count as shown in example on the right.
 - **Retrieved documents** (a.k.a Retrieval Augmented Generation, RAG). For RAG pipelines key paragraphs from the internal document storage are added to the prompt, before the user requests. Typically RAG systems target to use full available context length.
 - For 4K context typical 3500 input tokens, 500 output tokens
 - [What is RAG — NVIDIA blog](#)
 - **Chat history**: previous exchange of messages in the conversation

Custom instructions ⓘ

What would you like ChatGPT to know about you to provide better responses?

I work for NVIDIA as a Solutions Architect.

This sentence costs +9 input tokens

43/1500

How would you like ChatGPT to respond?

Respond concisely, unless asked to expand your thoughts.

This sentence costs +12 input tokens

56/1500


Enable for new chats ☒

Cancel

Save

Peak Requests Per Second

- Poisson distribution approximation
 - One knows the average, but would like to know the peak
- Find 95th percentile: [ChatGPT dialogue](#)



```
from scipy.stats import poisson
```

```
# Parameters
```

```
lambda_ = 64 # average number of requests per second
```

```
percentile = 0.95 # 95th percentile
```

```
# Calculate the 95th percentile value
```

```
k_95th_percentile = poisson.ppf(percentile, lambda_)
```

```
print(k_95th_percentile) # 77, 20% difference
```


```
print(poisson.ppf(0.95, 7)) # 12, 71% difference
```


LLM Inference Requires Multiple GPUs

Tensor Parallelism (TP) – so how to split your neural network across several GPUs

- Tensor Parallelism (TP) can be used for LLM Inference. One model gets split across several GPUs. Heavily relies on data exchange between GPUs.
 - Lower latency, but lower throughput
 - $TP \geq 2$ required for bigger models like LLaMa-70B
- If $TP > 2$ we strongly recommend NVLink-enabled servers for inference, such as HGX and DGX systems
- We normalize all the results for servers with 8 GPUs (even for L40s)
 - $(\# \text{ of instances}) * TP = 8$
 - 8 instances with TP1, 2 instances with TP4

TP8	Instance 1							
TP4	In. 1				In. 2			
TP2	In. 1		In. 2		In. 3		In. 4	
TP1	In. 1	In. 2	In. 3	In. 4	In. 5	In. 6	In. 7	In. 8



Time = \$



Tools Available


Publicly Available Performance Benchmarking

- <https://github.com/NVIDIA/TensorRT-LLM/tree/main/benchmarks/cpp> — TensorRT-LLM C++
 - TensorRT-LLM provides users with an easy-to-use Python API to define Large Language Models (LLMs) and build TensorRT engines that contain state-of-the-art optimizations to perform inference efficiently on NVIDIA GPUs. TensorRT-LLM also contains components to create Python and C++ runtimes that execute those TensorRT engines.
 - Some results: <https://github.com/NVIDIA/TensorRT-LLM/blob/main/docs/source/performance.md>
- https://github.com/triton-inference-server/client/blob/main/src/c%2B%2B/perf_analyzer/docs/llm.md
 - Triton Performance Analyzer is CLI tool which can help you optimize the inference performance of models running on Triton Inference Server by measuring changes in performance as you experiment with different optimization strategies.

Inference Performance Exploration Tools

- Simpler, less precise, benchmarks-based <https://nemo-inference-sizing.nvidia.com/> (to be published)
- Very precise, more complex, benchmarks + simulation: [Nemo Inference Microservice](#) (available Early Access)


<https://nemo-inference-sizing.nvidia.com/>



Recommended Configurations Within Limits

Metric	Throughput-optimal ▼	Latency-optimal ▲	Clicked data
Latency for first token (ms)	258.3	42.2	[Click on a point]
Latency per generated token (ms)	14.8	8.7	[Click on a point]
Prompts per second per 8 GPUs	65.6	14.3	[Click on a point]
Tensor Parallelism	1	1	[Click on a point]
Batch size	64	8	[Click on a point]

Nemo Inference Microservice


NVIDIA

[CATALOG](#)
[CONSOLE](#)
[PRIVATE REGISTRY](#)

Private Registry
Containers
NeMo Microservice Inference

NeMo Microservice Inference

Get Container


Overview

Tags

Layers

Security Scanning

Related Collections



NVIDIA
NEMO

Description

NeMo Microservice Inference GPU accelerated LLM capabilities through OpenAI compatible APIs

Publisher

NVIDIA

NVIDIA

Latest Tag

23.12

Modified

January 8, 2024

Compressed Size

13.49 GB

Multinode Support

No

Multi-Arch Support

No

23.12 (Latest) Security Scan Results

No results available.

Overview

The NeMo Microservice Inference (NMI) brings the power of state-of-the-art large language models (LLM) to your applications, providing unmatched natural language processing and understanding capabilities. Whether you're developing chatbots, content analyzers—or any application that needs to understand and generate human language—NMI has you covered. Built on the NVIDIA software platform incorporating CUDA, TRT, TRT-LLM, and Triton, NMI brings state of the art GPU accelerated Large Language model serving.

High Performance Features

NMI enables optimized scheduling technique called in-flight batching. This technique takes advantage of the fact that the overall text generation process for an LLM can be broken down into multiple execution iterations on the model. With in-flight batching, rather than waiting for the whole batch to finish before moving on to the next set of requests, the NMI runtime immediately evicts finished sequences from the batch. The runtime then begins executing new requests while other requests are still in flight.

Scalable Deployment: Whether you're catering to a few users or millions, the microservice scales seamlessly to meet your demands.

Advanced Language Models: Built on cutting-edge LLM architectures, NMI provides optimized and pre-generated engines for a variety of popular models. The tooling to create GPU optimized models is also included.

Flexible Integration: Easily incorporate the microservice into existing workflows and applications, thanks to multiple API endpoints.

Secure Processing: Your data's privacy is paramount. NMI ensures that all inferences are processed securely, with rigorous data protection measures in place.

Applications

Chatbots & Virtual Assistants: Empower your bots with human-like language understanding and responsiveness.

Content Generation & Summarization: Generate high-quality content or distill lengthy articles into concise summaries with ease.

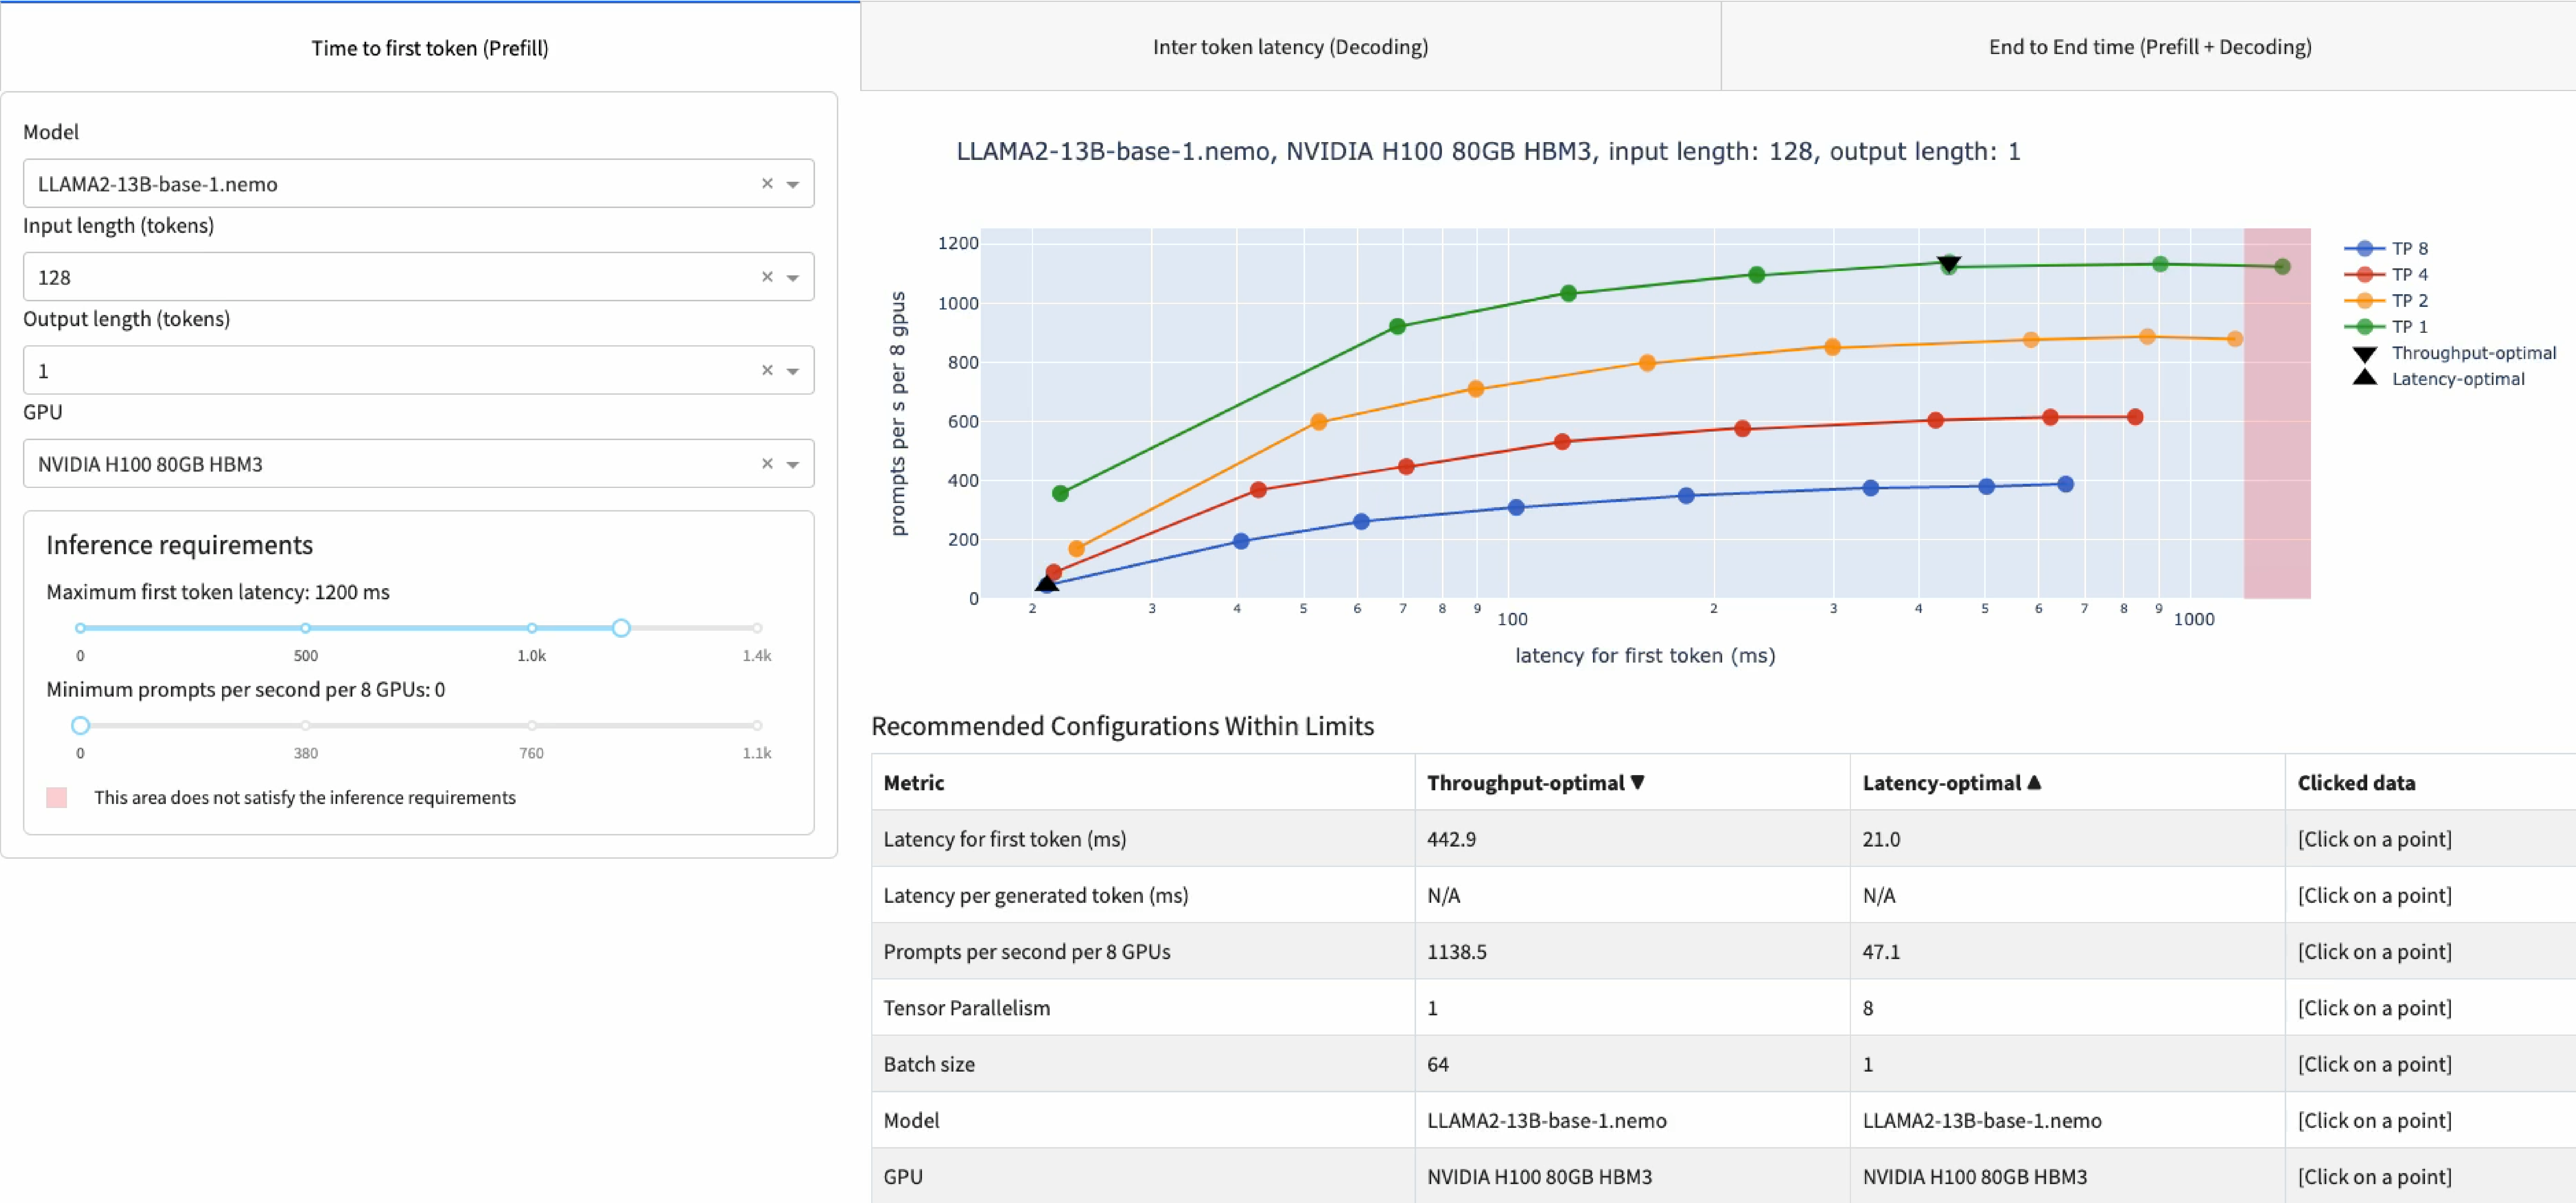
Sentiment Analysis: Understand user sentiments in real-time, driving better business decisions.

Language Translation: Break language barriers with efficient and accurate translation services.

And many more... The potential applications of NMI are vast, spanning across various industries and use-cases.

*name on the page will be updated

Demo of NeMo inference sizing




Example 1

Smaller model – for auxiliary task

LLAMA2-7B-base-1.nemo, NVIDIA H100 80GB HBM3, input length: 128, output length: 512

- We are looking for a sizable use case of **Llama-7B**. 128 in, 512 out.
- For input 128, output 512 we have **65.6** peak prompts per second per **one** DGX H100
- That's 53.2 requests per second on average
- That's 1.5M requests per working day (8 hours)
- 3 requests per person → **500k daily active users**
- 192M input, 768M output tokens per day



Recommended Configurations Within Limits


Metric	Throughput-optimal ▼	Latency-optimal ▲	Clicked data
Latency for first token (ms)	258.3	42.2	[Click on a point]
Latency per generated token (ms)	14.8	8.7	[Click on a point]
Prompts per second per 8 GPUs	65.6	14.3	[Click on a point]
Tensor Parallelism	1	1	[Click on a point]
Batch size	64	8	[Click on a point]

Example 2

Larger model (ChatGPT like)

LLAMA2-70B-base-1.nemo, NVIDIA H100 80GB HBM3, input length: 128, output length: 512

- We are looking for a sizable use case of **Llama-70B**. 128 in, 512 out.
- For input 128, output 512 we have **11.2** peak prompts per second per **one** DGX H100
- That's 6.92 requests per second on average
- That's 200k requests per working day (8 hours)
- 3 requests per person → **66k daily active users**
- 25.6M input, 102M output tokens per day
- \$38.4 + \$204 GPT 3.5 turbo per day (fair comparison) = \$7.2K/month on OpenAI



Recommended Configurations Within Limits

Metric	Throughput-optimal ▼	Latency-optimal ▲	Clicked data
Latency for first token (ms)	1516.7	37.3	[Click on a point]
Latency per generated token (ms)	41.6	17.2	[Click on a point]
Prompts per second per 8 GPUs	11.2	0.1	[Click on a point]
Tensor Parallelism	4	8	[Click on a point]
Batch size	128	1	[Click on a point]

Example 3

A100 FP16 vs H100 FP8

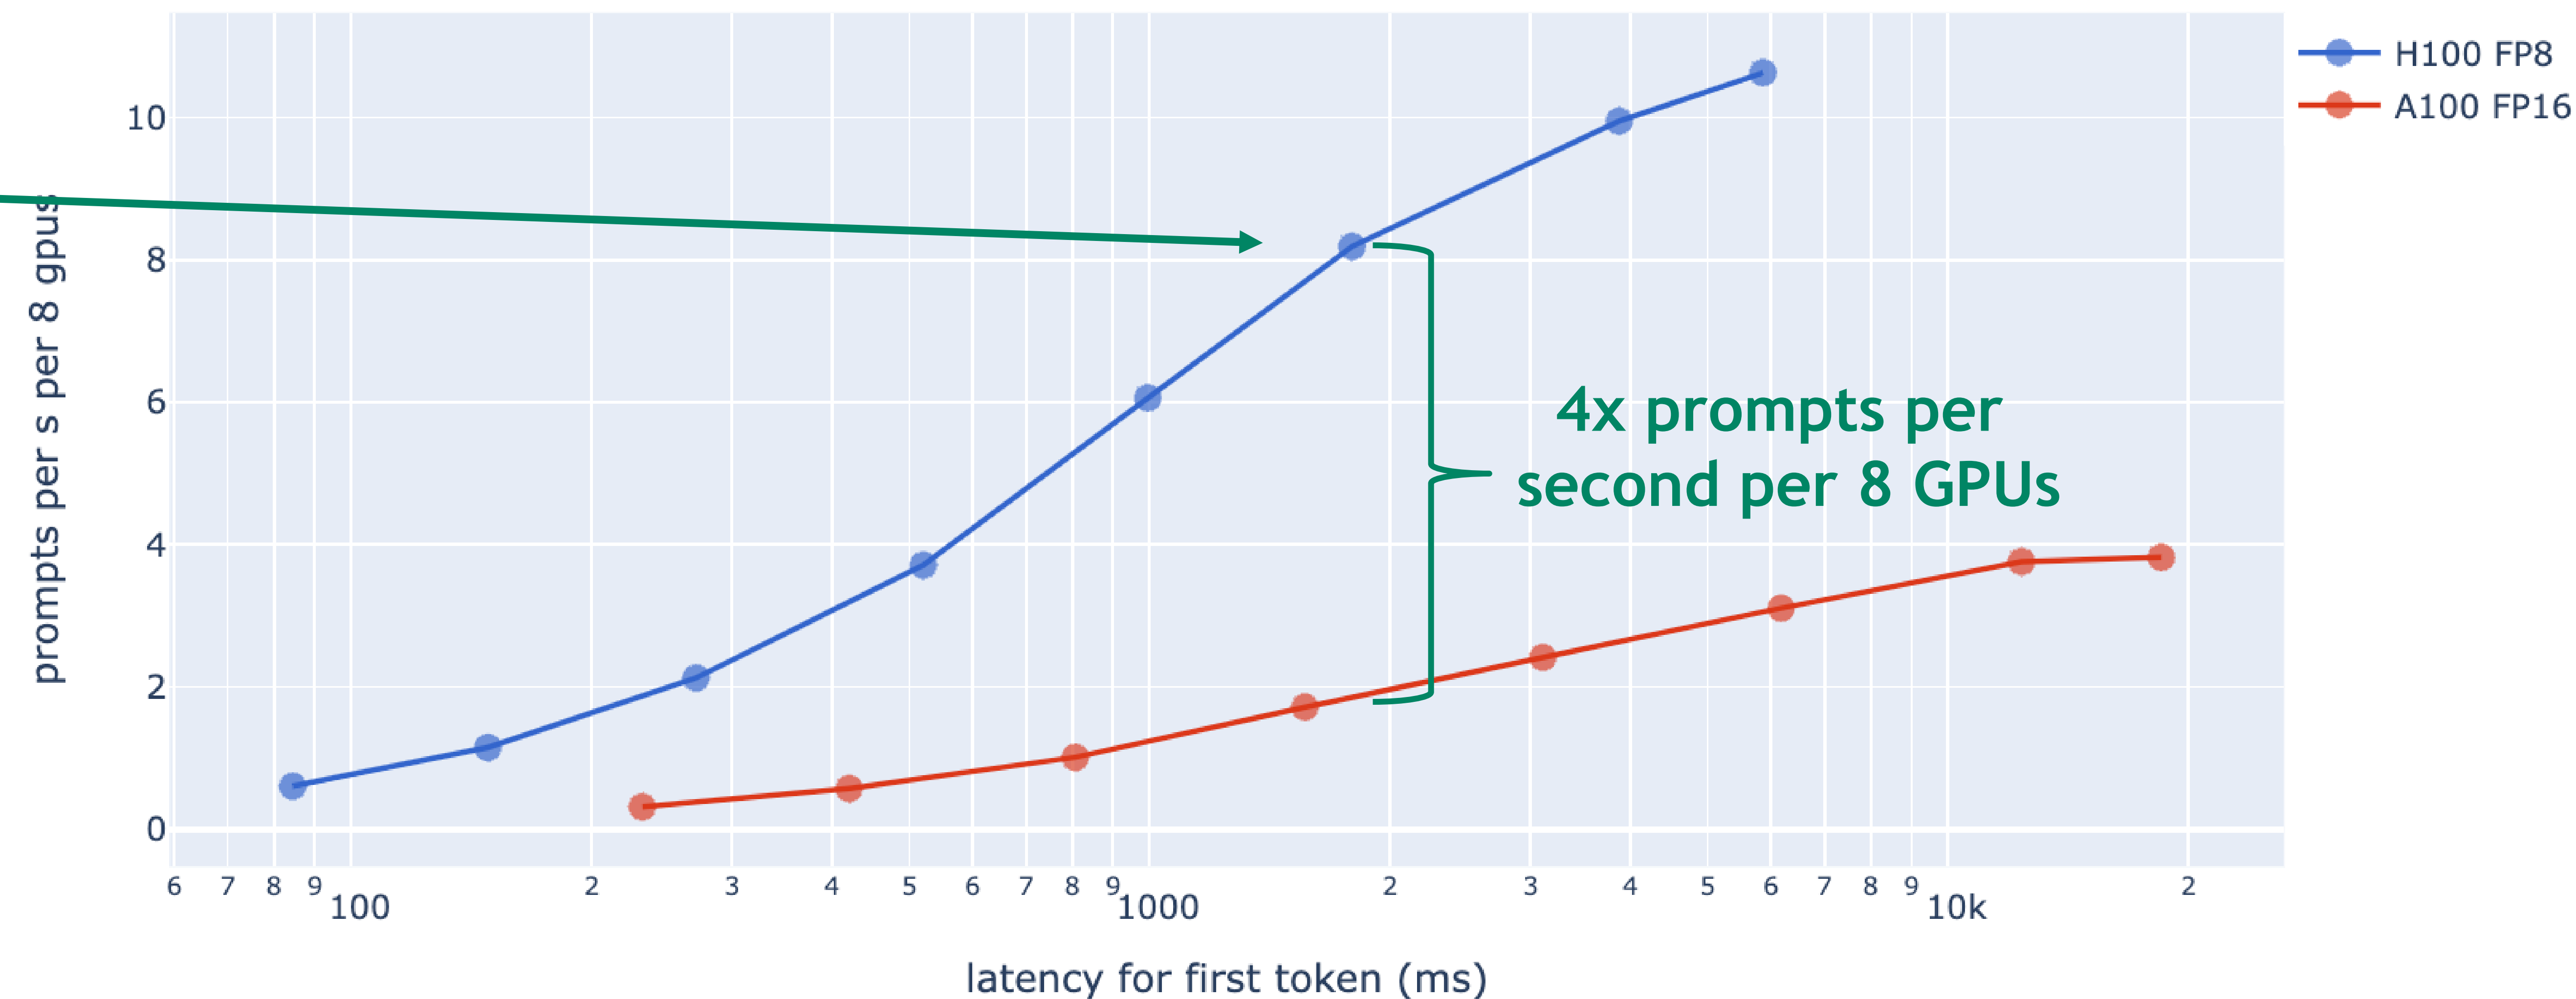
- H100 features a Transformer Engine with FP8 precision support

LLAMA2-70B-base-1.nemo, input length: 2048, output length: 128, TP: 8

- For input 2048, output 128 and TP8, H100 with FP8 delivers **x4 prompts per second** compared to A100 with FP16, for the same latency


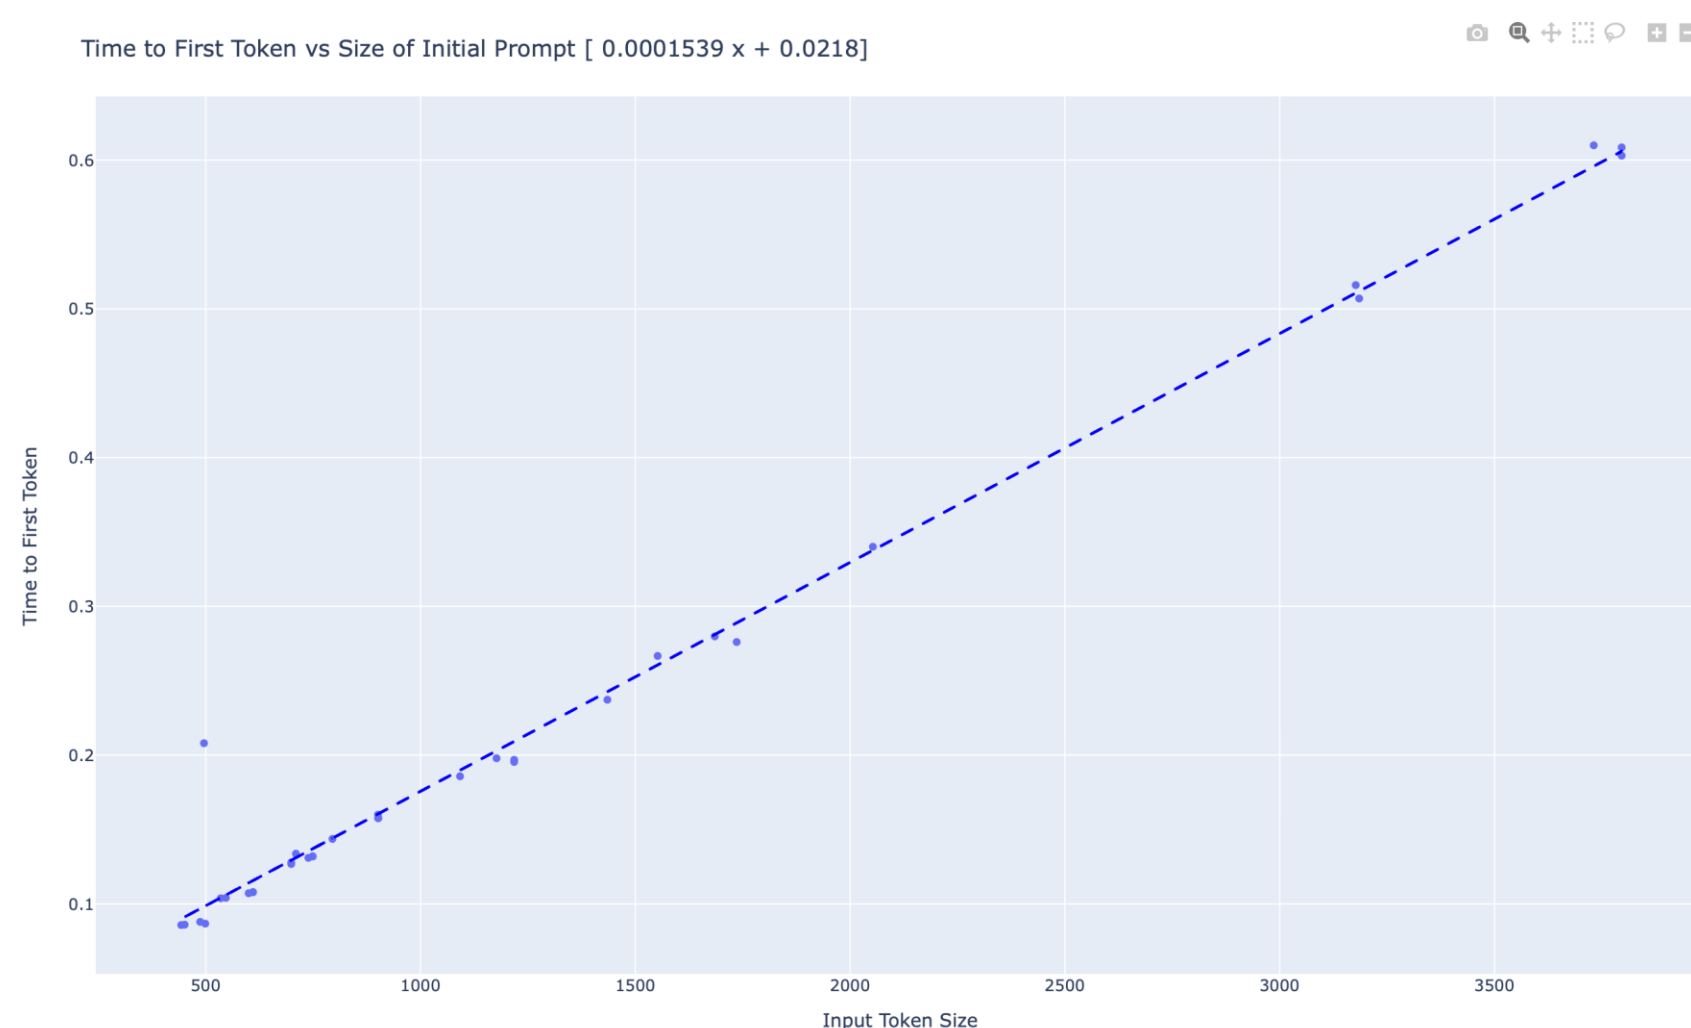
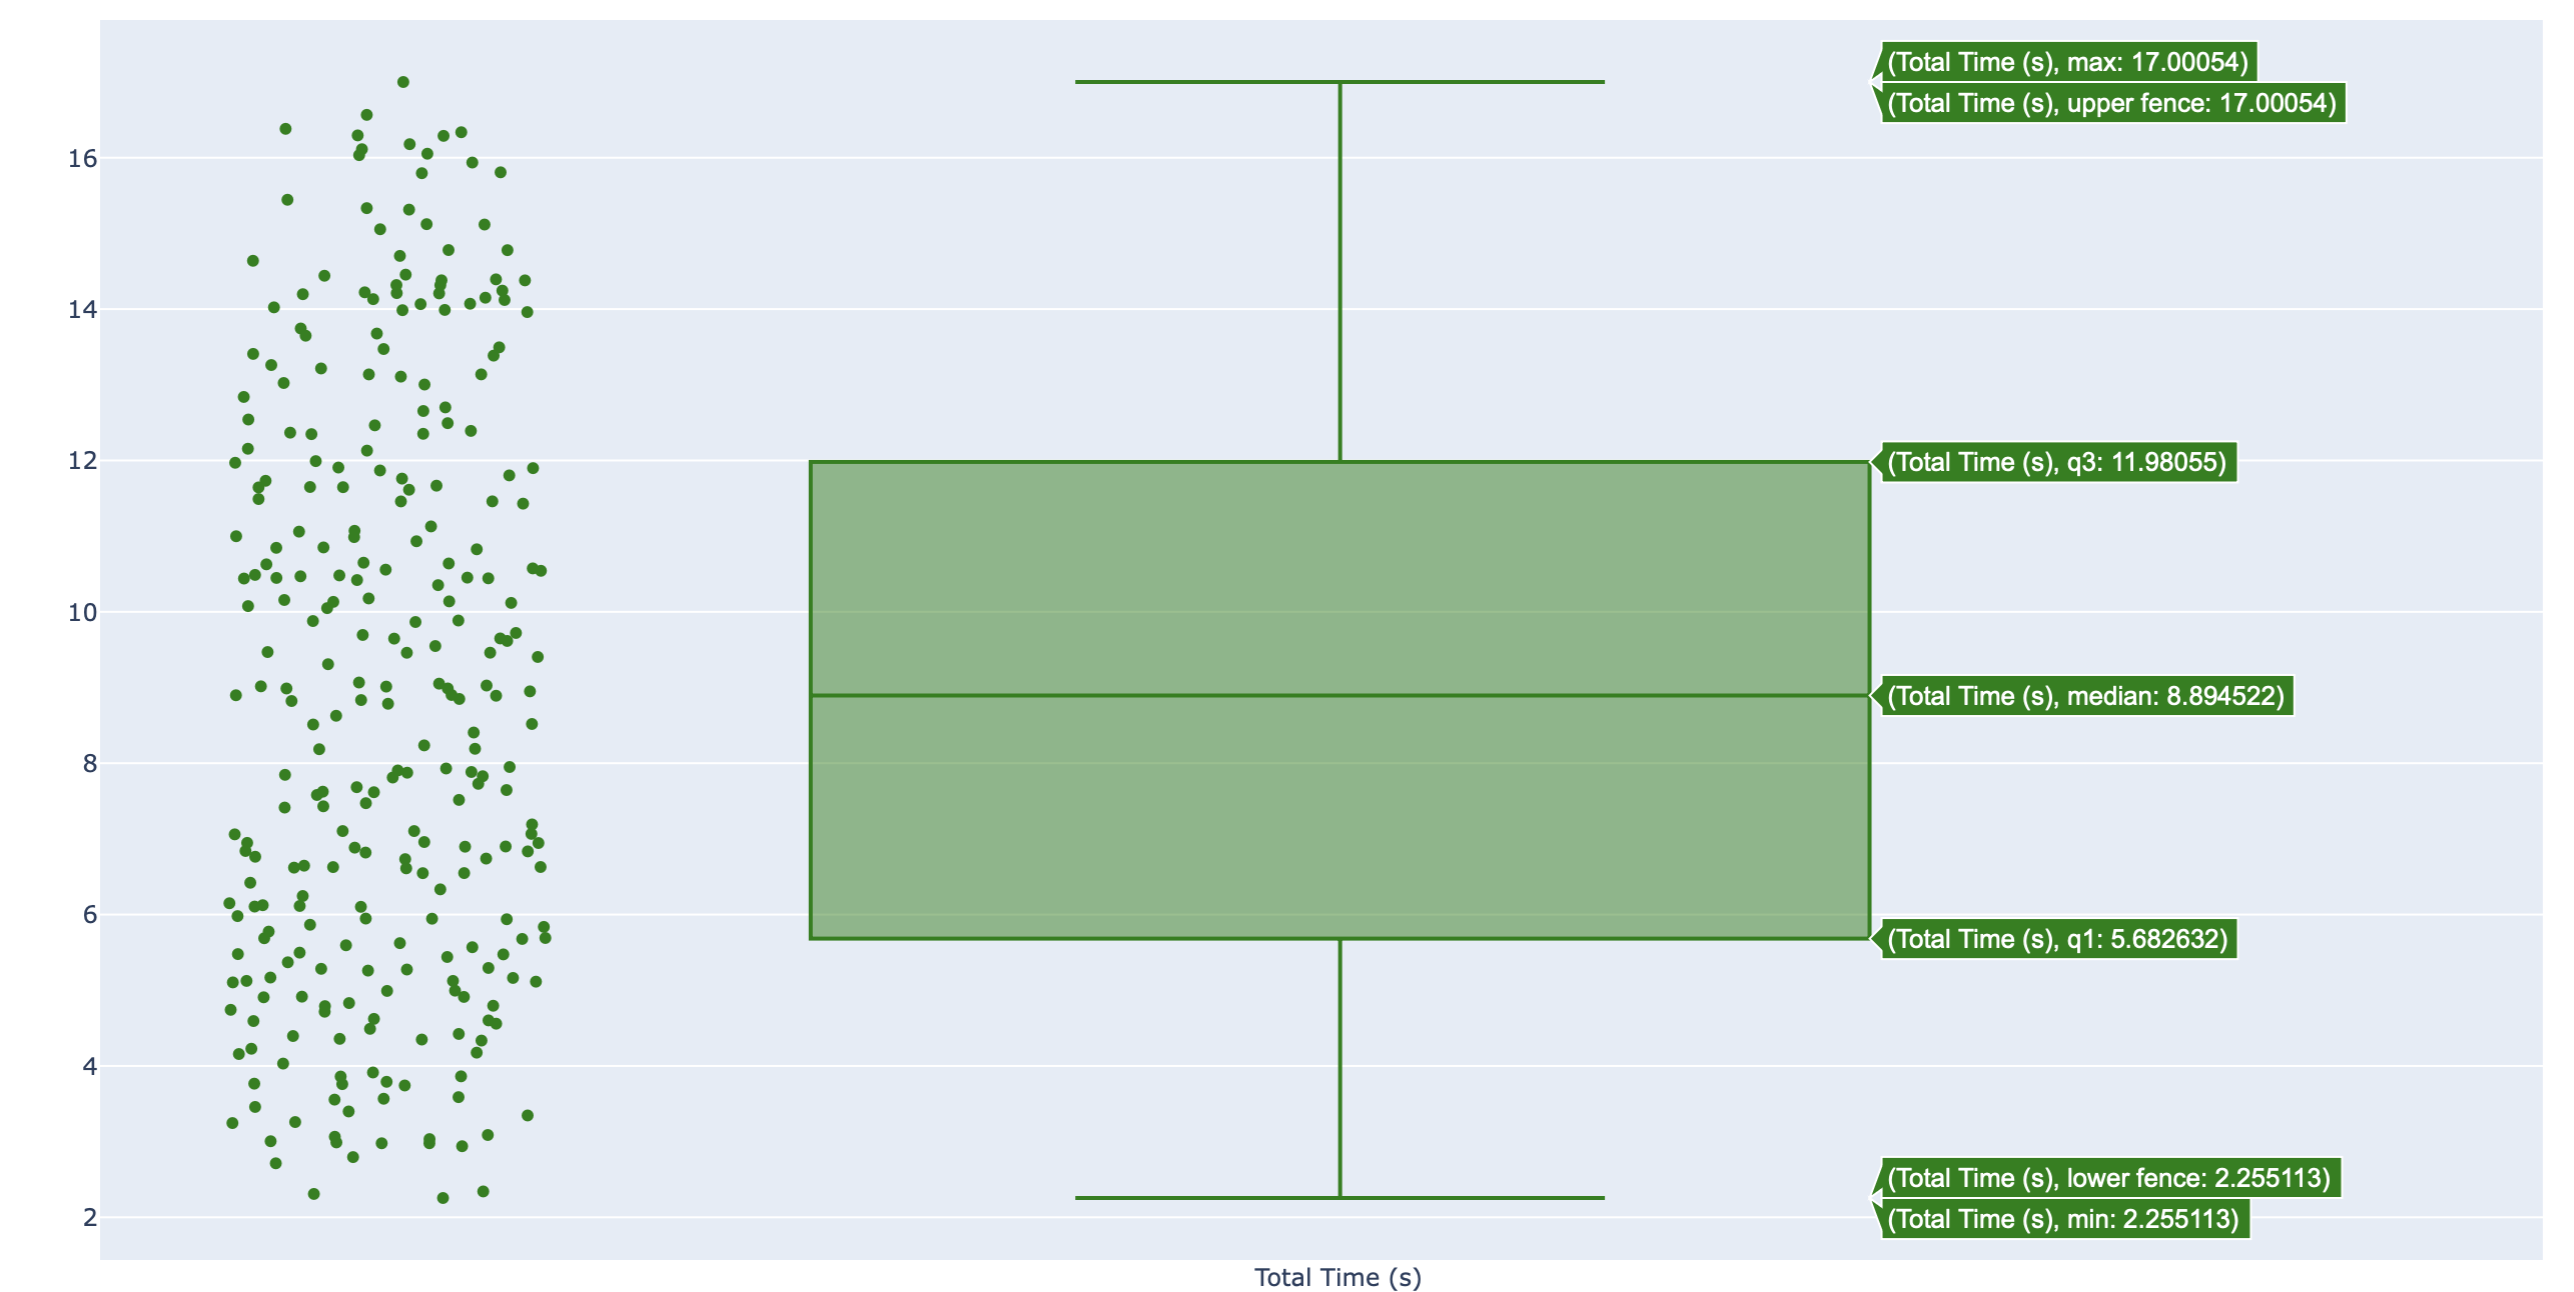

- Part of that increase is due to the FP8 format and part due to the better performance of H100 vs A100

- Other efficient techniques like pruning, distillation or sparsification can increase performance



NeMo Inference Microservice Performance Tools

- Measure, Plan, Deploy
 - Load Generators
 - Concurrent Loadgen = MLPerf MultiStream Scenario or LLM Perf tool
 - Poisson Loadgen = MLPerf Server Scenario
 - Trace Analysis
 - Collect runtime statistics from Load Generators
 - Generate Visualizations of the collected statistics
 - Performance Models
 - Use collected statistic to build a predictive performance model
 - Digital Twin of the Production Deployment
 - Evaluate Scenario in Simulation
 - Tweak arrival rates and/or input/output and evaluate the performance of the virtual deployment.
 - Simulation is accurate to within 10-15% for small models and to within 3-5% for large models




```
In [25]: print(f"actual walltime: {full_loadgen_report.walltime:.2f}; simul  
actual walltime: 163.49; simulated walltime: 170.15; ratio: 0.96
```

Rules of Thumb for Sizing

- We estimate the sizing based on NVIDIA SW stack: NeMo, TensorRT-LLM (=TRT-LLM) and Triton Inference Server
- For models greater than 13B, that need more than 1 GPU, prefer NVLink-enabled systems.
- In the streaming mode, when the words are returned one by one, first-token latency is determined by the input length.
- The cost and the latency are usually dominated by the number of output tokens
 - Example below: H100 SXM, Llama 70B, BS 8, TP 4, FP 16.
Input of 3500 tokens takes the same amount of time as generating 99 tokens
(2.6 seconds each stage, 26.8 ms/generated token)
 - However, generating is almost always faster than human reading speed
 - Thus, input tokens are much cheaper
- Introducing latency limit can significantly decrease available throughput
- Larger models require more memory and have higher latency, scaling approximately with the model size.
- New apps should be developed in streaming mode . To introduce LLMs into the older apps, one may use sequential mode.
- Locality of compute is not too important for the cloud deployments of the LLM.
 - Consider cheapest deployment across the world due to latency in seconds

Input processing: 3500 tokens

Generating 99 tokens out



Inference Containers

- NeMo Inference Microservice — fresh release
 - Supports OpenAI-compatible API — killer feature
 - Supports NeMo LLM Service compatible API
 - Accelerated by TRT-LLM
- Triton + TRT-LLM
 - Part of NVAIE and can be supported
 - Works with HF models
 - Open Source

Catalog > Containers > Triton Inference Server

Triton Inference Server

Get Container

Features

☆ NVIDIA AI Enterprise Supported

Description

Triton Inference Server is an open source software that lets teams deploy trained AI models from any framework, from local or cloud storage and on any GPU- or CPU-based infrastructure in the cloud, data center, or embedded devices.

Publisher

NVIDIA

Latest Tag

23.08.04-py3-igpu

Modified

January 12, 2024

Compressed Size

4.6 GB

Multinode Support

Yes

Multi-Arch Support

Yes

23.08.04-py3-igpu

01/11/2024 11:38 PM 4.6 GB 1 Architecture

nvcr.io/nvidia/tritonserver:23.08.04-py3-igpu

23.12-py3-igpu-sdk

12/20/2023 3:21 AM 4.46 GB 1 Architecture

nvcr.io/nvidia/tritonserver:23.12-py3-igpu-s...

23.12-py3-igpu

12/20/2023 3:21 AM 5.1 GB 1 Architecture

nvcr.io/nvidia/tritonserver:23.12-py3-igpu

23.12-py3

12/20/2023 3:11 AM 7.19 GB 2 Architectures

nvcr.io/nvidia/tritonserver:23.12-py3

23.12-vllm-python-py3

12/20/2023 3:03 AM 10.68 GB 1 Architecture

nvcr.io/nvidia/tritonserver:23.12-vllm-pytho...

23.12-trtllm-python-py3

12/20/2023 2:58 AM 10.83 GB 2 Architectures

nvcr.io/nvidia/tritonserver:23.12-trtllm-pyt...

Paper to Understand Inference

Chunked Prefill

- <https://arxiv.org/pdf/2308.16369.pdf>

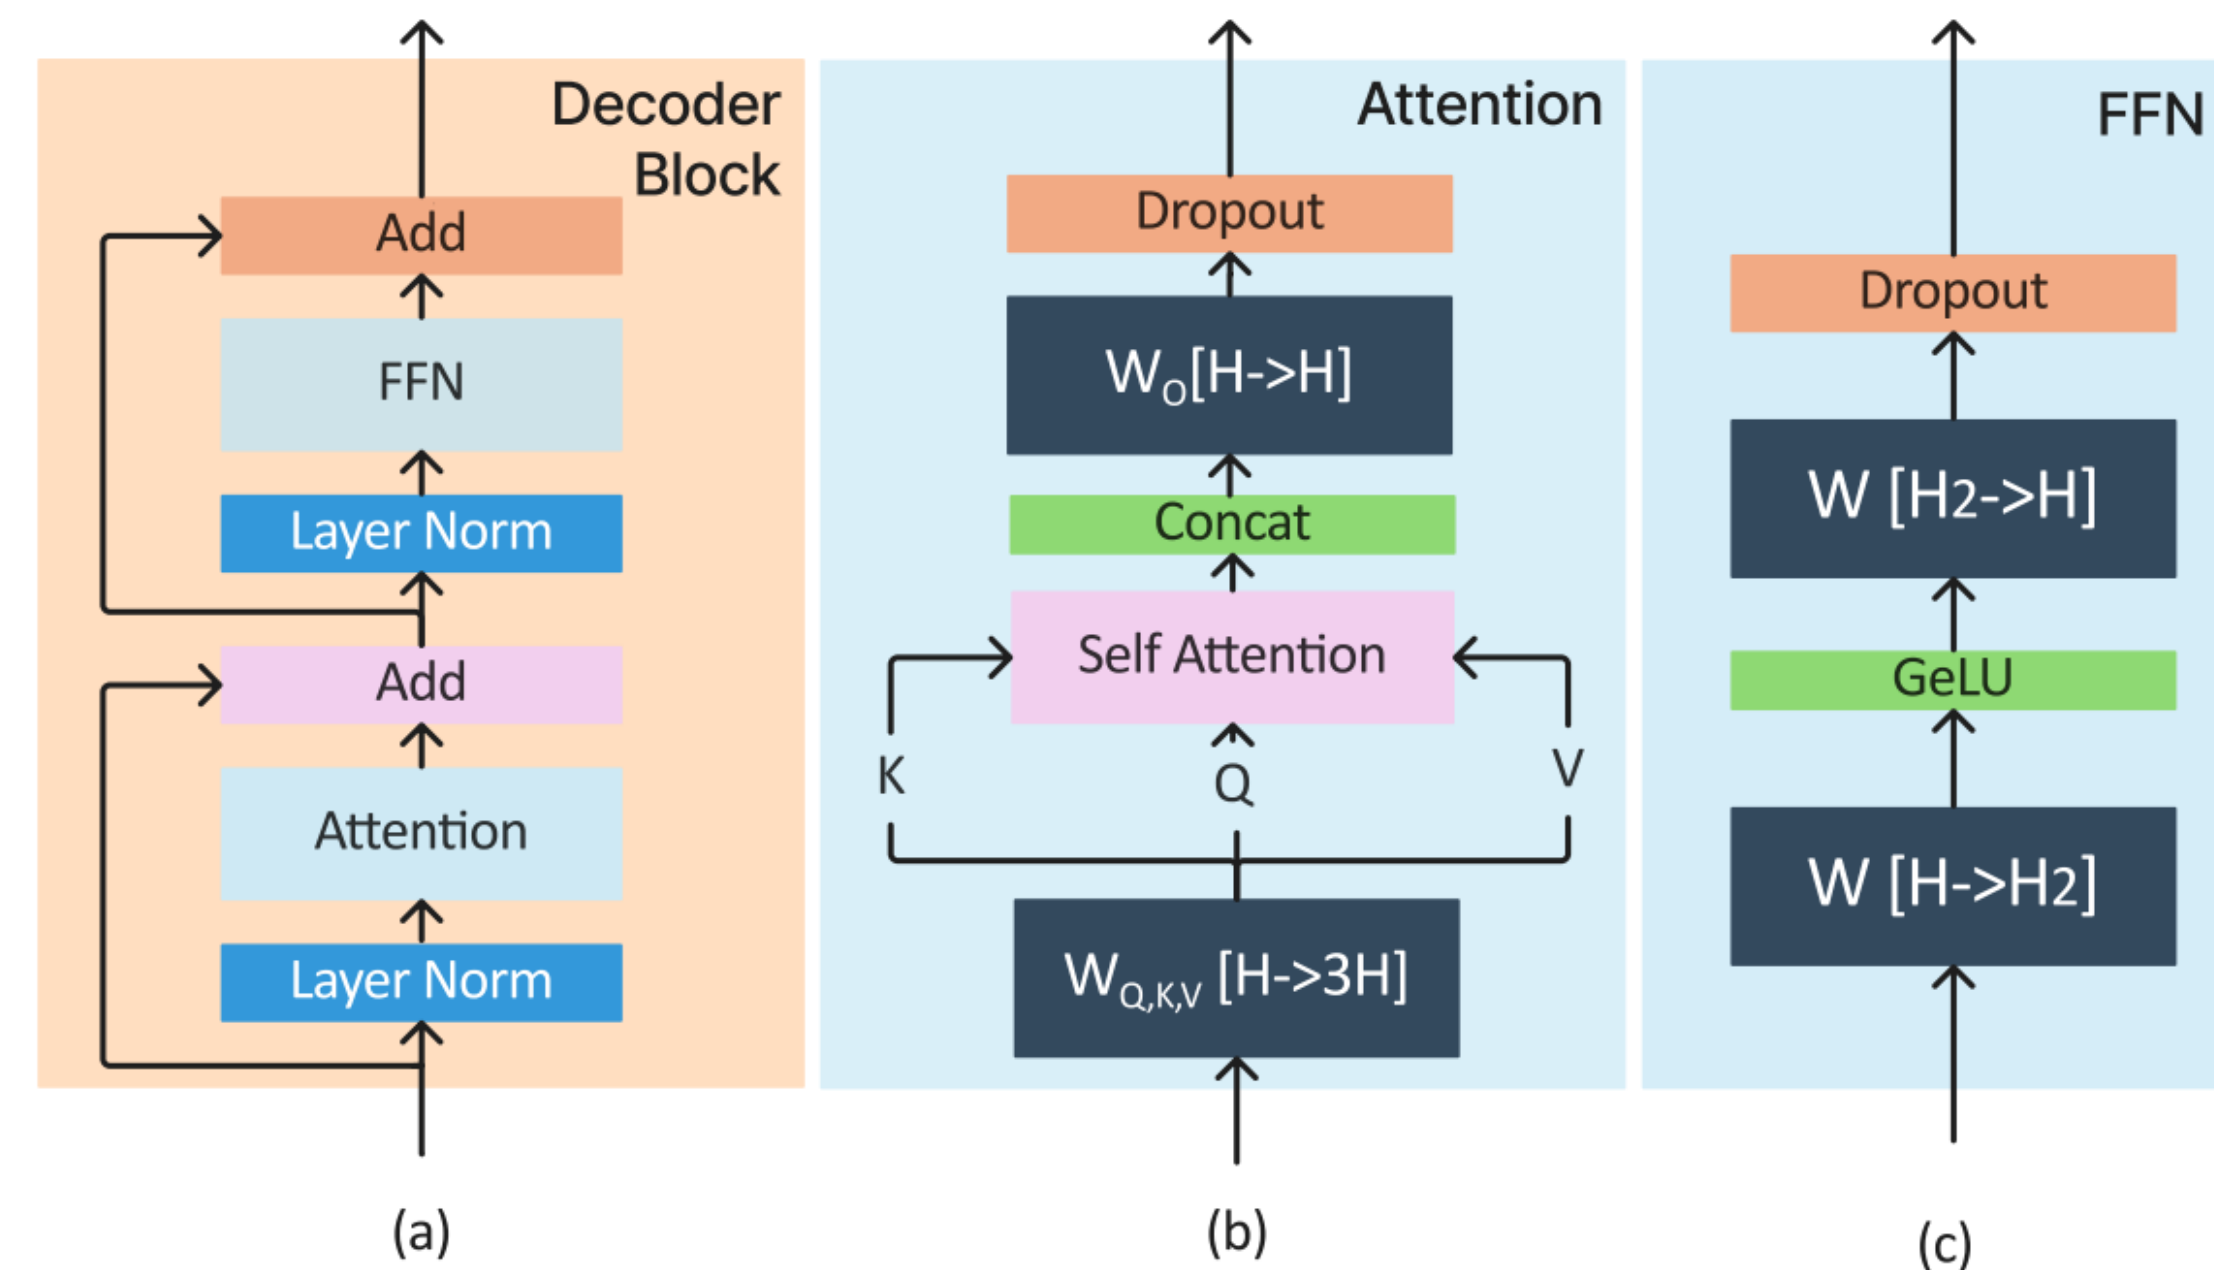


Figure 2: High-level architecture of a decoder block.

Operation	Shapes of tensors		
	Input(s)	Weight(s)	Output(s)
<i>preproj</i>	$[B, L, H]$	$[H, H]$	$[B, L, H]$
<i>attn</i>	$[B, L, H]$	-	$[B, L, H]$
<i>postproj</i>	$[B, L, H]$	$[H, H]$	$[B, L, H]$
<i>ffn_ln1</i>	$[B, L, H]$	$[H, H_2]$	$[B, L, H_2]$
<i>ffn_ln2</i>	$[B, L, H_2]$	$[H_2, H]$	$[B, L, H]$

Table 1: Shapes of the input, weight, and output tensors in a transformer decoder block. B, L and H denote batch size, embedding (aka hidden) size and sequence length (L=1 during decode, except for attention).





Figure 3: Per-token prefill and decode time with different batch sizes (sequence length = 1024) for LLaMa-13B on A6000 GPU. Prefill saturates GPU compute even at batch size of 1 and results in almost constant per-token time across batch sizes. Decode under-utilizes GPU compute and costs as much as $200\times$ prefill for batch size 1. The incremental cost of linear operators for decode is almost zero as batch size increases. The attention cost does not benefit from batch size as it is memory-bound.

Prefill Is Compute Bound in FFN

```
df_measured['sequence_position'] = df_measured['input_len']*df_measured['batch_size']
```

LLAMA2-7B-base-1.nemo, NVIDIA H100 80GB HBM3 Prefill




Decoding Is Memory Bound in KV Cache

But not so straightforward 😞

```
df_measured['sequence position'] = (df_measured['input len'] + (df_measured['output len'] - 1)/2) * df_measured['batch_size']
```

LLAMA2-7B-base-1.nemo, NVIDIA H100 80GB HBM3, TP 1 Decoding



Current Limitations and Extrapolations

- Extensive measurement are currently available only for NeMo Inference Container
 - PyTriton + TRT-LLM Python Backend
 - Only BF16 (expect approx. 30% improvement from switching to FP8 on H100 and L40s)
 - Client-side batching
- We expect the in-flight batching results to be compatible with the results of the NeMo Inference Sizing Calculator benchmarks
 - We expect the throughput to be the same
 - We expect the decoding latency to be the same
 - We expect the prefill latency to be close to the measured latency of prefill batch size 1
 - Results in significant improvement in first token latency

Call to Action

Use the Personal Checklist to understand your requirements

- Use the [sizing checklist](#) for your use cases
- Get a ballpark estimate of required HW using the calculator: <https://nemo-inference-sizing.nvidia.com/> (when it gets published)
- Need clarifications or help with sizing? Drop message to me and my colleague: Dmitry Mironov dmitrym@nvidia.com, Sergio Perez sergiop@nvidia.com




Dmitry Mironov
dmitrym@nvidia.com



Sergio Perez
sergiop@nvidia.com


Inference Resources

- <http://nemo-inference-sizing.nvidia.com/>
- [NVIDIA NeMo Microservices Early Access](#)
- [OpenAI pricing + token counter](#)
- [What is RAG — NVIDIA blog](#)
- [Mastering LLM Techniques: Inference Optimization — NVIDIA Blog](#)



Recommended Configurations Within Limits

Metric	Throughput-optimal ▼	Latency-optimal ▲	Clicked data
Latency for first token (ms)	1516.7	37.3	[Click on a point]
Latency per generated token (ms)	41.6	17.2	[Click on a point]
Prompts per second per 8 GPUs	11.2	0.1	[Click on a point]
Tensor Parallelism	4	8	[Click on a point]
Batch size	128	1	[Click on a point]



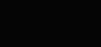
Apply for NVIDIA NeMo Microservices Early Access

NVIDIA NeMo™ microservices is a collection of containerized software to easily and rapidly build and deploy large language model (LLM) workloads for enterprise use cases.

NeMo microservices is in private, early access and provides the easiest, most performant way of deploying LLMs on your preferred infrastructure (on-prem or cloud), and also supports inference on embedding models for retrieval-augmented generation applications.

The container comes with 12 optimized prebuilt models (NVIDIA TensorRT™ engines) that can be deployed out of the box: Llama 2 (7B, 13B, and 70B), NVIDIA Nemotron-3 8B and 43B, StarCoder, StarCoder plus, and NVIDIA NeMo Retriever text QA embedding model.

All models are curated with the optimal hyperparameters, and inference APIs are provided that are compatible with OpenAI APIs.

 NVIDIA

NGC | PRIVATE REGISTRY

CATALOG

CONSOLE

PRIVATE REGISTRY

Collections

Containers

Helm Charts


Models

Resources

Private Registry > Containers > NeMo Microservice Inference

NeMo Microservice Inference

Get Container

 NVIDIA

NEMO

Description

NeMo Microservice Inference GPU accelerated LLM capabilities through OpenAI compatible APIs

Publisher

NVIDIA

Latest Tag

23.12

Modified

January 8, 2024

Compressed Size

13.49 GB

Multinode Support

No

Multi-Arch Support

No

23.12 (Latest) Security Scan Results

No results available.

Overview

Tags

Layers

Security Scanning

Related Collections

Overview

The NeMo Microservice Inference (NMI) brings the power of state-of-the-art large language models (LLM) to your applications, providing unmatched natural language processing and understanding capabilities. Whether you're developing chatbots, content analyzers—or any application that needs to understand and generate human language—NMI has you covered. Built on the NVIDIA software platform incorporating CUDA, TRT, TRT-LLM, and Triton, NMI brings state of the GPU accelerated Large Language model serving.

High Performance Features

NMI enables optimized scheduling technique called in-flight batching. This technique takes advantage of the fact that the overall text generation process for an LLM can be broken down into multiple execution iterations on the model. With in-flight batching, rather than waiting for the whole batch to finish before moving on to the next set of requests, the NMI runtime immediately evicts finished sequences from the batch. The runtime then begins executing new requests while other requests are still in flight.

Scalable Deployment

Whether you're catering to a few users or millions, the microservice scales seamlessly to meet your demands.

Advanced Language Models

Built on cutting-edge LLM architectures, NMI provides optimized and pre-generated engines for a variety of popular models. The tooling to create GPU optimized models is also included.

Flexible Integration

Easily incorporate the microservice into existing workflows and applications, thanks to multiple API endpoints.

Secure Processing

Your data's privacy is paramount. NMI ensures that all inferences are processed securely, with rigorous data protection measures in place.

Applications

Chatbots & Virtual Assistants

Empower your bots with human-like language understanding and responsiveness.

Content Generation & Summarization

Generate high-quality content or distill lengthy articles into concise summaries with ease.

Sentiment Analysis

Understand user sentiments in real-time, driving better business decisions.

Language Translation

Break language barriers with efficient and accurate translation services.

And many more... The potential applications of NMI are vast, spanning across various industries and use-cases.

