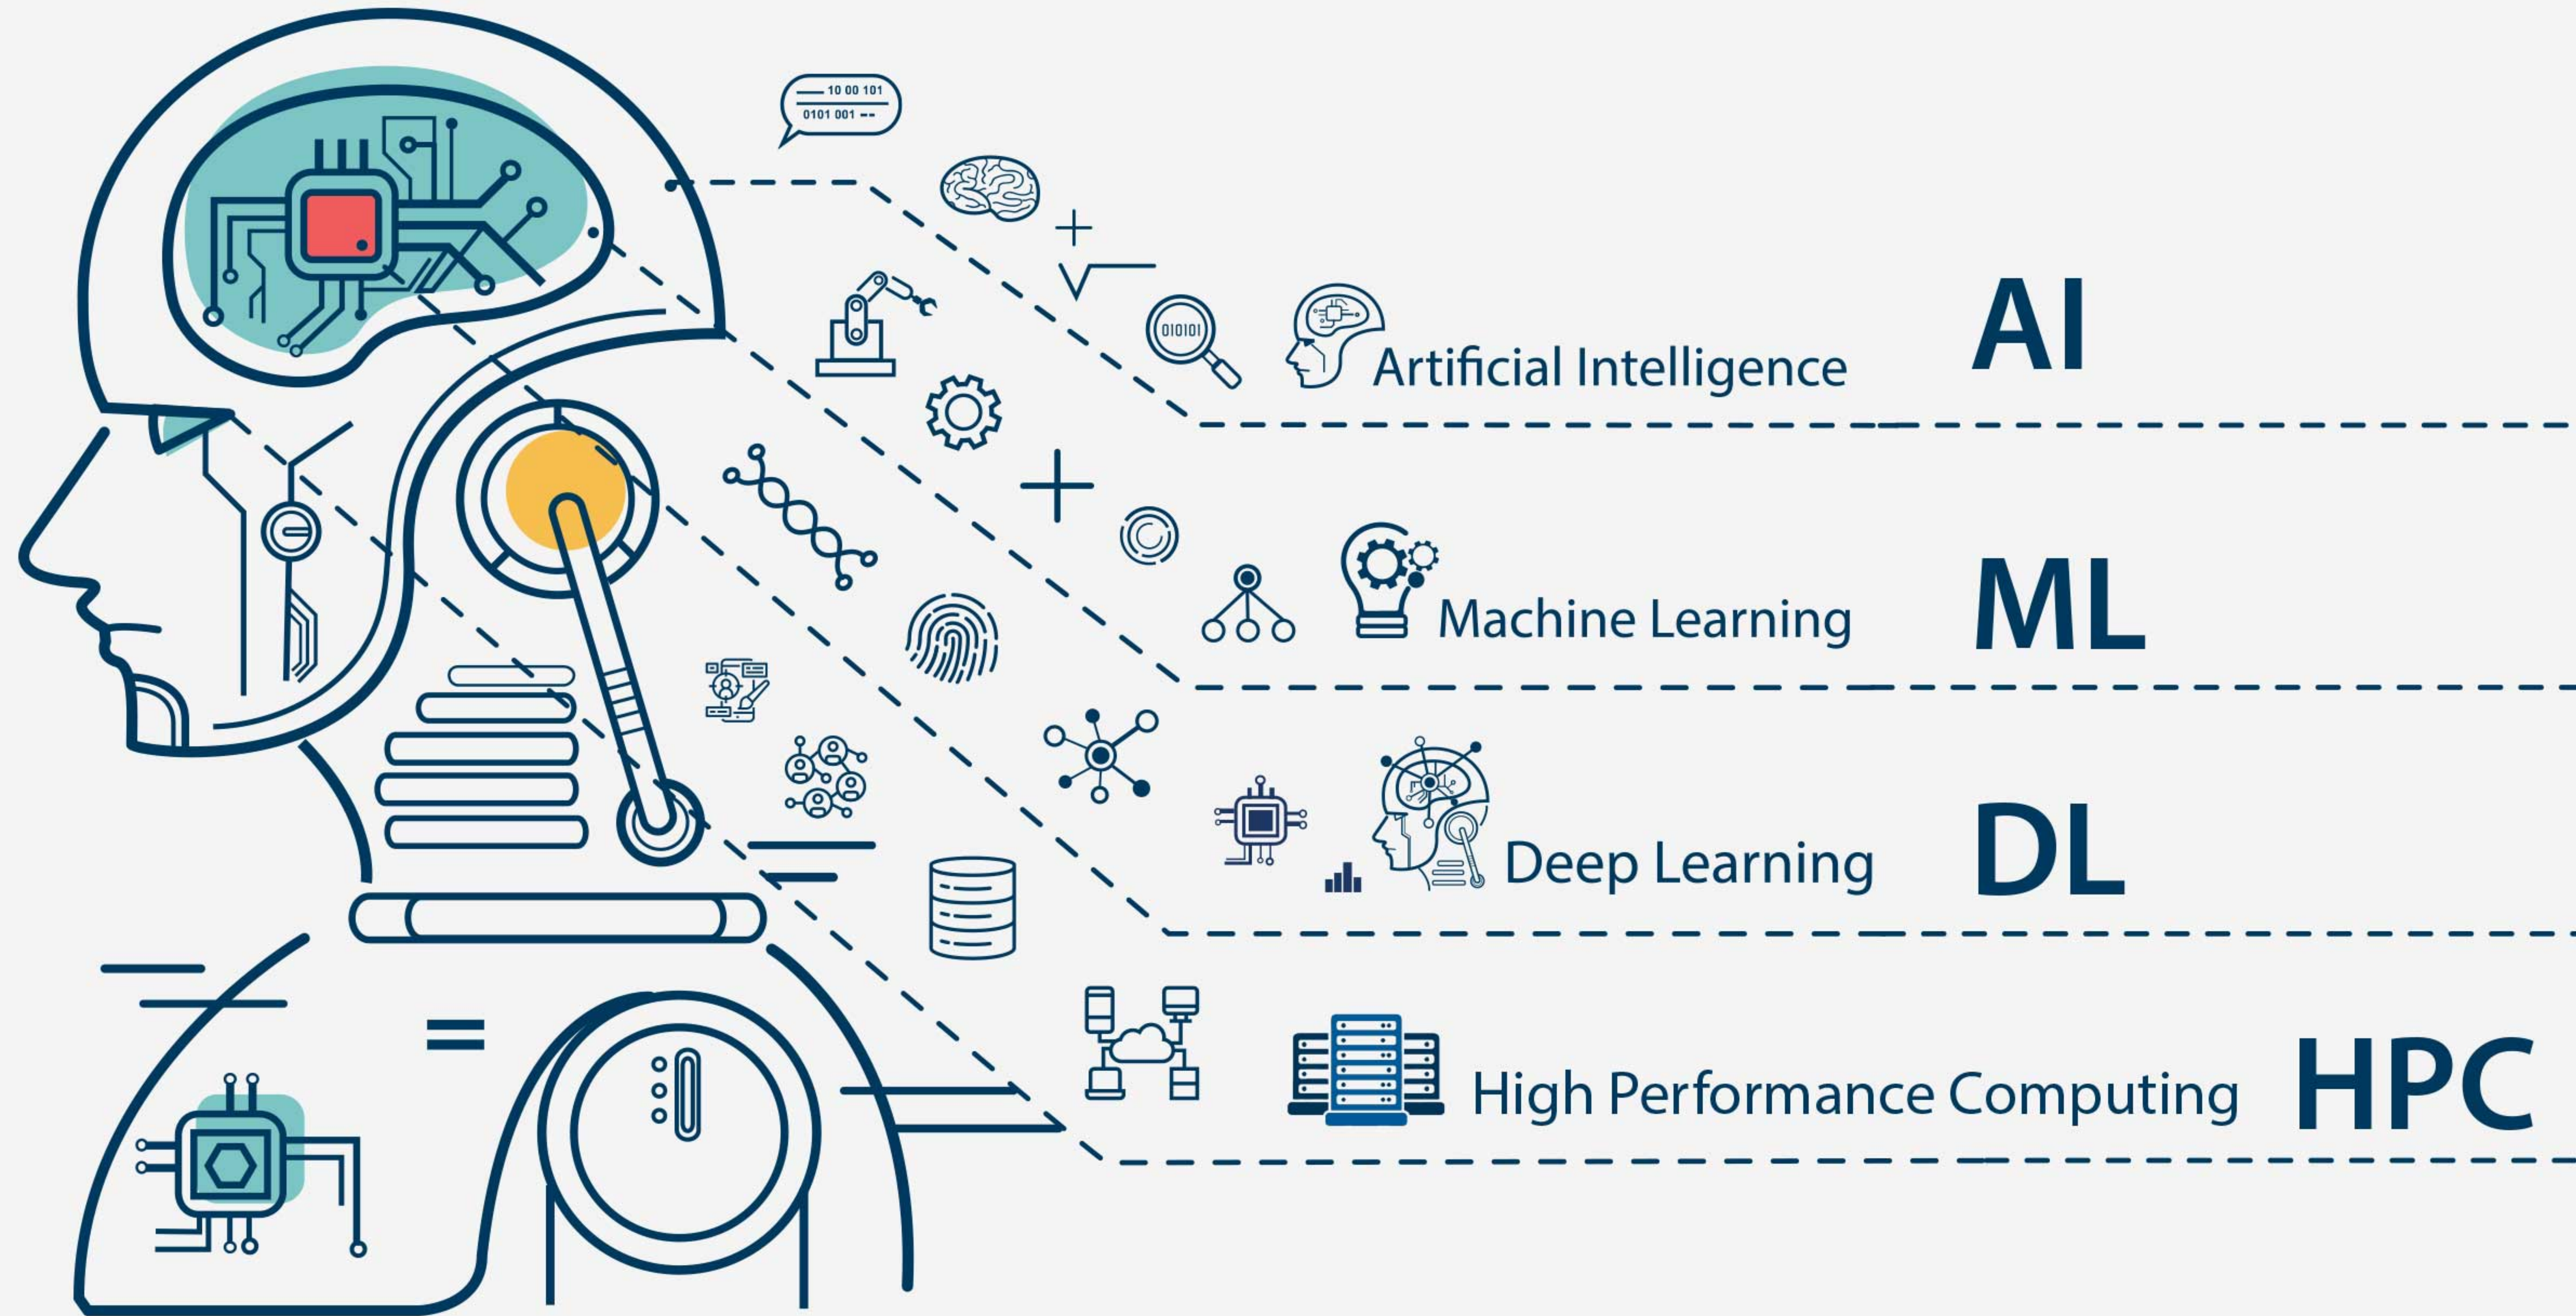



GPU Cloud

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant

Powered By:



State Of The Art GPU Cloud | Virtual + Physical Server Available.



Use Cases GPU Cloud

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant

Powered By:



Machine Learning/AI:



Natural language Processing



Image and Video Recognition



Autonomous Vehicle Systems



Recommendation Systems



Customer sentiment analysis



Predictive analysis using sensor data


High Performancing Compute:



Computational Fluid Dynamics



Financial & Data Analytics



Weather Simulation



Computational Chemistry



Creative and Technical



Knowledge Workers

Virtual Graphics:

Scientific Applications:



Modular Simulations



Seismic Mapping



Weather Forecasting



Real-time analytics



Motion Detection



Augmented Detection

Consumer Internet & Industry Applications:



Speech



Translate



Recommender




Healthcare




Manufacturing



Supply Chain



Time Series




NVIDIA DATA Center Platform

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



Single Platform Drives Utilization and Productivity

CUSTOMER USE CASES



CONSUMER INTERNET & INDUSTRY APPLICATIONS

SCIENTIFIC APPLICATIONS

VIRTUAL GRAPHICS

APPS & FRAMEWORKS



CUDA-X & NVIDIA SDKs

MACHINE LEARNING: cuDF, cuML, cuGRAPH

DEEP LEARNING: cuDNN, CUTLASS, TensorRT


HPC: OpenACC, cuFFT

VIRTUAL GPU: vDWS, vPC, vAPPS

CUDA & CORE LIBRARIES - cuBLAS | NCCL

TESLA GPUs & SYSTEMS





NVIDIA Tesla V100

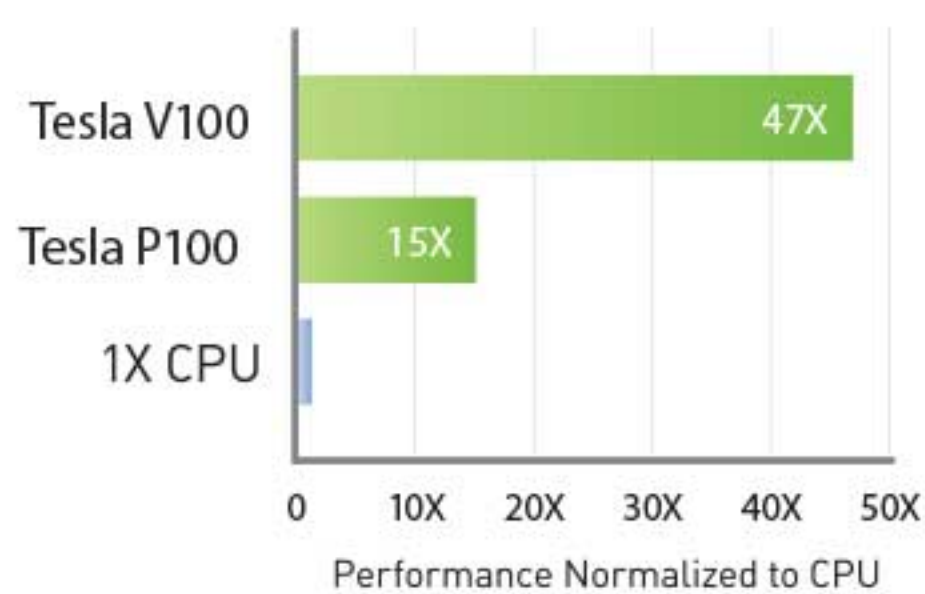
Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



THE MOST ADVANCED DATA CENTER GPU EVER BUILT.


NVIDIA® Tesla® V100 is the world's most advanced data center GPU ever built to accelerate AI, HPC, & graphics. Powered by NVIDIA Volta, the latest GPU architecture, Tesla V100 offers the performance of up to 100 CPUs in a single GPU—enabling data scientists, researchers, and engineers to tackle challenges that were once thought impossible.

47X Higher Throughput than CPU Server on Deep Learning Inference



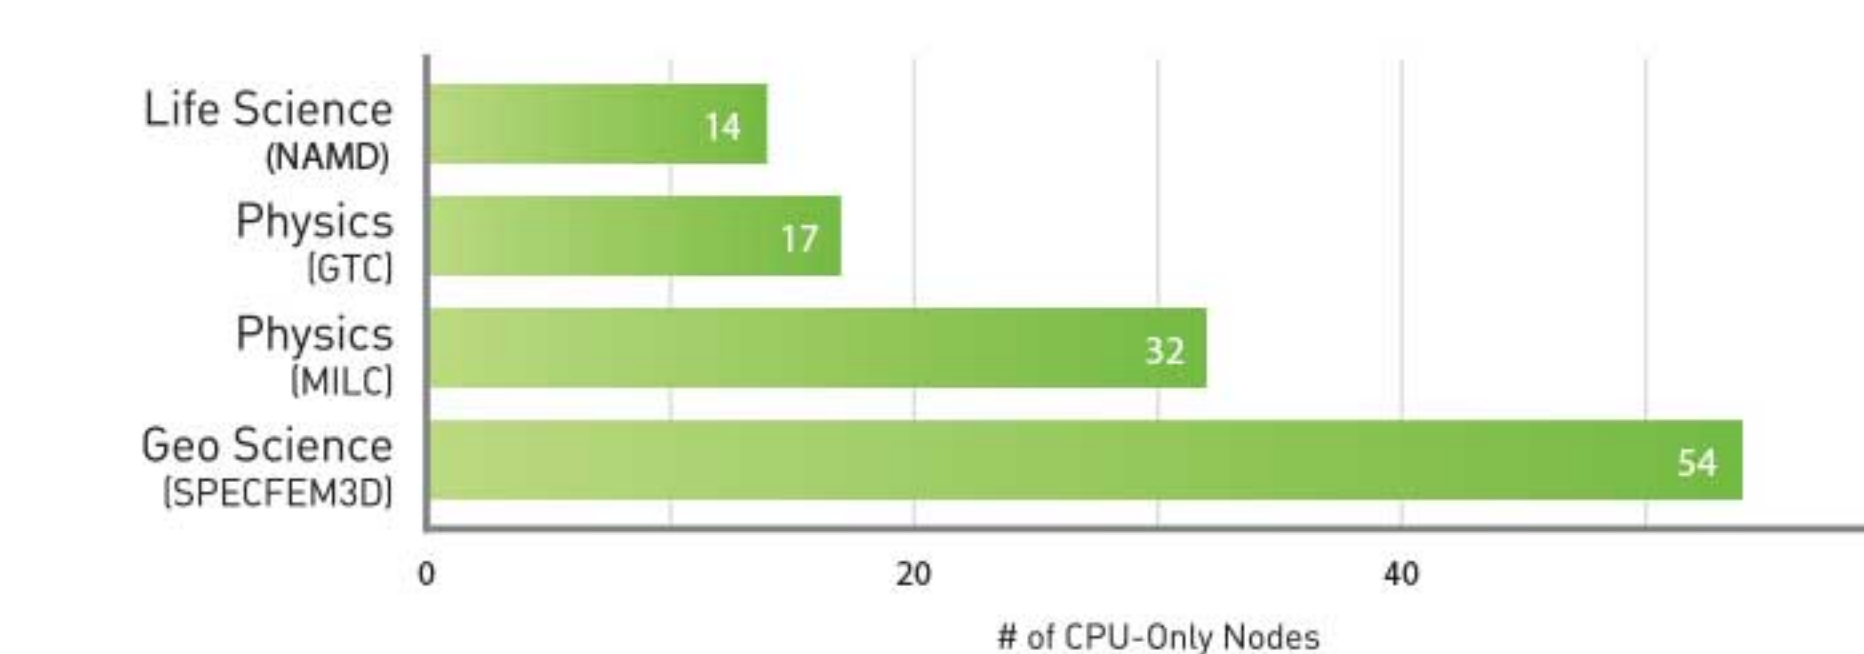
Workload: ResNet-50 | CPU: 1X Xeon E5-2690v4 @ 2.6GHz | GPU: add 1X NVIDIA* Tesla* P100 or V100

Deep Learning Training in Less Than a Workday



Server Config: Dual Xeon E5-2699 v4 2.6GHz | 8X Tesla P100 or V100 | ResNet-50 Training on MXNet for 90 Epochs with 1.28M ImageNet dataset.

1 GPU Node Replaces Up To 54 CPU Nodes



CPU Server: Dual Xeon Gold 6140@2.30GHz, GPU Servers: same CPU server w/ 4x V100 PCIe | CUDA Version: CUDA 9.x | Dataset: NAMD (STMV), GTC (mpi#proc.in), MILC (APEX Medium), SPECFEM3D (four_material_simple_model) | To arrive at CPU node equivalence, we use measured benchmark with up to 8 CPU nodes. Then we use linear scaling to scale beyond 8 nodes.

SPECIFICATIONS



Tesla V100 PCIe



Tesla V100 SXM2


GPU Architecture	NVIDIA Volta	
NVIDIA Tensor Cores	640	
NVIDIA CUDA * Cores	5,120	
Double-Precision Performance	7 TFLOPS	7.8 TFLOPS
Single-Precision Performance	14 TFLOPS	15.7 TFLOPS
Tensor Performance	112 TFLOPS	125 TFLOPS
GPU Memory	32GB /16GB HBM2	
Memory Bandwidth	900GB/sec	
ECC	Yes	
Interconnect Bandwidth	32GB/sec	300GB/sec
System Interface	PCIe Gen3	NVIDIA NVLink
Form Factor	PCIe Full Height/Length	SXM2
Max Power Consumption	250 W	300 W
Thermal Solution	Passive	
Compute APIs	CUDA, DirectCompute, OpenCL™, OpenACC	

Tesla V100 is the flagship product of Tesla data center computing platform for deep learning, HPC, & graphics. The Tesla platform accelerates over 550 HPC applications & every major deep learning framework. It is available everywhere from desktops to servers to cloud services, delivering both dramatic performance gains and cost savings opportunities.

EVERY DEEP LEARNING FRAMEWORK



550+ GPU-ACCELERATED APPLICATIONS





XcellHost NxtGen Cloud Infrastructure

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



Tired of Slow Performing Servers.

Get XcellHost Screaming Fast Cloud Servers - Powered by Intel SSD




NxtGen Cloud Infrastructure



Intel Gold Xeon CPU



nVME SSD



DDR4 RAM



Tier 4 Datacenter



Advanced DDOS




25G Network Fabric



Control Panel




99.95 % Uptime



Powered By:





XcellHost GPU Cloud

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



Get easy & swift access to the processing power needed for Deep Learning with our GPU Cloud.



Leverage XcellHost GPU on Cloud for **Machine Learning**, **Scientific Computing**, **3D visualization**, **Artificial Intelligence**, **Deep Learning**.

Features



Several GPU types available



Flexible GPU counts per instance



Bare metal performance



GPU application frameworks



Full SSH Access



Virtual Workstations in the Cloud



Preemptible GPUs



Attach GPUs to any machine type

Benefits



Cutting-edge hardware



Save weeks of setup time



Secure & Reliable



Save hundreds of thousands



Simple & scalable



Per-hour billing



User-Friendly Interface

Order your dream Cloud Server now

<https://www.xcellhost.cloud/gpu>

Runs Graphic Intensive workloads



Supported Software




GPU Cloud Pricing



NVIDIA Tesla v100

Virtual Server NVIDIA Tesla V100	GPU	vCPU (in core)	vRAM (in GB)	vDisk (in GB)	Total Data Transfer*	Monthly Cost ₹
gpucloudv100	1 x NVIDIA Tesla V100	8	61	500	10 TB	53,568.00
gpucloud4v100	4 x NVIDIA Tesla V100 NVLink	32	244	500	10 TB	59,520.00
gpucloud8v100	8 x NVIDIA Tesla v100 NVLink	64	488	500	10 TB	29,520.00



GPU Dedicated Servers

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



Speed up the computing of Deep Learning tasks with powerful multi-GPU dedicated servers



Leverage GPUs on Cloud for **Machine Learning**, **Scientific Computing**, **3D visualization**, **Artificial Intelligence**, **Deep Learning**.

Features



Performance Network

Bare metal performance



DDOS Protection



GPU application frameworks



40 Gb/s Ports



Custom Congiurations



Managed Support & Server Monitoring



Hardware Testing



10TB Traffic

Benefits



Powerfull



Cutting-edge hardware



Scalable



Save hundreds of thousands



Secure & Reliable



Efficient



Choice of GPU types



Dedicated Resources

Order your dream Cloud Server now

<https://www.xcellhost.cloud/gpu>

Dedicated GPU Cloud Pricing



NVIDIA Tesla V100

Dedicated Server NVIDIA Tesla V100	GPU	Physical CPU (in core)	Physical RAM (in GB)	SSD Storage (in GB)	Total Data Transfer*	Monthly Cost ₹
GPU.T4	NVIDIA T4	12	60	900	10 TB	29,200.00
GPU.V100	NVIDIA Tesla V100	12	60	900	10 TB	91,250.00
1 x V100	1 x NVIDIA Tesla v100 16 GB Skylake Gold 6148v5	20	284	900	10 TB	1,55,485.44
2 x V100 PCI	2 x NVIDIA Tesla v100 PCI 16 GB 2 x Intel Xeon E5-2630v4 2.2 Ghz	20	128	960	10 TB	1,76,117.04
2 x V100 NVLink	2 x NVIDIA Tesla v100 NVLink 16 GB 2 x Intel Xeon Silver 4114 2.2 Ghz	20	128	960	10 TB	1,77,344.64
2 x V100	2 x NVIDIA Tesla V100 16 GB Skylake Gold 6148v5	20	768	960	10 TB	2,83,219.20
4 x V100 NVLink32	4 x NVIDIA Tesla V100 NVLink 32 GB 2 X Intel Xeon Gold 6138 Processor 3.70 GHZ	40	384	2048	10 TB	4,27,242.96




NVIDIA GeForce RTX 2080 Ti

Dedicated Server NVIDIA GeForce RTX 2080 Ti	GPU	Physical CPU (in core)	Physical RAM (in GB)	SSD Storage (in GB)	Total Data Transfer*	Monthly Cost ₹
2xRTX2080Ti	2x NVIDIA RTX 2080 Ti / Intel Xeon E5 - 2630V4 2.2 GHz	10	64	480	10 TB	65,072.88
2xRTX2080Ti9	2x NVIDIA RTX 2080 Ti / Intel Core i9-9980XE	18	64	1024	10 TB	78,939.36
4xRTX2080Ti	4x NVIDIA RTX 2080 Ti / 2x Intel Xeon E5-2630 V4 2.2 GHz	20	128	960	10 TB	1,64,060.64
8xRTX2080Ti	2x NVIDIA RTX 2080 Ti / Intel Xeon E5 - 2630V4 2.2 GHz	10	64	480	10 TB	1,97,310.96
8xRTX2080Ti256	2x NVIDIA RTX 2080 Ti / Intel Xeon E5 - 2630V4 2.2 GHz	10	64	480	10 TB	2,05,883.28
8xRTX2080Ti384	2x NVIDIA RTX 2080 Ti / Intel Xeon E5 - 2630V4 2.2 GHz	10	64	480	10 TB	2,17,826.64



NVIDIA GeForce GTX 1080 Ti

Dedicated Server NVIDIA GeForce GTX 1080 Ti	GPU	Physical CPU (in core)	Physical RAM (in GB)	SSD Storage (in GB)	Total Data Transfer*	Monthly Cost ₹
2x1080Ti32	2x NVIDIA GeForce GTX 1080 Ti / Intel Xeon E5-2609 v4 1.7 GHz	8	32	480	10 TB	41,564.88
2x1080Ti64	2x NVIDIA GeForce GTX 1080 Ti / Intel Xeon E5-2609 v4 1.7 GHz	8	64	480	10 TB	44,172.72
2x1080Ti64HF	2x NVIDIA GeForce GTX 1080 Ti / Intel Xeon E5-1630 v4 3.70 GHz	4	64	480	10 TB	46,073.52
4x1080Ti64	4x NVIDIA GeForce GTX 1080 Ti / 2 x Intel Xeon E5-2609 v4 1.7 GHz	16	64	480	10 TB	82,751.04
4x1080Ti128	4x NVIDIA GeForce GTX 1080 Ti / 2 x Intel Xeon E5-2609 v4 1.7 GHz	16	128	480	10 TB	94,457.52
8x1080Ti	8x NVIDIA GeForce GTX 1080 Ti / 2x Intel Xeon E5-2630 v4 2.2 GHz	20	128	480	10 TB	1,67,235.12
8x1080Ti512	NVIDIA GeForce GTX 1080 Ti / 2x Intel Xeon E5-2630 v4 2.2 GHz	20	512	480	10 TB	2,16,884.16
8x1080Ti1TB	8x NVIDIA GeForce GTX 1080 Ti / 2x Intel Xeon E5-2630 v4 2.2 GHz	20	128	960	10 TB	1,72,462.32
8x1080Ti2TB	8x NVIDIA GeForce GTX 1080 Ti / 2x Intel Xeon E5-2630 v4 2.2 GHz	20	128	2000	10 TB	1,82,046.24
10x1080Ti	10x NVIDIA GeForce GTX 1080 Ti / 2x Intel Xeon E5-2630 v4 2.2 GHz	20	128	480	10 TB	2,06,314.56



GPU Cloud On AWS

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant

Powered By:



Build powerful machine learning applications on the most advanced & highest performing GPU-accelerated cloud infrastructure

Features

- Optimize Compute Performance
- Pause and Resume Your Instances
- Bare Metal Instance
- High I/O Instances
- Dense HDD Storage Instances
- Virtual Workstations in the Cloud
- Optimized CPU Configurations
- Attach GPUs to any machine type

Benefits

- Elastic Web-Scale Computing
- Flexible, Powerful HPC
- Inexpensive
- Completely Controlled
- Simple & scalable
- Per-second billing
- User-Friendly Interface

Order your dream Cloud Server now <https://www.xcellhost.cloud/gpu>


Supported Software



GPU Cloud On AWS Pricing




EC2 P3 Instances - GPUs - Tesla V100	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
p3.2xlarge	1	16	8	61	5 TB	27,884
p3.8xlarge	4	64	32	244	5 TB	108,781
p3.16xlarge	8	128	64	488	5 TB	216,644



EC2 P2 Instances - GPUs - NVIDIA K80	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
p2.xlarge	1	12	4	61	10 TB	23,685
p2.8xlarge	8	96	32	488	10 TB	183,051
p2.16xlarge	16	192	64	732	10 TB	317,158




EC2 G3 Instances - NVIDIA Tesla M60	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
g3.4xlarge	1	8	16	122	10 TB	37,791
g3.8xlarge	2	16	32	244	10 TB	74,664
g3.16xlarge	4	32	64	488	10 TB	148,409



GPU Cloud On Azure

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



Bringing AI & High Performance Computing to Enterprise A Cloud With Improved GPU Performance, Speed & Accuracy.

Features



Optimize VM Performance



Azure Resource Manager (ARM) templates



Bare Metal Server



Attach GPU To Any Machine Type



Open, with options



Virtual Workstations In The Cloud



Hybrid Connections



High-Density Multi-Tenancy

Benefits



Connect From Any Machine



Flexible, Powerful HPC



Scaling Of VM Is Effortless



You Get First Class Download & upload speeds



Per Second Billing



Disks can be detached & restored on another machine



User-Friendly Interface

Order your dream Cloud Server now

<https://www.xcellhost.cloud/gpu>

GPU Cloud on Azure Pricing




NVIDIA Tesla K80




NVIDIA Tesla P100



NVIDIA Tesla V100



NVIDIA Tesla P40



NVIDIA Tesla M60


NC-series (NVIDIA Tesla K80)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NC6	1	16	8	61	5 TB	22,998.93
Standard_NC12	4	64	32	244	5 TB	45,217.10
Standard_NC24	8	128	64	488	5 TB	90,367.28

NCv2-series (NVIDIA Tesla P100)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NC6s_v2	1	12	4	61	10 TB	38,132.27
Standard_NC12s_v2	8	96	32	488	10 TB	75,501.63
Standard_NC12s_v2	16	192	64	732	10 TB	150,222.49

NCv3-series (NVIDIA Tesla V100)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NC6s_v3	1	8	16	122	10 TB	38,132.27
Standard_NC12s_v3	2	16	32	244	10 TB	75,501.63
Standard_NC24s_v3	4	32	64	488	10 TB	150,222.49

ND-series (NVIDIA Tesla P40)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_ND6s	1	12	4	61	10 TB	42,593.75
Standard_ND12s	8	96	32	488	10 TB	84,424.59
Standard_ND24s	16	192	64	732	10 TB	168,068.41

NV (NVIDIA Tesla M60)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NV6	1	8	16	122	10 TB	20,768.19
Standard_NV12	2	16	32	244	10 TB	40,755.62
Standard_NV24	4	32	64	488	10 TB	81,444.32



GPU Cloud on Google Cloud Platform

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



Leverage GPUs on Google Cloud for machine learning, scientific computing, & 3D visualization.

Features



Several GPU types available



Flexible GPU counts per compute engine



Bare metal performance



GPU application frameworks



Preemptible GPUs



Virtual Workstations in the Cloud



Attach GPUs to any machine type

Benefits



Cutting-edge hardware



Flexible, Powerful HPC



Save weeks of setup time



Completely Controlled



Simple & scalable



Per-second billing




User-Friendly Interface

Order your dream Cloud Server now / <https://www.xcellhost.cloud/gpucloud>

GPU Cloud on GCP Pricing




NVIDIA Tesla K80




NVIDIA Tesla P100



NVIDIA Tesla V100



NVIDIA Tesla P40



NVIDIA Tesla M60


NC-series (NVIDIA Tesla K80)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NC6	1	16	8	61	5 TB	22,998.93
Standard_NC12	4	64	32	244	5 TB	45,217.10
Standard_NC24	8	128	64	488	5 TB	90,367.28

NCv2-series (NVIDIA Tesla P100)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NC6s_v2	1	12	4	61	10 TB	38,132.27
Standard_NC12s_v2	8	96	32	488	10 TB	75,501.63
Standard_NC12s_v2	16	192	64	732	10 TB	150,222.49

NCv3-series (NVIDIA Tesla V100)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NC6s_v3	1	8	16	122	10 TB	38,132.27
Standard_NC12s_v3	2	16	32	244	10 TB	75,501.63
Standard_NC24s_v3	4	32	64	488	10 TB	150,222.49

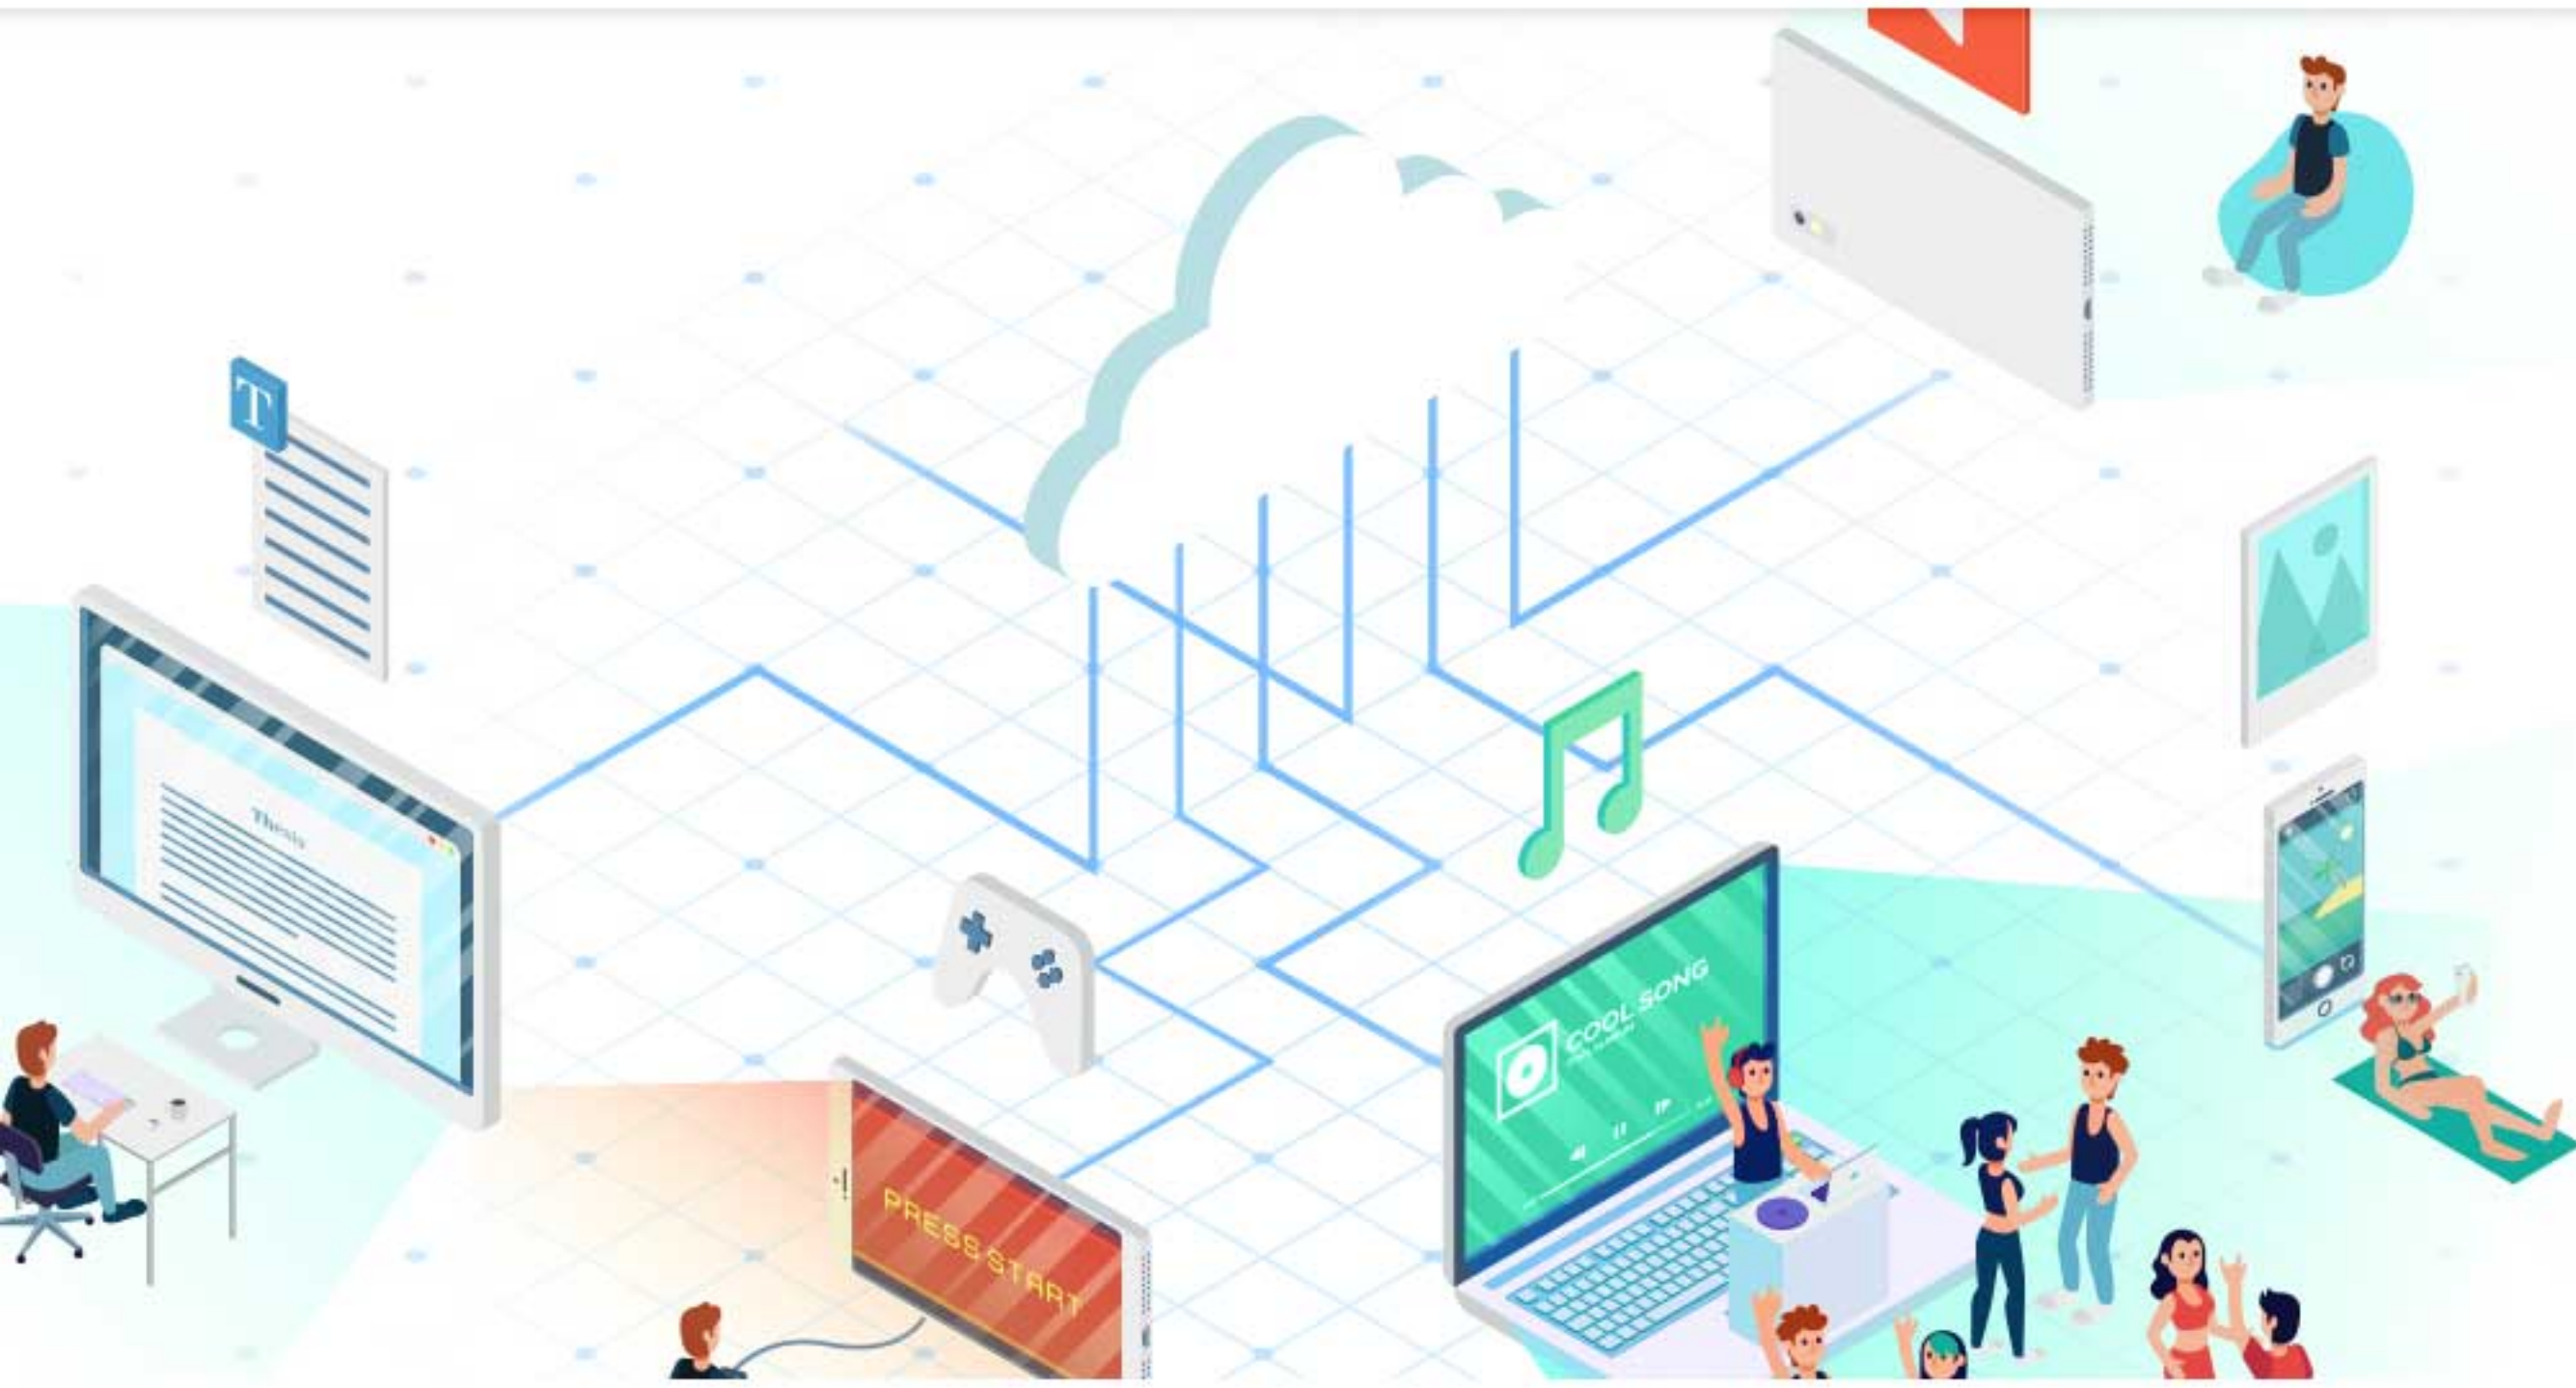
ND-series (NVIDIA Tesla P40)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_ND6s	1	12	4	61	10 TB	42,593.75
Standard_ND12s	8	96	32	488	10 TB	84,424.59
Standard_ND24s	16	192	64	732	10 TB	168,068.41

NV (NVIDIA Tesla M60)	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
Standard_NV6	1	8	16	122	10 TB	20,768.19
Standard_NV12	2	16	32	244	10 TB	40,755.62
Standard_NV24	4	32	64	488	10 TB	81,444.32




GPU Cloud on Alibaba

Secure Cloud Services for your online business since 1999.
.Reliable .Secure .Speed .Scalable .Manageable .Compliant



The GPU Computing on Alibaba Cloud can work with NGC to significantly Reduce the development environment provisioning time.

Features



Running high-performance NVMe drives



Computing capabilities



Extraordinary general network performance



Attach GPU To Any Machine Type



Multiple payment methods



Virtual Workstations In The Cloud



Hybrid connections



High-Density Multi-Tenancy


Benefits



Simple & Scalable



Flexible , powerfull HPC



Video Processing



Scientific Computing



Per Hour Billing



Disks can be detached & restored on another machine



User-Friendly Interface

Ask for GPU Cloud Virtual + Physical Servers

<https://www.xcellhost.cloud/gpu>

Common Scenarios



Online rendering in the cloud (GA1)



General-purpose GPU computation (GN4)



Outstanding computation acceleration (GN5)



Outstanding computation acceleration (GN5)

Upgraded Support For You



1 on 1 Presale Consultation



24/7 Technical Support



6 Free Tickets per Quarter




Faster Response

Alibaba GPU Cloud Pricing



EC2 P3 Instances - GPUs - Tesla V100	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
p3.2xlarge	1	16	8	61	5 TB	27,884
p3.8xlarge	4	64	32	244	5 TB	108,781
p3.16xlarge	8	128	64	488	5 TB	216,644



EC2 P2 Instances - GPUs - NVIDIA K80	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
p2.xlarge	1	12	4	61	10 TB	23,685
p2.8xlarge	8	96	32	488	10 TB	183,051
p2.16xlarge	16	192	64	732	10 TB	317,158



EC2 G3 Instances - NVIDIA Tesla M60	GPUs	GPUs RAM	vCPU (in core)	vRAM (in GB)	Total Data Transfer*	Monthly Cost ₹
g3.4xlarge	1	8	16	122	10 TB	37,791
g3.8xlarge	2	16	32	244	10 TB	74,664
g3.16xlarge	4	32	64	488	10 TB	148,409



• INDIA • DUBAI • SINGAPORE

209, Laxmi Plaza, Bldg. No. 9
Laxmi Industrial Estate,
Andheri (W), Mumbai - 400053
MAHARASHTRA, INDIA