



How LLMs are Unlocking New Opportunities for Enterprises

Part 2

Contributors:

Annamalai Chockalingam

Ankur Patel

Shashank Verma

Tiffany Yeung




Table of Contents

- Overview 3
- Traditional NLP Tasks Performed by Large Language Models..... 5
 - Content Generation..... 6
 - Summarization..... 7
 - Translation 9
 - Classification 11
 - Chatbot Support 12
- Case Study: Korea Telecom X NeMo Megatron 15

Overview

Most enterprises need to perform many language-related tasks every day. They need to organize documents based on common themes (*text classification*), write emails (*content generation*), remove harmful posts from their online communities (*toxicity classification*), and much more. Applications powered by large language models can help enterprises automate these and many other tasks, helping them to streamline their operations, decrease expenses, and increase productivity.

Alternatively, enterprises can use LLM-powered apps to provide innovative and improved services to clients or strengthen customer relationships. For example, enterprises could provide customer support via AI companions or use sentiment analysis apps to extract valuable customer insights.

This section will examine both applications of LLMs – in internal *and* external operations by reviewing the six LLM-powered text apps listed in Table 1.

Table 1. LLM-powered Text Applications

Text Application Category	Example Use Cases
Content generation	Marketing content and copy, paraphrasing, email composition
Summarization	Legal summarization, news summarization
Translation	Language translation, code translation
Classification	Toxicity classification, sentiment analysis, fraud analytics
Chatbot support	EQA and ODQA chatbots, AI companions
Tools and technologies	Character generation

It's worth noting that large language models can be applied to other content categories besides text. They're currently widely used for speech-, image-, and video-related tasks, such as image generation or video classification. Some of the applications presented in this document may be capable of performing such tasks as well.

However, the following sections primarily aim to familiarize enterprises with *text-related* LLM solutions. They'll explain what makes LLMs superior to traditional solutions, such as rule-based systems or manual approaches, as well as show how enterprises can implement them depending on their industry. This should clarify the immense value of LLMs for enterprises, as well as inspire organizations to develop their own LLM variants.

Traditional NLP Tasks Performed by Large Language Models

Traditional NLP tasks can be thought of as the foundational and more basic iterations of contemporary use cases. Chatbot support, for example, can be considered a sophisticated combination of two traditional NLP tasks, text generation and question-answering. These and some other basic NLP tasks are listed in Table 2.

Table 2. Traditional NLP Tasks

NLP Task	Function
Machine translation	Translations of text from one language to another
Text summarization	Generation of summaries of texts
Sentiment analysis	Analysis of text sentiment (e.g., positive, negative, and neutral)
Text classification	Classification of texts into categories
Text generation	Generation of texts
Question answering	Answering questions
Part-of-speech tagging	Assignment of part-of-speech tags, like “noun” or “verb,” to words or tokens
Named entity recognition	Identification and classification of named entities in text, such as people, organizations, and locations

This brief, albeit incomplete overview of traditional NLP tasks may help with understanding how to fine-tune proprietary models or select well-suited base models, a topic further explored in *Part 3: How Enterprises Can Build Large Language Models*. It should also be helpful in understanding the mechanics behind more advanced use cases analyzed in the remaining sections.

Content Generation

Content generation refers to producing any type of content, from social media posts to annual reports. Most enterprises need to create significant amounts of various types of content for a wide range of purposes and for both *internal* and *external* use.

Table 3 list common content types created for internal and external use within enterprises.

Table 3. Content Types Sorted by External and Internal Use

External Use	Internal Use
Blog and social media posts	Training materials
Press releases	Performance reports
Customer presentations	Internal memos
Web copy	Company announcements
Customer newsletter	Employee surveys

At the same time, enterprises face multiple challenges that can hinder their content-generation efforts and make it harder to achieve tangible results. Main content-related issues include:

- > **Increasing content production demands.** Today’s digital landscape requires businesses to consistently produce vast amounts of content. This is underscored by the fact that 4.4 million blog posts were published daily in 2022, with 69% of businesses planning to increase their content marketing budgets in 2023. Enterprises need to keep up with these standards to stay competitive but may not be able to do so solely via manual methods, especially as other businesses start embracing AI content generation.
- > **Increasing email composition demands.** Reading and writing emails consumes as much as 28% of workweeks, or 13 hours every week, that employees could spend on more important tasks. Additionally, email composition is also one of the major stressors within enterprises, with approximately 92% of employees experiencing elevated blood pressure and heart rate while handling work emails.
- > **The lack of real data within enterprises.** Many enterprises lack the data they need to conduct important progress-making activities like research or staff training. Some don’t have the resources to collect real data, while others cannot use real data without compromising client privacy. This can be especially problematic in industries like finance or healthcare, where data often contains sensitive information.

Enterprises can overcome these challenges by supplementing or replacing manual methods with LLM-powered content generation applications. This can help employees focus on more important tasks.

Depending on what they specialize in, content generation applications can help enterprises accelerate their content production and email writing processes, develop safe-to-use synthetic data, or even produce new, creative types of content to distinguish themselves from competitors.

- > **Marketing content and copy** generation apps help enterprises automate their copy and content production and customize the texts according to their brand and target audiences.
- > **Paraphrasing** apps enable enterprises to quickly rewrite existing content. This can help them create new, plagiarism-free content, ensure data privacy by removing sensitive information, adapt texts to a specific audience, and more.
- > **Email composition** apps can help enterprises accelerate email writing by providing context-dependent auto-completions in real time or generating complete emails from scratch based on user prompts.
- > **Creative writing** generation apps can enhance brand storytelling and help diversify an enterprise's content offerings by generating creative texts, such as books or poems.
- > **(Other) data generation** apps can equip enterprises with synthetic data, such as artificial financial reports, health records, or customer profiles, and help them avoid time-consuming data collection or privacy-preserving procedures.

Examples of Content Generation

- > **Jasper** is a GPT-3-powered AI writing assistant claiming to serve over 100,000 teams at well-known companies like Mayo Clinic, Intel, and Zoom. It can generate various types of content, including social media posts, blog posts, and email content, based on natural language prompts entered by users. [One case study](#) showed that Jasper-generated content can help companies grow website traffic by 166% in just two months. It also demonstrated that publishing AI-generated content doesn't necessarily cause backlash from search engines.
- > **SyngatorTron** is the largest clinical language generator to date developed jointly by the University of Florida and NVIDIA. The 5-billion parameter model was trained on data representing more than 2 million patients, with the purpose of creating synthetic patient profiles that mimic the health records from which it has learned. Unlike real patient data, SyngatorTron-generated health records don't have to be de-identified before use because they don't contain real, sensitive information. Enterprises can use them to create new models, tools, or tasks more rapidly, such as smarter chatbots that improve patient experience or summarization tools that reduce physician burnout.

Summarization

As enterprises grow and generate more data, it becomes increasingly difficult for employees and decision-makers to keep up with relevant information. Summarization can help them stay up to date by transforming long texts into concise summaries and extracting key information.

Enterprises have previously summarized texts by employing human workers or using smaller, mainly template-based models. However, both of these methods have some disadvantages that may limit their effectiveness and benefits for enterprises.

- > **Manual summarization** is tedious and prone to human error, as it highly depends on the skills and judgment of the summarizer. For example, if the summarizer is not skilled in a particular domain or the process of summarization itself, enterprises may run the risk of making flawed decisions based on inaccurate summaries.
- > **Template-based summarization models** are limited by human-crafted and often rigid and inflexible templates. The templates hinder their ability to adapt to new information or types of input, as well as to produce varied output. They require new templates to be made for each new use case, making them time-consuming and inefficient in the long run. Lastly, the quality of their summaries highly depends on the accuracy and completeness of the templates. If the templates are inaccurate, the summaries may be inaccurate as well.

LLM-powered summarization apps minimize the risk of human error, as they don't have personal biases or opinions that could influence their summarization processes. They can also process and learn from vast amounts of data much faster than humans, which makes them more efficient and accurate over time.

This ability to learn from data also makes LLM-based summarizers much more versatile than template-based models, as they don't require human agents to craft templates for each new or edge-case scenario. They can summarize various types of text, from financial reports to conversation transcripts, with no human intervention.

While summarization apps can be applied to all types of texts, legal and news summarization may have the most benefits for an average enterprise.

- > **Legal summarization** apps help enterprises convert lengthy legal documents into concise, easy-to-read summaries. It may also help them extract key information from texts to make it more accessible, such as referenced laws or main arguments. Enterprises can, for example, use them to ensure their compliance with various laws and regulations.
- > **News summarization** apps summarize news articles in order to help enterprises stay on top of events that may influence their operations or decision-making processes. For example, airlines can use it to stay up-to-date with weather issues or incidents that may impact their routes.

Examples of Summarization

- > **Instruct Davinci** is, according to OpenAI, the most capable GPT-3 model. [A study](#) suggests that it produces more extractive summaries than human writers, which may decrease the originality of summaries but also increase their accuracy. Additionally, the study found that human evaluators equally like summaries created by freelance writers and Instruct Davinci, suggesting that summaries generated by high-quality LLMs can be on par with those created by human agents.
- > **genei** is a GPT-3-powered tool specializing in summarizing research papers, paraphrasing complex ideas, and helping users find key information faster. According to the company, 95% of users say they find greater answers and insights from their work when using genei, while 98% of users believe genei's keyword extraction and summarization features save them time.

Translation

Translation typically refers to translating texts from one natural language to another, for example from French to English. As such, it's especially important for enterprises that serve global markets or have international teams.

However, translation can also be applied to other types of languages, such as programming languages like Python, JavaScript, or Ruby. In this context, translation is particularly important for tech enterprises but is often necessary for other industries as well. For example, even non-tech enterprises frequently need to translate code from outdated to more modern programming languages to make it more efficient and secure.

In the past, many enterprises tried meeting these translation-related needs by hiring human experts. Some also tried automating the process by using simple machine translation systems. Unfortunately, both methods are proving to be inadequate for contemporary demands.

- > **Pre-LLM machine translation systems** often had limited vocabulary, due to being trained on smaller amounts of data. They also lacked the contextual understanding needed to produce accurate translations. This often resulted in grammatically correct but overly literal or even nonsensical outputs.
- > **Manual translation** is usually more accurate but has the usual disadvantages of manual methods. It's time-consuming, prone to human error, and can lead to inconsistencies in tone or style, especially when conducted by more than one translator.

LLM-powered translation apps can help enterprises overcome these challenges but also unlock new possibilities, such as translating text to code or code to text.

- > **Natural language translation** apps powered by LLMs can understand more context than earlier systems and consider subtler linguistic nuances, such as text sentiment, when generating their predictions. This usually results in more accurate translations. They're also more flexible and versatile thanks to being trained with massive amounts of data.
- > **Code conversion** apps help enterprises convert code from one programming language to another, thus reducing the workload for developers and potentially helping them prevent security vulnerabilities and bugs.
- > **Code translation** apps unlock previously unheard-of use cases by allowing users to translate code to natural language or, vice versa, translate natural-language prompts into corresponding code. Non-programmers can use them to learn how to code or even write code themselves using prompts in plain English or another human language. Programmers, on the other hand, can use them to accelerate their coding processes, learn new programming languages, or quickly review the functions of different pieces of code.

Examples of Translation

- > **DeepL Translator** is a natural-language translation app that runs on DeepL's proprietary-developed models. The company claims their models can accurately convey meaning even when

inputs contain industry-specific jargon, and [a blind study](#) conducted in 2021 seems to confirm its claims. The study found that DeepL is 3x more accurate than translation apps developed by tech giants like Microsoft and Google.

- > **Tabnine** is an AI coding app boasting clients like Ericsson, Amazon, and SpaceX. Its service runs multiple code-native LLMs on A100 GPUs with NVIDIA Triton to handle more than 20 programming languages and 15 code editors. According to [the company](#), the app automates up to 30% of code for over a million developers globally.
- > **GitHub Copilot** is an AI pair programmer that writes code alongside users. It provides coding suggestions for users based on natural-language prompts or previous inputs, as well as natural-language explanations of selected pieces of code. It has already [shown promise](#) in classroom settings, where it can help students learn by generating code explanations, creating tests, and fixing bugs.

Classification

In the context of NLP, classification refers to the categorization of texts according to pre-defined criteria. For example, texts can be organized into groups based on topics, format, length, author, and so on.

The goal of such tasks may be classification itself, as can be the case with document or email classification, or another purpose, such as identifying indications of fraud. An overview of common classification tasks and their purposes is provided in Table 4.

Table 4. Common Classification Tasks

Classification Task	Explanation	Common purpose
Toxicity classification	Classifying texts as toxic and non-toxic based on whether the texts violate predefined guidelines and rules	Community moderation, monitoring employee communication
Customer classification	Classifying customers into groups based on shared traits, like purchasing habits or interests	Marketing, customer churn prediction
Sentiment analysis	Interpreting and classifying the emotional tone of text data, such as customer reviews or social media posts	Product development, increased customer or voter understanding
Email classification	Classifying emails into different folders or categories based on their content, such as spam, social, promotions, or personal	Email organization, minimizing the risk of phishing or other malicious attacks
Document classification	Classifying documents based on their topic, type, domain, or another criterion	Easier document management and retrieval, compliance with regulations such as GDPR
Product categorization	Organizing products into different categories or subcategories based on shared attributes	Personalized product recommendations, improving inventory systems and e-commerce platforms, streamlining supply chains
Fraud analytics	Classifying patterns in data as anomalies or non-anomalies	Preventing fraud and future fraudulent behavior

As in previous cases, enterprises mainly performed these tasks using manual methods or rule-based systems.

- > **Manual classification** is a labor-intensive process that distracts employees from performing higher-priority tasks.
- > **Rule-based systems** are often unreliable because they rely on rules and exact matches to accurately classify texts. For example, rule-based toxicity classifiers can allow huge amounts of toxic content to go undetected as long as they don't contain exact "toxic" keywords listed in their databases.

LLM-powered apps, on the other hand, classify texts into groups based on the overall context and meaning rather than straightforward features like used keywords. This usually results in more accurate classifications and makes LLMs more versatile.

Examples of Classification

- > **SambaNova's GPT model** can be applied to a variety of NLP tasks, including document classification. It's been shown to predict the topic of customer complaint emails with 90% accuracy, outperforming BERT by almost 5%. The company attributes this achievement to GPT's higher contextual understanding and ability to process a higher sequence length compared to BERT. As such, SambaNova's model eliminates the need for manual email reviews, reducing the time required to resolve customer complaints by an average of 2-5 minutes per email.
- > **Blue Silk™ GPT model** powers Talkwalker's customer classification and social listening tools. Among other functions, these tools are meant to help enterprises analyze public perception on different topics via sentiment analysis of customer-generated texts. Companies can use Blue Silk's insights to adapt their communications or identify viral topics they could capitalize on. [A case study](#) shows how Europe's leading temperature-controlled logistics company, STEF, leveraged Blue Silk-generated insights to double its LinkedIn followers and fine-tune its communication with stakeholders.

Chatbot Support

Enterprises began to use chatbots around the 1990s in order to provide adequate-yet-affordable customer service. Many would later also invest in building internal chatbots meant to support their employees and help them solve problems or find answers faster.

This helped enterprises reduce the workload of human agents and internal support staff, decrease their expenses, and ultimately improve customer experience by providing real-time support 24/7.

However, as in other cases, earlier chatbot systems do have prominent disadvantages that can make them inadequate for today's expectations and demands.

- > **Rule-based chatbots** can only handle a narrow range of requests that match the pre-defined, human-crafted rules. This prevents enterprises from providing comprehensive support to customers and employees with unique or complex requests.

- > **Early algorithms and approaches to NLP** also didn't allow earlier systems to adequately process unstructured data, such as customer reviews or emails. This limited their ability to understand and respond to natural-language requests and often resulted in unsatisfactory or irrelevant responses.

In contrast, LLM-powered chatbots – also called *conversational AI* – can handle a broad range of requests and adequately process unstructured data thanks to advanced ML algorithms. They can also detect subtle linguistic nuances and use them to draw conclusions about customer preferences or sentiment. This allows them to offer more personalized support and engaging responses.

On top of improving previously known use cases, LLM-powered systems also open up opportunities for new ones. For example, they can help companies provide emotional support to customers and employees. Some of these use cases are mentioned below.

- > **Extractive and open-domain question-answering (QA) systems** are capable of providing accurate answers by *extracting* them from given pieces of text rather than, for example, freely generating answers on their own. They can also answer questions from many domains, such as legal, medical, and financial, instead of just one, making them much more universal than earlier systems.
- > **Structured + unstructured data QA systems** can answer any type of question, regardless of whether the answer can be found in structured data like databases or unstructured data like customer reviews or social media posts. This makes them far more versatile than traditional systems that could mainly only process structured data. Additionally, they can also help enterprises extract valuable insights from customer questions or requests, such as their preferences or pain points.
- > **AI companions** are designed to engage in conversations with users as if they were friends. Enterprises can use them to provide emotional support to employees or to offer new, conversation-based services to customers. For example, healthcare enterprises can use AI companions to provide follow-up support to patients, while educational enterprises can use them to give students personalized academic advice.

Examples of Chatbot Support

- > **Replika** is an AI companion that converses with users and helps them build better habits to overcome emotional or mental health issues. It improves its abilities and user understanding over time by learning from previous inputs and feedback, making the support more personalized and effective over time.
- > **AskBrian** can be considered an AI companion, as it can converse with users, but also has many other functions primarily aimed to help consultants and business professionals. Besides answering users' questions or helping them generate new ideas, AskBrian can also perform tasks like converting files from PDF to Word, analyzing private and public companies, or quickly comparing a company with its peers.

Case Study: Korea Telecom X NeMo Megatron

Korea Telecom (KT) is South Korea's leading mobile operator with over 22 million subscribers.

In 2022, the telecom company decided to develop an LLM for the Korean language Hangul in order to improve their smart speaker, called GiGa Genie, and customer service platform, AI Contact Center (AICC).

- > **The goal for GiGa Genie** was to enhance its understanding of the Korean language and enable it to generate more human-like sentences.
- > **The goal for AICC** was to improve its handling, summarization, and classification of customer inquiries in order to reduce consultation times.

However, in order to achieve these goals, KT needed to overcome several challenges. This included the complexity of Hangul, the insufficiency of most LLM-development platforms, and the fact that KT needed to collect and use their own, Hangul-based datasets.

- > **Hangul's complexity.** Hangul often shows up in lists of the world's most challenging languages. Training an LLM to accurately process and generate such an intricate language requires cutting-edge training techniques and tools, which KT lacked and didn't want to develop on their own.
- > **The shortage of full-stack platforms.** KT also needed full-stack tools, from the inference to the hardware level. Most LLM-development platforms, however, offered only partial solutions. If KT used such platforms, the company would've needed to develop proprietary custom solutions, making the already complex process even more demanding.
- > **Time- and resource-consuming data collection and preparation processes.** KT also needed to collect and pre-process proprietary datasets to customize their LLM for Hangul. However, doing so would've been extremely time-consuming and expensive, especially if KT didn't have access to adequate data processing tools.

To overcome these challenges, KT decided to train, customize, and deploy their LLM with [NeMo Megatron](#).

NeMo is NVIDIA's end-to-end framework for training and deploying LLMs. It equipped KT with expert support and full-stack tools, allowed them to accelerate and improve model training, and helped them achieve high throughput.

- > **Expert support.** NVIDIA provided expertise from product to engineering teams and helped KT easily solve several technical issues. NeMo also helped them automatically find the best configurations for training and inference, enabling ML engineers to skip time-consuming trial-and-error processes.
- > **Access to full-stack environments.** NeMo equipped KT with all the necessary tools for training, customizing, and deploying LLMs, including tools for easy neural network development and data processing. This allowed KT to prepare and refine their Hangul-based datasets more easily.
- > **More effective LLM training.** Finally, NeMo gave KT access to state-of-the-art training methods, optimized algorithms, and hyperparameter optimization tools, which accelerated the training process and improved model accuracy.

After integrating the LLM into GiGa Genie and AICC, KT saw remarkable improvements with both services.

- > The newly-developed LLM enabled **GiGa Genie** to successfully interpret voice commands given in Hangul and complete a slew of home-assistance tasks. Its ability to generate human-like sentences has significantly improved as well.
- > The LLM also enabled **AICC** to more accurately classify and summarize customer inquiries, reducing consultation times by an impressive 15 seconds. AICC now manages more than 100,000 calls daily across Korea, providing customers with requested information or quickly connecting them with human agents.

Considering the success they had, KT is not planning to withdraw from the LLM arena anytime soon. In fact, the company is planning to further improve and diversify their LLM-based product offerings.

"Thanks to LLMs, KT can release competitive products faster than ever," said Hwijung Ryu, LLM development team lead at KT. "We also believe that our technology can drive innovation from other companies, as it can be used to improve their value and create innovative products," Ryu concluded.

The value of LLMs in the enterprise is only increasing. For developers interested in integrating LLMs into their workflows, stay tuned for Part 3 of our eBook, which will share best practices and hands-on tips for developing LLM-based enterprise applications.

Notice

This document is provided for information purposes only and shall not be regarded as a warranty of a certain functionality, condition, or quality of a product. NVIDIA Corporation ("NVIDIA") makes no representations or warranties, expressed or implied, as to the accuracy or completeness of the information contained in this document and assumes no responsibility for any errors contained herein. NVIDIA shall have no liability for the consequences or use of such information or for any infringement of patents or other rights of third parties that may result from its use. This document is not a commitment to develop, release, or deliver any Material (defined below), code, or functionality.

NVIDIA reserves the right to make corrections, modifications, enhancements, improvements, and any other changes to this document, at any time without notice.

Customer should obtain the latest relevant information before placing orders and should verify that such information is current and complete.

NVIDIA products are sold subject to the NVIDIA standard terms and conditions of sale supplied at the time of order acknowledgement, unless otherwise agreed in an individual sales agreement signed by authorized representatives of NVIDIA and customer ("Terms of Sale"). NVIDIA hereby expressly objects to applying any customer general terms and conditions with regards to the purchase of the NVIDIA product referenced in this document. No contractual obligations are formed either directly or indirectly by this document.

NVIDIA products are not designed, authorized, or warranted to be suitable for use in medical, military, aircraft, space, or life support equipment, nor in applications where failure or malfunction of the NVIDIA product can reasonably be expected to result in personal injury, death, or property or environmental damage. NVIDIA accepts no liability for inclusion and/or use of NVIDIA products in such equipment or applications and therefore such inclusion and/or use is at customer's own risk.

NVIDIA makes no representation or warranty that products based on this document will be suitable for any specified use. Testing of all parameters of each product is not necessarily performed by NVIDIA. It is customer's sole responsibility to evaluate and determine the applicability of any information contained in this document, ensure the product is suitable and fit for the application planned by customer, and perform the necessary testing for the application in order to avoid a default of the application or the product. Weaknesses in customer's product designs may affect the quality and reliability of the NVIDIA product and may result in additional or different conditions and/or requirements beyond those contained in this document. NVIDIA accepts no liability related to any default, damage, costs, or problem which may be based on or attributable to: (i) the use of the NVIDIA product in any manner that is contrary to this document or (ii) customer product designs.

No license, either expressed or implied, is granted under any NVIDIA patent right, copyright, or other NVIDIA intellectual property right under this document. Information published by NVIDIA regarding third-party products or services does not constitute a license from NVIDIA to use such products or services or a warranty or endorsement thereof. Use of such information may require a license from a third party under the patents or other intellectual property rights of the third party, or a license from NVIDIA under the patents or other intellectual property rights of NVIDIA.

Reproduction of information in this document is permissible only if approved in advance by NVIDIA in writing, reproduced without alteration and in full compliance with all applicable export laws and regulations, and accompanied by all associated conditions, limitations, and notices.

THIS DOCUMENT AND ALL NVIDIA DESIGN SPECIFICATIONS, REFERENCE BOARDS, FILES, DRAWINGS, DIAGNOSTICS, LISTS, AND OTHER DOCUMENTS (TOGETHER AND SEPARATELY, "MATERIALS") ARE BEING PROVIDED "AS IS." NVIDIA MAKES NO WARRANTIES, EXPRESSED, IMPLIED, STATUTORY, OR OTHERWISE WITH RESPECT TO THE MATERIALS, AND EXPRESSLY DISCLAIMS ALL IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY, AND FITNESS FOR A PARTICULAR PURPOSE. TO THE EXTENT NOT PROHIBITED BY LAW, IN NO EVENT WILL NVIDIA BE LIABLE FOR ANY DAMAGES, INCLUDING WITHOUT LIMITATION ANY DIRECT, INDIRECT, SPECIAL, INCIDENTAL, PUNITIVE, OR CONSEQUENTIAL DAMAGES, HOWEVER CAUSED AND REGARDLESS OF THE THEORY OF LIABILITY, ARISING OUT OF ANY USE OF THIS DOCUMENT, EVEN IF NVIDIA HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Notwithstanding any damages that customer might incur for any reason whatsoever, NVIDIA's aggregate and cumulative liability towards customer for the products described herein shall be limited in accordance with the Terms of Sale for the product.

Trademarks

NVIDIA, the NVIDIA logo are trademarks and/or registered trademarks of NVIDIA Corporation in the U.S. and other countries. Other company and product names may be trademarks of the respective companies with which they are associated.

Copyright

© 2023 NVIDIA Corporation. All rights reserved.