



# SUPERCHARGED COMPUTING FOR THE DA VINCIS AND EINSTEINS OF OUR TIME

We set out 26 years ago to transform computer graphics.

Fueled by the massive growth of the gaming market and its insatiable demand for better 3D graphics, we've evolved the GPU into a computer brain at the intersection of virtual reality, high performance computing, and artificial intelligence.

NVIDIA GPU computing has become the essential tool of the da Vincis and Einsteins of our time. For them, we've built the equivalent of a time machine.




# TWO FORCES SHAPING COMPUTING


For 30 years, the dynamics of Moore's law held true. But now CPU scaling is slowing while the demand for computing power surges ahead.

With AI, machines can learn. AI can solve grand challenges that have been beyond human reach. But it must be fueled by massive compute power.

Accelerated computing is the path forward beyond Moore's law, delivering 1,000X computing performance every 10 years.



40 YEARS OF CPU TREND DATA

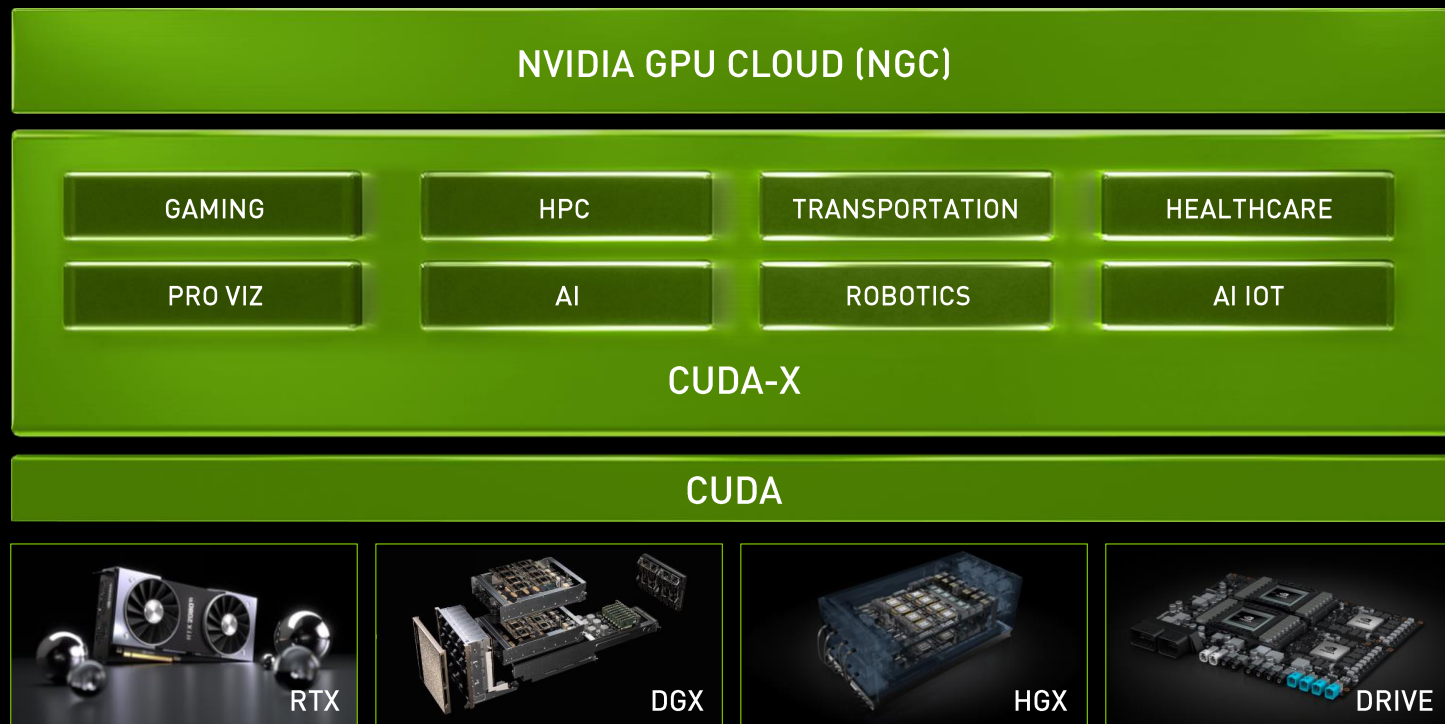


ALEXNET: THE SPARK OF THE MODERN AI ERA

# ONE ARCHITECTURE

NVIDIA is an accelerated computing company. It starts with a highly specialized parallel processor called the GPU and continues through system design, system software, algorithms, and optimized applications.

CUDA-X® is a suite of software libraries that accelerate applications for our growth markets — from gaming to transportation to healthcare — all based on a common CUDA architecture supported by more than 1.2 million developers today.


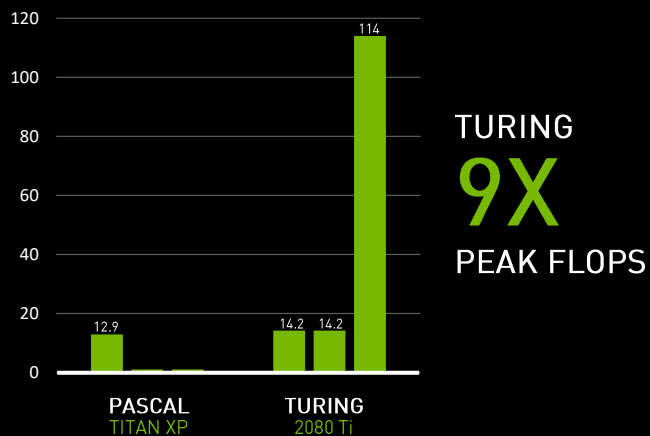




# TURING — A GIANT LEAP FOR GRAPHICS AND AI

The revolutionary NVIDIA Turing™ GPU architecture fuses together real-time ray tracing, artificial intelligence, and programmable shading to reinvent professional graphics and gaming.

With over 100 TFLOPS of performance, Turing can process deep learning models to create special effects, enhance images, and animate AI characters.





# NVIDIA REINVENTS MODERN COMPUTER GRAPHICS

Our invention of the GPU in 1999 made real-time programmable shading possible, giving artists an infinite palette for expression.

In 2018, the introduction of the Turing architecture and NVIDIA RTX™ ray-tracing technology fulfilled another vision of computer scientists, paving the way to new levels of art and realism in real-time graphics.

We've led the field of visual computing for decades.

# A NEW ERA FOR VISUAL EFFECTS

Turing-based Quadro® RTX delivers photoreal graphics that creators didn't expect for another 5-10 years.

Quadro RTX GPUs can now accelerate photoreal rendering for large industries that previously only used CPU server farms: film, animation, architecture, product design, and others.

NVIDIA has reinvented computer graphics, again.





# ONE 3D DESIGN APP TO UNITE THEM ALL

There are more than 200 animation studios in the world today – and none of them work alone. It's a collaboration-heavy industry without a way to work together in real time. Until now.

NVIDIA Omniverse integrates the pieces of the 3D production pipeline into a single virtual design studio.

Now the work of many can be seen by all, in real time and in the splendor of Omniverse's RTX-based viewport.





# RTX ARRIVES IN GRAPHICS DATA CENTERS

Designers and artists across industries are producing amazing digital content faster than ever, but CPU render farms aren't keeping up.

The NVIDIA RTX Server packs 40 Turing GPUs into a single machine capable of running graphics apps for 320 users at once.

RTX Server is software-optimized to deliver incredible performance and economics for rendering, remote workstations, and cloud gaming.



# RTX RESETS GAMING

GeForce® RTX has redefined what's possible in gaming. Real-time ray tracing and neural graphics processing come together to create eye-popping images and deliver a level of photorealism never before seen in PC gaming.

AAA games like *Battlefield V*, *Shadow of the Tomb Raider*, and *Metro Exodus* support RTX today. And with support in Microsoft DXR, Unreal Engine, and Unity, next-generation games can easily bring ray tracing to millions of gamers.

Microsoft®  
DirectX®





# GEFORCE — THE WORLD'S LARGEST GAMING PLATFORM

Gaming is the world's largest entertainment industry. With 200 million gamers, NVIDIA GeForce is its largest platform. GeForce RTX GPUs and the GeForce Experience™ application transform everyday PCs into powerful gaming machines.

G-SYNC

GEFORCE RTX

MAX-Q DESIGN





# HIGH-END PC GAMING FOR EVERYONE

One day, everyone will be a gamer. NVIDIA GeForce NOW™ is a gaming PC in the cloud that turns underpowered or incompatible computers – Macs and PCs alike – into powerful GeForce gaming rigs.

And through the GeForce NOW Alliance, global telcos like SoftBank and LGU+ are bringing the service to new markets around the world.

 SoftBank

 LGU+



# NVIDIA IS POWERING THE NEXT ERA OF COMPUTING

In 2006, the creation of our CUDA programming model and Tesla® GPU platform brought parallel processing to general-purpose computing. A powerful new approach to computing was born.

From the world's largest supercomputers to the vast data centers that power the cloud, this new computing model is helping to answer complex questions, discover new science, and bring amazing capabilities to our mobile devices.


Now, as the paths of high performance computing and AI innovation converge, data science is emerging as an essential tool for modern enterprises.



# POWERING THE WORLD'S FASTEST SUPERCOMPUTERS

GPU acceleration is the most accessible and energy-efficient path forward for the world's most powerful computers. More than 600 applications support CUDA today, including the top 15 in HPC.

NVIDIA powers U.S.-based Summit, the world's fastest supercomputer, as well as the fastest systems in Europe and Japan. 27,000 NVIDIA Volta Tensor Core GPUs accelerate Summit's performance to more than 200 petaflops for HPC and 3 exaops for AI.






# AI IS AUTOMATING THE WORLD

The demand for AI is all around us. In 2018, e-commerce surged to account for \$2.8 trillion of global retail sales. There are more than a billion automotive drivers in the world, driving a trillion miles a year. Half of the world has fewer than one doctor per 1,000 people.


AI can solve these challenges, spurring a wave of social progress unmatched since the Industrial Revolution.



TRANSPORTATION



SCIENTIFIC RESEARCH



HEALTHCARE



INTERNET / E-TAIL

# NVIDIA CUDA-X IS THE ENGINE OF AI

CUDA-X AI addresses the end-to-end development of AI. From data processing, to AI model training, to model deployment in the cloud.

NVIDIA is optimized for every AI model and every AI framework. We compute all forms of AI from machine learning to deep learning to reinforcement learning. The extensive software suite is available for free in the NVIDIA GPU Cloud (NGC) software hub.

CUDA-X AI is available in every cloud and from every computer maker.

## NVIDIA PARTNER ECOSYSTEM



## NVIDIA GPU CLOUD (NGC)

DATA ANALYTICS


GRAPH

MACHINE  
LEARNING


DL TRAINING

DL INFERENCE

## NVIDIA CUDA-X AI



HGX



DGX



DRIVE

# WORLD'S MOST POWERFUL AI TRAINING TOOL

Building amazing AI applications begins with training neural networks. NVIDIA DGX-2™ is the world's most powerful tool for AI training, uniting 16 GPUs to deliver 2 petaflops of training performance.

In December 2018, DGX-2 set six world records in the debut of MLPerf, a new set of industry benchmarks designed to test deep learning performance.

Image Classification

70 mins

Object Detection

176 mins

Object Instance Segmentation

14 mins

Translation (Recurrent)


10 mins

Translation (Non-Recurrent)

19 mins

Recommendation

0.4 mins





# POWERING AI INFERENCE FOR HYPERSCALE

Trained AI applications are deployed in large-scale, highly complex cloud data centers that serve voice, video, image, and recommendation services to billions of users. Hundreds of AI algorithms are in use today, making inference a big and costly challenge.

NVIDIA TensorRT™ software and the new T4 GPU converge to optimize, validate, and accelerate trained neural networks.



# ACCELERATING DATA SCIENCE

Data science is the fastest-growing area in computing. As it becomes a universal toolkit for businesses in every industry, data analytics and machine learning are tapping into the GPU.

NVIDIA RAPIDS™ is a platform for GPU-accelerated data science: analytics, machine learning, and, soon, data visualization.

NVIDIA powers a full range of data science computing, from hyperscale to the desktop. Data scientists can use the new NVIDIA Data Science Workstation to speed their predictive models and transform their businesses.



*Visualization: Fannie Mae loan portfolio risk  
based on US mortgage data 2000-2016*




# NVIDIA CLARA AI — A MEDICAL IMAGING SUPERCOMPUTER

The latest breakthroughs of AI and computational imaging are giving radiologists incredibly powerful tools for early detection and treatment.

The NVIDIA Clara AI Toolkit supercharges existing instruments with state-of-the-art image reconstruction, object detection and segmentation, and visualization capabilities.


The American College of Radiology chose Clara AI to empower its 38,000 members to build, share, and validate AI algorithms in thousands of hospitals in the U.S.

## TRADITIONAL MEDICAL IMAGING



## ENHANCED WITH NVIDIA CLARA AI

IMAGING AND VISUALIZATION APPS  
CUDA | CUDNN | TENSORRT | OGL | RTX  
GPU CONTAINERS | VGPU  
NVIDIA GPU SERVER

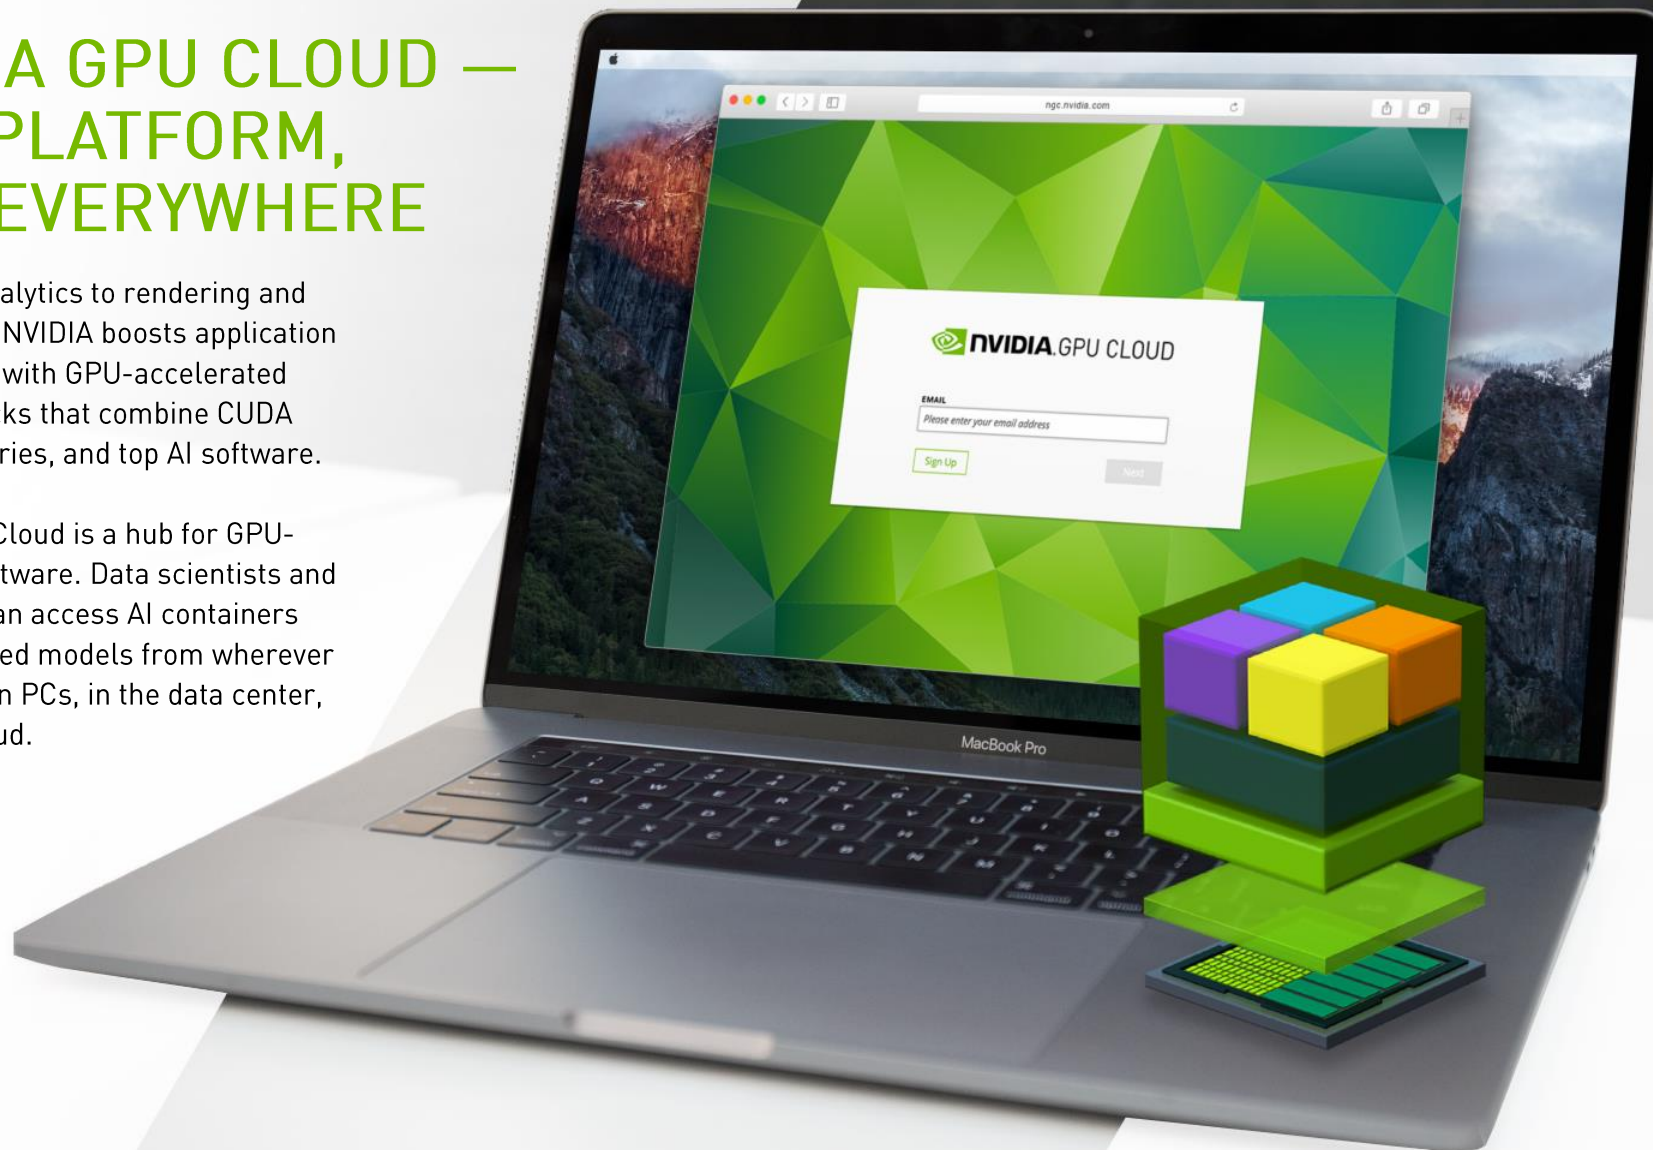




# NVIDIA GPU CLOUD — ONE PLATFORM, RUN EVERYWHERE

From data analytics to rendering and visualization, NVIDIA boosts application performance with GPU-accelerated software stacks that combine CUDA toolkits, libraries, and top AI software.

NVIDIA GPU Cloud is a hub for GPU-optimized software. Data scientists and developers can access AI containers and pre-trained models from wherever they want – on PCs, in the data center, or via the cloud.



Science



Medical Imaging



Deep Learning



Machine Learning



Rendering and  
Visualization



Hyperscale Inference



# AUTONOMOUS MACHINES ARE REVOLUTIONIZING INDUSTRIES


AI breakthroughs no longer come from scientific labs and hyperscale cloud providers alone. The era of robotics has arrived and is changing how people travel, the way products are made, and how business is done in trillion-dollar industries.

Self-driving cars, automated farm equipment, and autonomous factory robots have moved quickly from ideas to reality. And it's only the beginning. The fourth industrial revolution has begun.

# XAVIER — THE ENGINE OF AUTONOMOUS MACHINES

Xavier™, the first chip ever designed for autonomous AI systems, will enable automation in many large industries.

With 9 billion transistors, Xavier is the most complex SoC ever made. After more than 8,000 engineering years in development, Xavier is automotive-grade and in full production.





# NVIDIA DRIVE FOR AUTONOMOUS VEHICLES

Autonomous vehicles will revolutionize the \$10 trillion transportation industry.

NVIDIA DRIVE is an open platform and enables researchers and programmers to develop new algorithms or adapt them for specific vehicles.

To train the network, data from all over the world needs to be collected and fed into an NVIDIA DGX supercomputer.

Simulation expands the training set and covers dangerous scenarios that can't be captured on the road. The trained model is deployed on an in-car supercomputer, for capabilities like pedestrian detection and driver monitoring.



# CREATING THE FUTURE OF TRANSPORTATION

Hundreds of automakers, mobility services, truck makers, tier ones, mapping, and sensor companies are developing on NVIDIA DRIVE to create the future of transportation.

Daimler and Bosch are developing robotaxi fleets with DRIVE. Volkswagen is developing AI-infused cockpit experiences on DRIVE. And safety agencies like TÜV SÜD are using DRIVE Constellation to formulate their self-driving validation standards.



## AUTOMAKERS



Mercedes-Benz

TOYOTA

SUBARU



## MOBILITY SERVICES



NTierIV

UBER

Z E N U I T Y

## TRUCK MAKERS

ISUZU



SCANIA



## TIER ONES

Autoliv



BOSCH

Continental



DENSO



## MAPPING COMPANIES

Baidu 百度



TOMTOM



ZENRIN

## SENSOR COMPANIES

OMRON

Pioneer



Panasonic

SONY

# NVIDIA AND TOYOTA – A MODEL FOR AUTONOMOUS VEHICLE COLLABORATION

For decades, auto industry's R&D investments focused on engine design and materials fabrication. Now leading automakers are shifting investments toward the software-defined future of autonomous driving.

Toyota recently announced a new collaboration with NVIDIA to develop, train, and validate a single architecture to scale across its full line of next-generation autonomous vehicles.





# NVIDIA ISAAC — WHERE ROBOTS GO TO LEARN

Eventually, all machines that move will be autonomous, creating new opportunities in every industry from manufacturing and agriculture to consumer products and retail.

NVIDIA Isaac is a platform to accelerate the development and deployment of robotics. Isaac Sim is a virtual reality simulator where roboticists can create and train robots. Drop the software created in Isaac Sim into a robot with the Isaac SDK, and an intelligent machine is born.




*The "kitchen manipulator" is one of the robots in training at NVIDIA's AI Robotics Research Lab in Seattle.*

# JETSON AGX AND ISAAC DELIVER AI TO ROBOTICS AND IOT INDUSTRY

Jetson™ AGX Xavier delivers the energy-efficient computational power needed for embedded systems like robots, drones, and smart cities. And the new Jetson Nano™ will enable millions more small, low-power AI systems for embedded IoT apps.

From Xavier to Nano, all of NVIDIA's AI computers run on the same CUDA-X AI software stack.



The background is a dark, almost black, space filled with a network of thin, light green lines. These lines connect various points, some of which are highlighted as small, glowing green circles. The overall effect is reminiscent of a digital network or a molecular structure. In the upper left corner, there is a thin, light green L-shaped line that frames the top-left corner of the text area.

# NVIDIA CULTURE

NVIDIA is united by a unique culture — the operating system of our company. We dream big, take risks, and learn from our mistakes together. Speed is the key to our success. Craftsmanship is a passion. There are no org charts — the project is the boss.

These beliefs inform everything we do, from designing amazing products to building one of the world's great companies — a place where people can do their life's work.



# NVIDIA — A LEARNING MACHINE

NVIDIA has continuously reinvented itself over more than two decades.

Our invention of the GPU in 1999 sparked the growth of the PC gaming market, redefined modern computer graphics, and revolutionized parallel computing. More recently, GPU computing ignited the era of AI.

NVIDIA is a “learning machine” that constantly evolves by adapting to new opportunities that are hard to solve, that only we can tackle, and that matter to the world.







## INSPIRED TO GIVE TO OUR COMMUNITIES

NVIDIA's people share a strong sense of corporate responsibility. Our philanthropic giving exceeded \$5.3 million in 2018. To date, our NVIDIA Foundation's Compute the Cure initiative has directed more than \$5 million to the fight against cancer. And our Techsplorer program, which introduces underserved youth to AI, has reached more than 10,000 students.





“World’s Best  
Performing CEOs”

— Harvard Business Review

“100 Best Companies  
to Work For”

— Fortune

“World’s Best CEOs”

— Barron’s

“Most Innovative  
Companies”

— Fast Company

“Employees’ Choice:  
Highest Rated CEOs”

— Glassdoor

“50 Smartest  
Companies”

— MIT Tech Review

Founded in 1993 | Jensen Huang, Founder & CEO | 13,000 employees | \$11.7B in FY19



