

## Customer Success Story

Education  
Cloud computing

# Intel and AWS help educators launch online learning platform

**In the face of the pandemic lockdown, Career Launcher successfully delivers learning on a massive scale in India.**

*"It is indeed great to collaborate with tech giants such as Intel and AWS to help education stay uninterrupted in this period of crisis. We hope to take this partnership to a hundred such large projects globally and to be seen as practitioners of the do-good-do-well philosophy."*

—Satya Narayanan R, Chairman  
CL Educate

*"AWS is pleased to have worked closely with Career Launcher and Intel to build the aspiration.ai learning platform, which is focused on delivering learning, fun, and mentoring at scale. We are glad to be part of this effort to enable continued education for students in government schools in Delhi."*

—Rahul Sharma, President,  
World Wide Public Sector, India and  
South Asia, Amazon Internet Services  
Pvt. Ltd.

### Executive Summary

The COVID-19 global pandemic and subsequent lockdown left students across India unable to attend school in person. The Directorate of Education of National Capital Territory (NCT) Delhi needed a fast-scaling solution to not only ensure educational continuity but to also serve as a platform for long-term educational transformation. It engaged Intel's partner, Career Launcher, and soon a partnership was formed with Intel and [Amazon Web Services \(AWS\)](#). Together, they worked with educators to adapt aspiration.ai, Career Launcher's AWS-based learning portal, to address the needs of Delhi's schools. Within two weeks the trial was successfully launched, and Project Aspiration 2020 was born. Within two months it was scaled to serve over 160,000 students with powerful learning tools and the potential to accommodate many more.

### Meeting an immediate need to bring education home

In March 2020, the COVID-19 pandemic hit India and led to a nationwide lockdown that closed all government schools. With little warning, hundreds of thousands of students in the NCT of Delhi were left without a school to attend. The Directorate of Education of NCT Delhi needed to mount a swift intervention to limit the interruption for students. More than that, they wanted to launch a solution that would help them transform education in Delhi and India as a whole.

With a relatively high degree of broadband internet access compared with other parts of India, the city of Delhi was an ideal candidate to trial online learning for public schools, but the typical time for developing and deploying such a solution was not available. The government's educational leaders contacted Career Launcher, an Intel educational technology partner to launch a distance learning solution that could scale fast and had the potential to deliver value for students for years to come. To quickly acquire the necessary resources, Career Launcher enlisted the help of Intel and Amazon Web Services (AWS), and the initiative known as Project Aspiration 2020 began.

With offices around the world, Career Launcher partners with educational institutions and employs AWS resources to provide robust distance learning through its online portal, aspiration.ai. However, this portal was designed for the needs of adults seeking continuing education and required modification to fit the objectives and learning modes of K-12 school students. What's more, expanding from Aspiration.ai's existing audience of around 140,000 learners to create virtual classrooms for the targeted group of K-12 students meant increasing their daily active users by 60 percent. It was a monumental scale-up and they had to do it fast.



## Fast-moving collaboration, fast-scaling platform

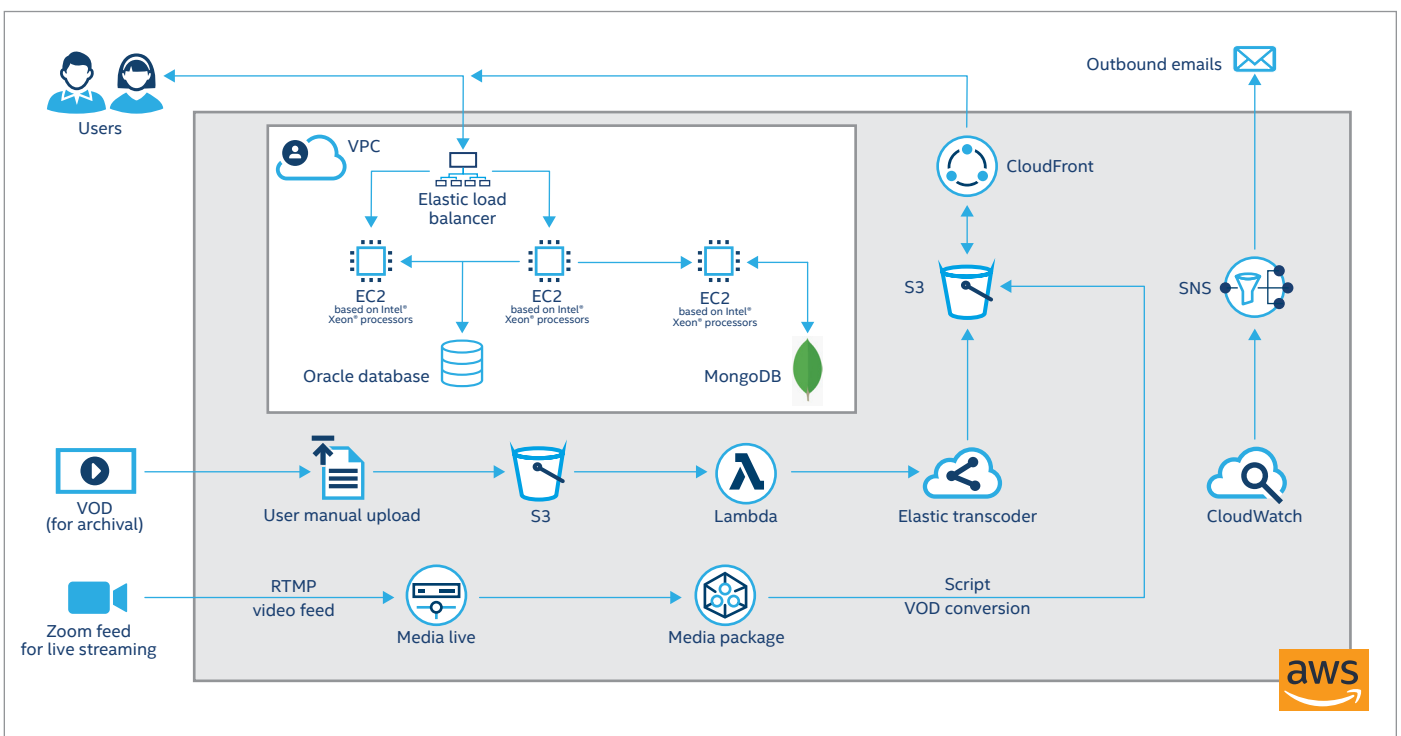
Intel and AWS realized the impact Project Aspiration 2020 could have and quickly created a nimble team to facilitate the implementation. The two companies provided funding for a trial and worked with Career Launcher to configure a solution that overcame their challenges and met their needs. The AWS team worked closely with Career Launcher to create a flexible and secure Amazon Virtual Private Cloud (Amazon VPC) for Project Aspiration 2020. This solution brought together the right mix of resources to establish the web portal, deliver the learning content, and manage their database. Within a week of launching Project Aspiration 2020, the revised solution was up and running.

There were many considerations that affected the configuration—the first was cost. The government school system, under financial pressure from the pandemic, needed a cost-effective solution. Amazon Elastic Compute Cloud (Amazon EC2) C4 instances, powered by Intel® Xeon® processors, helped keep costs low while still delivering the fast-scaling compute necessary for advanced features and a responsive experience. Amazon EC2 R4, also featuring Intel Xeon processors, gave them the memory capacity required for data layer functions like machine learning-based predictive analytics on their MongoDB and Oracle databases.

## Overcoming content challenges

The next challenge was content delivery. Career Launcher had previously livestreamed video from their studio. With the lockdown, that was no longer an option. They needed a content delivery network (CDN) that could support the livestreaming of 30 to 40 video classes per day, each of which can be up to three hours long and have up to 7,500 concurrent users. The solution needed to be low latency to give students and teachers the chance to meaningfully interact. Amazon CloudFront gave them a fast and secure CDN that can add or remove resources in response to changes in demand. Live video of classes is encoded and compressed in real time with AWS Elemental MediaLive. This service works in conjunction with AWS Elemental MediaPackage to prepare the video for delivery to many different types of screens over the internet. MediaPackage also enables DVR-like features that are very helpful to students, such as the ability to pause and rewind. It then goes out on Amazon CloudFront to the end user via the portal.

Another key concern with content delivery was access for students without the reliable broadband internet required for livestreaming lessons. Career Launcher developed a separate AWS-based video pipeline to take video of classes and upload them to Amazon Simple Storage Service (Amazon S3). Once there, an AWS Lambda function sends the video to Amazon Elastic Transcoder, which converts it to multiple formats—including those suited for low bandwidth and low-resolution playback devices. Offering both synchronous and asynchronous learning content in many sizes allows students to consume learning content based on their available resources. This library of archival video has now surpassed 7,000 lessons that are available whenever a student needs them.



Project Aspiration 2020's AWS solution architecture

## Driving success with analytics

Beyond delivering lessons that allow for a virtual face-to-face connection between teacher and student, the advanced features of AWS enable analytics that improve the program and outcomes for students. The solution tracks key indicators of student success and engagement in order to make sure that everyone's needs are being addressed. Career Launcher also employs machine learning to extract learnings that can predict student performance. By analyzing how students do on practice exams, they can predict how they will score on the high-pressure public exams that will affect their future. Career Launcher can analyze a student's performance and suggest an academic path that best suits their needs. These predictive analytics also help guide educators on the content team. They can predict whether students will be able to answer certain questions while writing exams.

Analytics are also employed to assess the performance of the program among stakeholders. Using Metabase and Mongo charts—which are integrated into their Amazon VPC—Career Launcher provides metrics to administrators, government officials, and educators. This gives the coalition of educational and technical partners the ability to assess what's working and what's not. They can share these data visualizations across their organizations to ensure that the best insights are brought to light. That way, the service can be evolved to continue to deliver the best learning experience possible for the students.

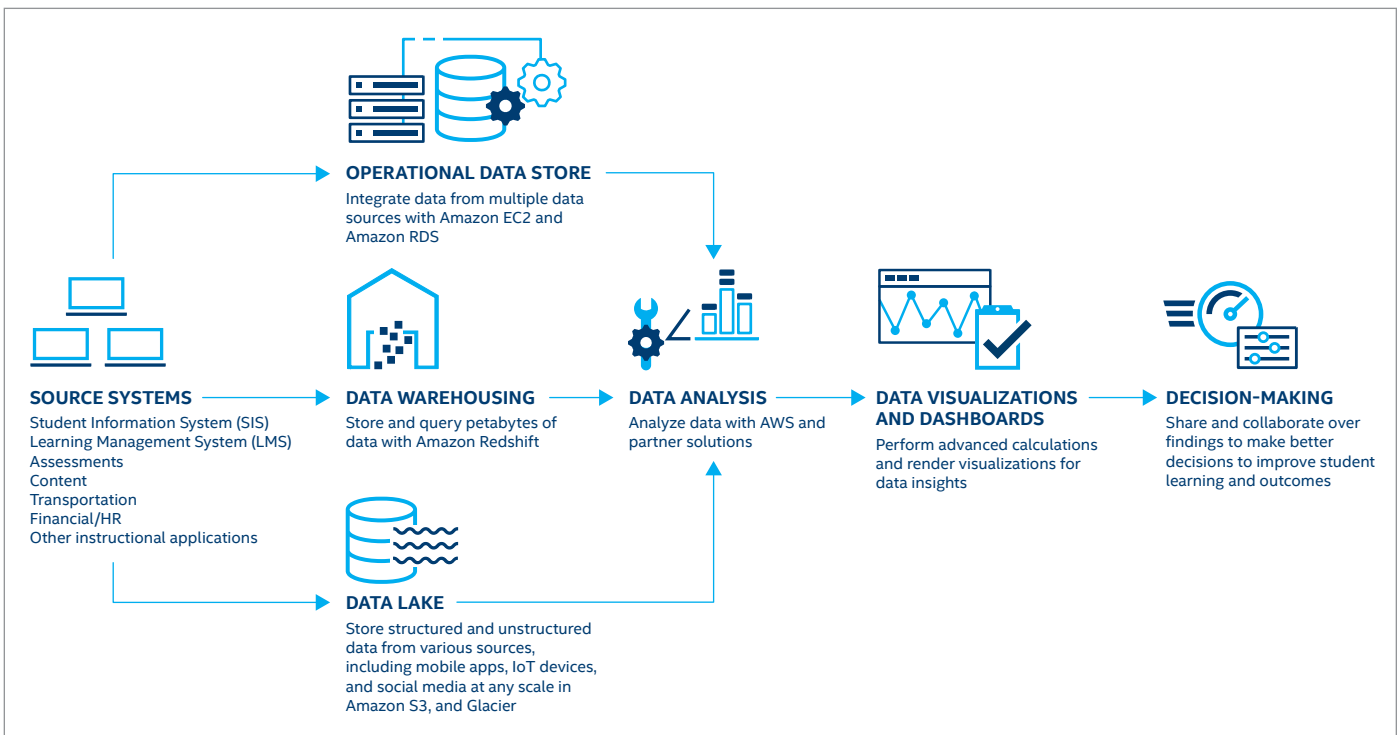
Going forward, Career Launcher is considering a move to Amazon EC2 C5 and R5 instances to gain the advantages of 2nd Generation Intel® Xeon® Scalable processors. This would greatly improve the performance of their machine learning algorithms with built-in artificial intelligence (AI) acceleration, such as Intel AVX-512.

## A powerful platform for learning

The scalable power of AWS allowed Career Launcher to offer special interactive features to keep students learning, despite the less than ideal circumstances. Teachers now have new tools to promote learning and make learning more engaging and fun. They can create learning games to make the material fun and keep students motivated during this difficult time. Teachers are also able to create quizzes to drive home the subject matter and incorporate a wider variety of media into their lessons. They are now able to show video clips to make history lessons real and relevant and play animations in biology classes that make it easier to explore complex topics.

Project Aspiration 2020 is enabling new ways for teachers to engage with students. For large classes, a team of teachers assists the main instructor by interacting with students in the chat to help them resolve questions. Students have also been inspired by the format to participate in new ways. Teachers report that students who rarely spoke up in class are active and engaged in the chat. One student, Lokesh Jha, commented about the experience, "The classes have become more interesting with graphics and animations. Quizzes galore. All of us wait for our names to be taken by the teachers when we answer the question." Jyotsna Davar, an accounting and business instructor noted, "All my students are terrifically enjoying these classes. They look forward to these classes."

The platform also enables the schools to offer mentoring and career guidance to both students and parents. This allows educators to enhance the quality of education student receive by bringing parents into the equation while also providing helpful resources for their families.



Data management analytics solution

## Speedy training, onboarding, and launch

Meeting the needs of the students required more than computing resources and a powerful portal. It also meant training and onboarding hundreds of teachers and establishing a system for creating and managing courses. Implementation of Project Aspiration 2020 began with fast-track onboarding program for educators. Career Launcher worked with the Department of Education to identify a group of top teachers who would spearhead the transition to online learning. From this initial group of 300 teachers, 60 were selected to be the core instructors for key classes. Within one week, this select group were certified to begin conducting classes with thousands of live participants.

All those who participated in the onboarding have benefited, as well. Not only were they trained in managing aspiration. ai classes, they also received a course in online education. Many of these teachers came to the program with very limited computing skills and experience. Through the training, they were able to learn how to leverage media resources and keep children deeply engaged.

## More than a solution: a transformation

From the initial group of 55 schools and 1,500 students, the program is now serving over 165,000 students with high rates of engagement: 90 percent of students enrolled are regularly active. With requests to participate in the program coming in from across the country, the project is planned to grow to reach about one million students within the next month.


But the goal for the Career Launcher program stretches well beyond the COVID-19 crisis. The fast scaling that Career Launcher, Intel, and AWS have achieved is inspiring government officials to think differently about education. They are seeing the potential for this program to have a lasting impact on students' lives and to help lead a transformation in the education system across the country that gives more children access to the best learning experience possible.

The platform also allowed the Department of Education to increase the scale of education. The AWS Cloud resources have made them able to deliver education with smaller, nimbler footprint. The public education system is one of the biggest pieces of the government's budget. Ministers are very optimistic about the potential for Project Aspiration 2020 to make their spending much more efficient by reducing the need for capital expenses and physical infrastructure, while still improving the quality of education that students receive.

## Scalable and powerful: Intel in AWS

Intel architecture is at the heart of Amazon EC2 cloud computing instances, providing strong data protection, fast processing of large data volumes, and service flexibility without a hit to performance. Intel processors feature:

- Intel® Advanced Vector Extension 512 (Intel® AVX-512), which offers accelerated application performance up to 2x better than previous-generation technologies<sup>1</sup>, enabling significant improvements in speed for data application workloads.
- Intel® Trusted Execution Technology (Intel® TXT), which remains Intel's technology for establishing more secure platforms. One-Touch Activation brings an added level of protection for geographic needs like regional and county-specific data-sovereignty regulations.
- Intel® Deep Learning Boost (Intel® DL Boost) offers built-in AI acceleration, with up to 14x better inference<sup>2</sup> performance on image classification in 2nd Generation Intel Xeon Scalable processors, compared to competing processors.
- Intel® Turbo Boost Technology accelerates processor and graphics performance for peak loads, automatically allowing processor cores to run faster than the rated operating frequency if they're operating below power, current, and temperature specification limits.
- Intel® AES New Instructions (Intel® AES-NI) improves upon the original Advanced Encryption Standard (AES) algorithm to provide faster data protection and greater security. All current-generation Amazon EC2 instances support this processor feature.



## Learn more

Explore Intel's global response to the COVID-19 crisis here:

[intel.com/content/www/us/en/corporate-responsibility/covid-19-response](https://intel.com/content/www/us/en/corporate-responsibility/covid-19-response) >

To see more of what Intel's partner, **Career Launcher**, is doing to deliver the future of education, please visit: [careerlauncher.com](https://careerlauncher.com) >

Learn more about the platform driving **Project Aspiration 2020** at: [aspiration.ai](https://aspiration.ai) >

See what **AWS** is doing to enable innovation for enterprises of all type and sizes:

[aws.amazon.com](https://aws.amazon.com) >



1. FSI Kernels- Baseline: Intel® Xeon® Platinum 8268 processor configuration: Intel "Wolf Pass" platform with 2-socket Intel® Xeon® Platinum 8268 processors (2.9GHz, 24C), 12x16GB DDR4-2933, 1 SSD, BIOS: SE5C620.86B.02.01.0008.031920191559; Microcode: 0x500001c, Red Hat Enterprise Linux\* 7.7, kernel 3.10.0-1062.1.1. FSI Kernels v2.0: Geomean (3 workloads: Binomial Options, Black Scholes, Monte Carlo), AVX2\_256 build, Intel® Compiler 2019u5, Intel® Math Kernel Library (Intel® MKL) 2019u5, BIOS: Binomial (HT ON, Turbo ON, SNC OFF, 2 threads/core), Black Scholes (HT OFF, Turbo ON, SNC OFF, 1 threads/core), Monte Carlo (HT ON, Turbo ON, SNC OFF, 2 threads/core). Test by Intel as of 11/1/2019, w/AVX-512: Intel® Xeon® Platinum 8268 processor configuration: Intel "Wolf Pass" platform with 2-socket Intel® Xeon® Platinum 8268 processors (2.9GHz, 24C), 12x16GB DDR4-2933, 1 SSD, BIOS: SE5C620.86B.02.01.0008.031920191559; Microcode: 0x500001c, Red Hat Enterprise Linux\* 7.7, kernel 3.10.0-1062.1.1. FSI Kernels v2.0: Geomean (3 workloads: Binomial Options, Black Scholes, Monte Carlo), AVX-512 build, Intel® Compiler 2019u5, Intel® Math Kernel Library (Intel® MKL) 2019u5, BIOS: Binomial (HT ON, Turbo ON, SNC OFF, 2 threads/core), Black Scholes (HT OFF, Turbo ON, SNC OFF, 1 threads/core), Monte Carlo (HT ON, Turbo ON, SNC OFF, 2 threads/core). Test by Intel as of 11/1/2019
  2. Up to 14X AI Performance Improvement with Intel® DL Boost compared to Intel® Xeon® Platinum 8180 Processor (July 2017). Tested by Intel as of 2/20/2019. 2 socket Intel® Xeon® Platinum 8280 Processor, 28 cores HT On Turbo ON Total Memory 384GB (12 slots/ 32GB/ 2933 MHz), BIOS: SE5C620.86B.0D.01.0271.120720180605 (ucode: 0x200004d), Ubuntu 18.04.1 LTS, kernel 4.15.0-45-generic, SSD 1x sda INTEL SSDSC2BA80 SSD 745.2GB, nvme1n1 INTEL SSDPE2KX040T7 SSD 3.7TB, Deep Learning Framework: Intel® Optimization for Caffe version: 1.1.3 (commit hash: 7010334f159da247db3fe3a9d96a3116ca06b09a), ICC version 18.0.1, MKL DNN version: v0.17 (commit hash: 830a10059a018cd2634d94195140cf2d8790a75a, model: [https://github.com/intel/caffe/blob/master/models/intel\\_optimized\\_models/int8/resnet50\\_int8\\_full\\_conv.prototxt](https://github.com/intel/caffe/blob/master/models/intel_optimized_models/int8/resnet50_int8_full_conv.prototxt), B5=64, DummyData, 4 instance/2 socket, Datatype: INT8 vs Tested by Intel as of July 11th 2017: 2S Intel® Xeon® Platinum 8180 CPU @ 2.50GHz (28 cores), HT disabled, turbo disabled, scaling governor set to "performance" via intel\_pstate driver, 384GB DDR4-2666 ECC RAM. CentOS Linux release 7.3.1611 (Core), Linux kernel 3.10.0-514.10.2.el7.x86\_64. SSD: Intel® SSD DC S3700 Series (800GB, 2.5in SATA 6Gb/s, 25nm, MLC).
- Performance measured with: Environment variables: KMP\_AFFINITY=granularity=fine,compact; OMP\_NUM\_THREADS=56, CPU Freq set with cpupower frequency-set -d 2.5G -u 3.8G -g performance. Caffe: (<http://github.com/intel/caffe/>), revision f96b759f71b2281835f690af267158b82b150b5c. Inference measured with "caffe time --forward\_only" command, training measured with "caffe time" command. For "ConvNet" topologies, dummy dataset was used. For other topologies, data was stored on local storage and cached in memory before training. Topology specs from [https://github.com/intel/caffe/tree/master/models/intel\\_optimized\\_models](https://github.com/intel/caffe/tree/master/models/intel_optimized_models) (ResNet-50), Intel C++ compiler ver. 17.0.2 20170213, Intel MKL small libraries version 2018.0.20170425. Caffe run with "numactl -l".

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others.