

INVESTOR RELATIONS NEWSLETTER



Welcome to the NVIDIA Quarterly Investor Newsletter

This is the inaugural edition of our new quarterly investor newsletter. We like to say at NVIDIA that we move “at the speed of light.” But in these pages, our goal is to pause and reflect on each quarter’s key developments and milestones.

GTC 2018 and Investor Day: Just 10 days away

A computing platform is only as strong as its developer ecosystem, as students of tech history from Wintel to iOS know well. As such, our annual GPU Technology Conference (GTC) is integral to our strategy of building the leading accelerated computing platform for the era of AI. We’re kicking off our [GTC 2018](#) global tour on March 26-29 in San Jose. This year’s event is expected to be the biggest yet, with more than 8,000 attendees – up 3.5X in 5 years. GTC features hundreds of sessions on the most important topics in computing today, including AI/deep learning, HPC, virtual reality, self-driving cars, smart cities, and much more.

NVIDIA CEO Jensen Huang will deliver the opening keynote at 9am on March 27, followed by an investor-only track featuring in-depth presentations by our executives. If you can’t join in person, tune in to the [Investor Day](#) webcast to gain insights into NVIDIA’s growth strategies and how we’re leveraging our investments to build our businesses.

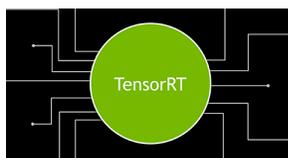
Strong 4Q results capped a great year

We had an outstanding fourth quarter and fiscal 2018, led by strong growth in our gaming and data center businesses. Q4 revenue reached \$2.91 billion, up 34% year on year, and all measures of profitability set records. Our gaming business was up 29% y/y to \$1.74 billion, and our datacenter business grew 105% y/y to \$606 million. We highlight a handful of product and customer examples below to help illustrate some of the drivers of that growth.



All major cloud providers are ramping up AI workloads on NVIDIA Tesla V100 GPUs

Our latest data center GPU, the Tesla V100 built on the Volta architecture, is off to a great start. Volta, [introduced](#) in May 2017, offers 5x the deep learning performance of its predecessor, Pascal, or the equivalent of 100 CPUs. It also adds mixed-precision computation formats (i.e. floating point 16, 32 and 64 in addition to 8-bit integer), which means it can handle inference in addition to training. All major cloud providers have announced plans to adopt Volta, including [AWS](#) (available since Oct. 2017), [IBM Cloud](#) (available since Jan. 2018), [Microsoft Azure](#) (available since Mar. 2018). Other CSPs on the way include Alibaba Cloud, Baidu Cloud, Google Cloud Platform, Oracle Cloud and Tencent Cloud.



NVIDIA’s GPU platform is getting traction in AI inference, a new opportunity for GPUs

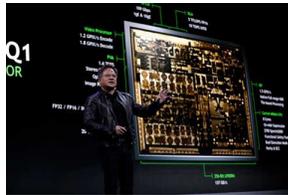
While our leadership in AI training is well established, AI inference is a new multi-billion market opportunity for our GPUs. In addressing this opportunity, the most important development was the introduction last fall of our Tensor RT 3 programmable inference accelerator. With this sophisticated software, NVIDIA’s AI inference platform can improve performance and efficiency by orders of magnitude over CPUs. We have seen increased adoption of our AI inference platform and are working with over 1,000 companies across a wide spectrum of industries, such as Pinterest, Twitter, KLM, Tencent and JD.com, to discover

new insights from data and deploy intelligent services to businesses and consumers.



Increasingly social gameplay attracts more players and requires higher-end GPUs

PC gaming is becoming increasingly social. Look no further than the success of PUBG and Fortnite, games that have attracted over 80 million players combined in just a few months in a new "Battle Royale" genre, where 100 players square off in a last-man-standing competition. Or witness the rise of eSports – multiplayer videogame competitions that are on track to amass close to 600 million spectators by 2020, up almost 2x from 2016. This sets in motion a virtuous cycle: more social gameplay expands the universe of gamers, encouraging the development of more and better games, which in turn attracts more gamers. This dynamic not only drives increasing GPU demand, but also a favorable mix shift, as higher-production value games tend to require higher-performance GPUs.



CES 2018: Xavier processor ready for auto-makers to start building self-driving cars

Jensen kicked off CES, announcing that the first samples of our Xavier processors are being delivered to customers this quarter, with production planned for later this year. Xavier is the world's first autonomous machine processor, enabling Level 2 to 4 autonomous vehicles. We believe Xavier will be in production two years ahead of the competition. We also announced automotive customer wins with Aurora, VW, and Uber, and a partnership with Baidu and ZF for autonomous vehicles in China, which makes up 30% of the global passenger vehicle market.

Notable Links

Microsoft blog: [Microsoft Azure announces general availability of NVIDIA Tesla V100](#)
The Next Platform: [The Engine of HPC and Machine Learning](#)
NVIDIA blog: [NVIDIA Captures Pole Position at CES Amid Flurry of Auto News](#)
NVIDIA blog: [PC Gaming: More Than 1 Billion Players ... and Counting](#)
Fortune: [Jensen Huang: Business Person of the Year](#)

Legal Notices

FOLLOW US



[Unsubscribe](#)

© 2018 NVIDIA Corporation. All rights reserved.
NVIDIA Corporation, 2788 San Tomas Expressway, Santa Clara, CA 95051.