Intel® Unnati

Data-Centric Labs in Emerging Technologies

Give Your Students the Intel Edge.

intel.

High Performance Computing

Today, there is a wide, and growing, skill gap between technical graduates and IT industry expectations. To propel India's digital economy transformation, it is imperative that the higher education system in the country bridges this gap by developing new curricula and offering courses in emerging technologies. The National Education Policy 2020[1] recognises this, and stresses the need for greater industry-academic linkages, and for higher education institutions to focus on research and innovation.

With the **Intel® Unnati Program**, you can keep pace with fast changing industry needs and expectations. It will help you:

Equip your students with industry relevant data-centric skills

In this age of data explosion, there is immense opportunity. Give your students the edge by equipping them with data-centric skills that will help them glean better insights and develop high-value solutions.

Unleash your students' creative potential

We, in India, have an incredible opportunity to unleash the creative potential of the largest student population in the world by training them in the right skills to drive India's digital transformation.

Build a strong reputation

With an Intel co-branded lab, you can be recognised as an institute that is committed to train your students in the latest technology to prepare them for industry, and focus on faculty development.

Build capability for the long term

Establish and maintain your leadership with the help of our **System Integrator Associates**. From Intel's recommendations for end-to-end technology labs set up and course content to training, customisations of your lab set up, or your maintenance and support requests, you can rely on them for all your needs.

[1]Ministry of Human Resource Development, Government of India, National Education Policy 2020, *https://static.pib.gov.in/WriteReadData/userfiles/NEP_Final_English_0.pdf*

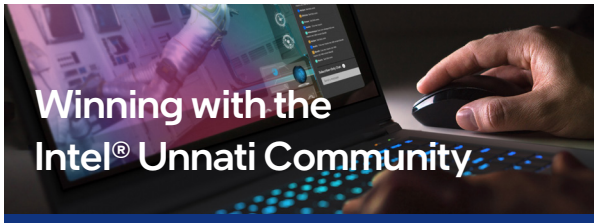

Winning with the Intel® Unnati Community

With an Intel® Unnati Lab, you—and your faculty and students—become part of the **Intel® Unnati Community**, and get exclusive benefits:

- **Intel® Unnati Grand Challenge**, where students solve industry relevant, high impact problems through technology, with cash awards and the opportunity to be evaluated for internships at Intel
- **Intel® Unnati Ignite** workshops that offer hands-on experience with Intel technologies
- **Intel® Unnati Catalyst** co-sponsorships of events focused on new technologies
- **Intel® Unnati Industrial Training**, where qualifying students work for a month on industry relevant problems under the guidance of industry mentors
- **Intel® Unnati Research Launchpad**, which offers grants to faculty members for original research in new and emerging technologies

Intel® Unnati

High Performance Computing

What is High Performance Computing?

High performance computing (HPC) is a class of applications and workloads that perform computationally intensive operations across multiple resources, and has been key to research and innovation in both academia and industry for several years.

HPC is fueling breakthroughs in areas like autonomous driving, climate change, seismic engineering, oil and gas production, precision medicine, and financial risk assessment.

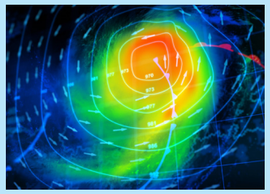


HPC at Work

(Clockwise from top left) Weather Forecasting; Automotive Industry; Genetic Engineering; and Aerospace

While HPC can be run on a single system, its real power comes from connecting multiple HPC nodes in a cluster. This provides the capabilities to compute large-scale simulations that are not feasible on a single system.

Parallel Programming

Modern HPC clusters can include various nodes that contain multi-core processors and include accelerators such as GPUs and FPGAs. To take full advantage of these compute resources, parallel programming is required.

In parallel programming, a complex problem is first broken down into smaller tasks that can be run at the same time using multiple compute resources. While there is initial effort required to set up this parallelism, the benefits in speed and performance are usually significantly higher, justifying the investment.

Heterogenous Computing and Intel® oneAPI

The increasingly heterogenous HPC compute environment requires a standard, simplified programming model that can run seamlessly across architectures while getting the maximum available performance out of the hardware. The Intel® oneAPI toolkits, targeting Intel CPUs and accelerators, provides this balance between optimal use of underlying hardware features and code that is portable, easily maintainable, and power-efficient.

HPC and AI: A Powerful Combination

- The architecture required for HPC has many similarities with AI implementations. Both use high levels of compute and storage, large memory capacity, and high-bandwidth fabrics to achieve results—typically by processing massive data sets of increasing size.
- The promise of AI in HPC is that AI models can augment expert analysis of data sets to produce results faster at the same level of accuracy.
- On the other hand, HPC can help address the challenges of AI at scale, such as the need for more parallelism, fast I/O for massive data sets, and efficient navigation of distributed computing environments.

HPC Made Simple

Students will be introduced to key foundational HPC concepts, parallel programming, software frameworks and toolkits, that simplify implementation of HPC applications on Intel® XPUs and accelerators with Intel's industry-leading compiler technology and libraries.

Suitable for Students and Researchers

With two available variant configurations, the lab can be designed to cater to students who are just getting started in HPC as well as students and researchers who require a powerful compute cluster for advanced HPC workloads.

Learn the toolkits the industry uses for building and analysing HPC workloads

Students will learn about and use the Intel® oneAPI HPC Toolkit, used by developers the world over to build, analyse, optimise, and scale HPC applications.

Taking on the world's biggest challenges through HPC and AI

The combination of HPC and AI is being increasingly used to work on some of the world's biggest challenges. Deep learning is a great match for problems addressed by HPC that involve very large, multidimensional data sets. Provide your institution with the infrastructure to enable your students and researchers take the first steps towards making a strong impact.

High Level Design

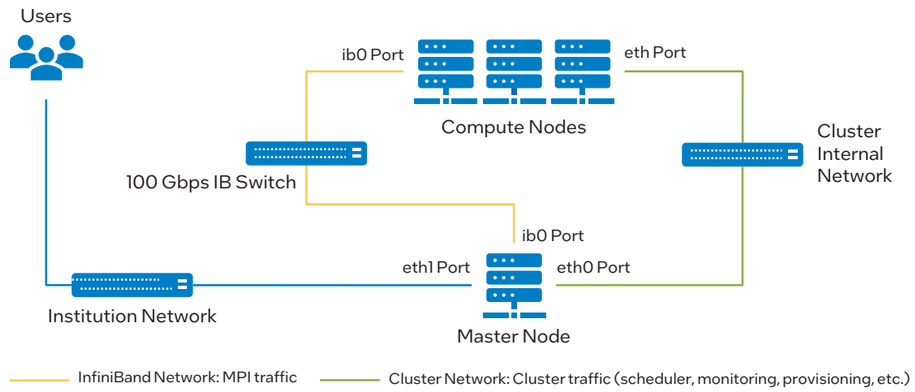

HPC Developer Lab

For institutions that are seeking to introduce HPC to their students with their first small, yet capable, cluster

Important notes relating to all AI Lab Configurations

1. If Network File System (NFS) is installed, then all Intel software tools need to be installed only once through any of the compute nodes on to a network location visible to all nodes. If you don't have NFS, then only the runtime components of Intel tools need to be installed on the hard drive of every node.
2. Initialise MPI environment first before installing Horovod*. (source setvars.sh)

Minimum Suggested Specifications

Hardware				Software
Server/Workstation	Component	Product Description	Quantity per System	
▪ 1 x Master Node	Memory	32GB RDIMM, 1 Rank, 4800MHz	8	▪ OS: CentOS* or Ubuntu*
	Processor	Intel® Xeon® Silver 4410, 12 Core, 2 GHz, 150W	2	▪ Intel® oneAPI Base and HPC Toolkit (Multi-Node) – Workgroup – Up to 10 Developers Academic (ESD)
	Management Key Options	Advanced System Management Key	1	▪ At least stable release of reputed Open Source workload and cluster management software suite (Rocks Cluster Manager)
	Power Supply	1300W Power Supply	2	▪ Scheduling and cluster management software should support policy-based workload management graphical cluster administration interface, monitoring and reporting tools, job scheduler (Slurm)
	Dual Hyper Hybrid Controller	M.2 Card RT3EX020E with HW RAID	1	
	TPM	Version 2.0	1	
	M.2 Drive	480GB SATA	2	
	Front Bay 1, Drive 01	1.92T 2.5 NVME4	1	
	Onboard NVME Cable 1	Cable Kit CBLMCSL1204KIT	1	
	Networking IO Modules	X710-T2L for OCP 3.0	1	
	+ Heat sink, Riser Card, Rail Option (Please discuss with System Integrator for options)			
▪ Minimum 2 x Compute Nodes	Memory	32GB RDIMM, 1 Rank, 4800MHz	8	
	Processor	Intel® Xeon® Gold 6430, 32 Core, 2.1 GHz, 270W	2	
	Management Key Options	Advanced System Management Key	1	
	Power Supply	1300W Power Supply	2	
	Dual Hyper Hybrid Controller	M.2 Card RT3EX020E with HW RAID	1	
	TPM	Version 2.0	1	
	M.2 Drive	480GB SATA	2	
	Front Bay 1, Drive	Front Bay 1, Drive	1	
	SATA Cable	1.92T 2.5 NVME4	1	
	Networking IO Modules	X710-T2L for OCP 3.0	1	
	+ Heat sink, Riser Card, Rail Option (Please discuss with System Integrator for options)			

HPC Developer Lab (continued)

Hardware				Software
Server/Workstation	**Component**	**Product Description**	**Quantity per System**	
▪ 1 x Storage Server	2 TB space (or as per requirements)		1	
▪ Network Router with 10Gbps Ports, Network Switch, Rack Cabinet, Power Delivery Unit (PDU), Patch Cables and Power Cable ▪ Intel® OpenVino™ labs will be run on Intel® DevCloud for the Edge				Software for job scheduling and queueing

HPC Research Lab

For institutions that are working with HPC applications and are seeking to upgrade their HPC infrastructure to a mid-size or large cluster

Minimum Suggested Specifications

Hardware				Software
Server/Workstation	**Component**	**Product Description**	**Quantity per System**	
▪ 1 x Master Node	Memory	32GB RDIMM, 1 Rank, 4800MHz	8	▪ Ubuntu* 18.04 ▪ Intel® oneAPI Base Toolkit ▪ Intel® AI Analytics Toolkit ▪ Horovod* + Intel® MPI (for distributed DLtraining with TensorFlow*) ▪ Intel® Extension for PyTorch* (IPEX) ▪ Intel® Distribution of OpenVINO™ Toolkit ▪ JupyterHub* and JupyterLab* ▪ Keras*, ipykernel*, Seaborn* + other libs as required by exercises **Note:** Check *https://software.intel.com/* containers for available AI containers
	Processor	Intel® Xeon® Silver 4410, 12 Core, 2 GHz, 150W	2	
	Management Key Options	Advanced System Management Key	1	
	Power Supply	1300W Power Supply	2	
	Dual Hyper Hybrid Controller	M.2 Card RT3EX020E with HW RAID	1	
	TPM	Version 2.0	1	
	M.2 Drive	480GB SATA	2	
	Front Bay 1, Drive 01	1.92T 2.5 NVME4	1	
	Onboard NVME Cable 1	Cable Kit CBLMCSL1204KIT	1	
	Networking IO Modules	X710-T2L for OCP 3.0	1	
	+ Heat sink, Riser Card, Rail Option (Please discuss with System Integrator for options)			
▪ Minimum 6 x Compute Nodes	Memory	32GB RDIMM, 1 Rank, 4800MHz	8	
	Processor	Intel® Xeon® Gold 6430, 32 Core, 2.1 GHz, 270W	2	
	Management Key Options	Advanced System Management Key	1	
	Power Supply	1300W Power Supply	2	
	Dual Hyper Hybrid Controller	M.2 Card RT3EX020E with HW RAID	1	
	TPM	Version 2.0	1	
	M.2 Drive	480GB SATA	2	
	Front Bay 1, Drive 01	1.92T 2.5 NVME4	1	
	SATA Cable	Cable Kit CYPCBLHDHDXXX1	1	
	Networking IO Modules	X710-T2L for OCP 3.0	1	
	+ Heat sink, Riser Card, Rail Option (Please discuss with System Integrator for options)			

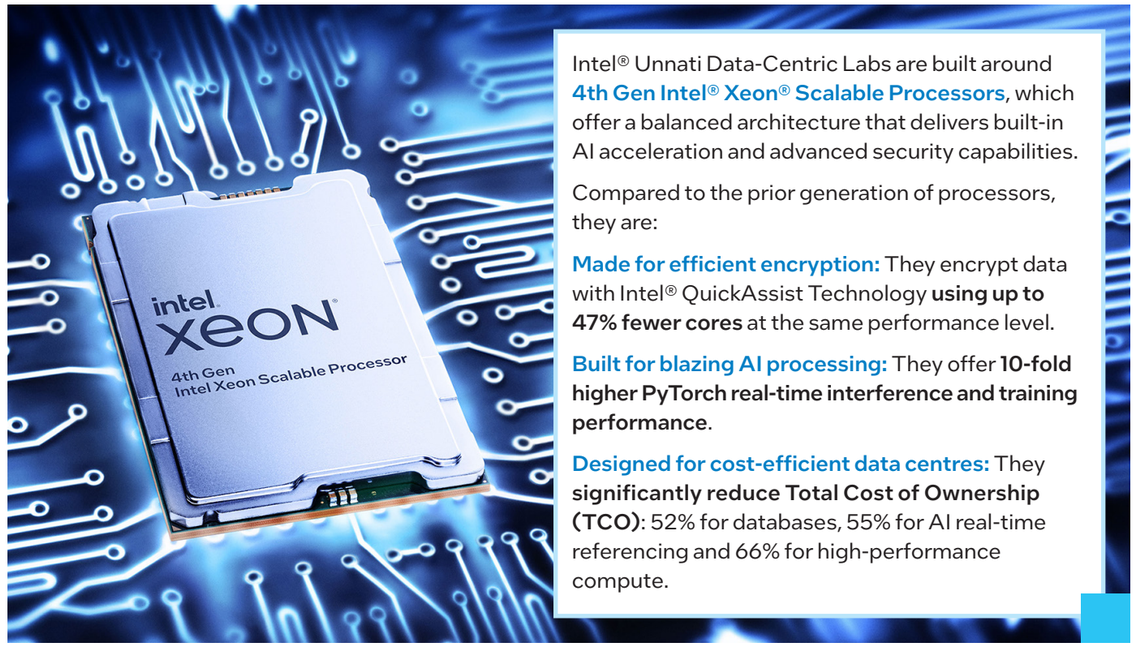

Intel® Unnati Data-Centric Labs are built around **4th Gen Intel® Xeon® Scalable Processors**, which offer a balanced architecture that delivers built-in AI acceleration and advanced security capabilities.

Compared to the prior generation of processors, they are:

Made for efficient encryption: They encrypt data with Intel® QuickAssist Technology **using up to 47% fewer cores** at the same performance level.

Built for blazing AI processing: They offer **10-fold higher PyTorch real-time interference and training performance**.

Designed for cost-efficient data centres: They **significantly reduce Total Cost of Ownership (TCO)**: 52% for databases, 55% for AI real-time referencing and 66% for high-performance compute.

HPC Research Lab (continued)

Minimum Suggested Specifications

Hardware				Software
Server/Workstation	**Component**	**Product Description**	**Quantity per System**	
▪ 1 x GPU Node	Memory	32GB RDIMM, 1 Rank, 4800MHz	8	▪ Ubuntu* 18.04 ▪ Intel® oneAPI Base Toolkit ▪ Intel® AI Analytics Toolkit ▪ Intel® Distribution of OpenVINO™ Toolkit
	Processor	Intel® Xeon® Gold 6438, 32 Core, 2.3 GHz, 205W	2	
	GPU	Intel® Data Center GPU Max 1100	1 or 2	
	Management Key Options	Advanced System Management Key	2	
	Power Supply	1300W Power Supply	2	
	Dual Hyper Hybrid Controller	M.2 Card RT3EX020E with HW RAID	1	
	TPM	Version 2.0	1	
	M.2 Drive	480GB SATA	2	
	Front Bay 1, Drive 01	1.92T 2.5 NVME4	1	
	SATA Cable	Cable Kit CYPCBLHDHDXXX1	1	
	Networking IO Modules	X710-T2L for OCP 3.0	1	
	+ Heat sink, Riser Card, Rail Option (Please discuss with System Integrator for options)			
▪ 1 x Storage Server	50 TB space (or as per requirements)		1	
▪ Primary Interconnect: 36-port 100 Gbps Infiniband/OPA switch with rail kit and 2PSU ▪ Secondary Interconnect: 24-Port 10 GbE Ethernet switch with required number of 2m CAT6 cables				▪ Software for job scheduling and queueing

To know more about how your institution can benefit from the Intel® Unnati Program, please contact:

Copyright © 2023 Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Agilex, Arria, Movidius, OpenVino and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others.

Intel_Unnati_HPC_Jul 2023