सत्यमेव जयते

NITI Aayog

AIRAWAT – ESTABLISHING AN AI SPECIFIC CLOUD COMPUTING INFRASTRUCTURE FOR INDIA

An Approach Paper

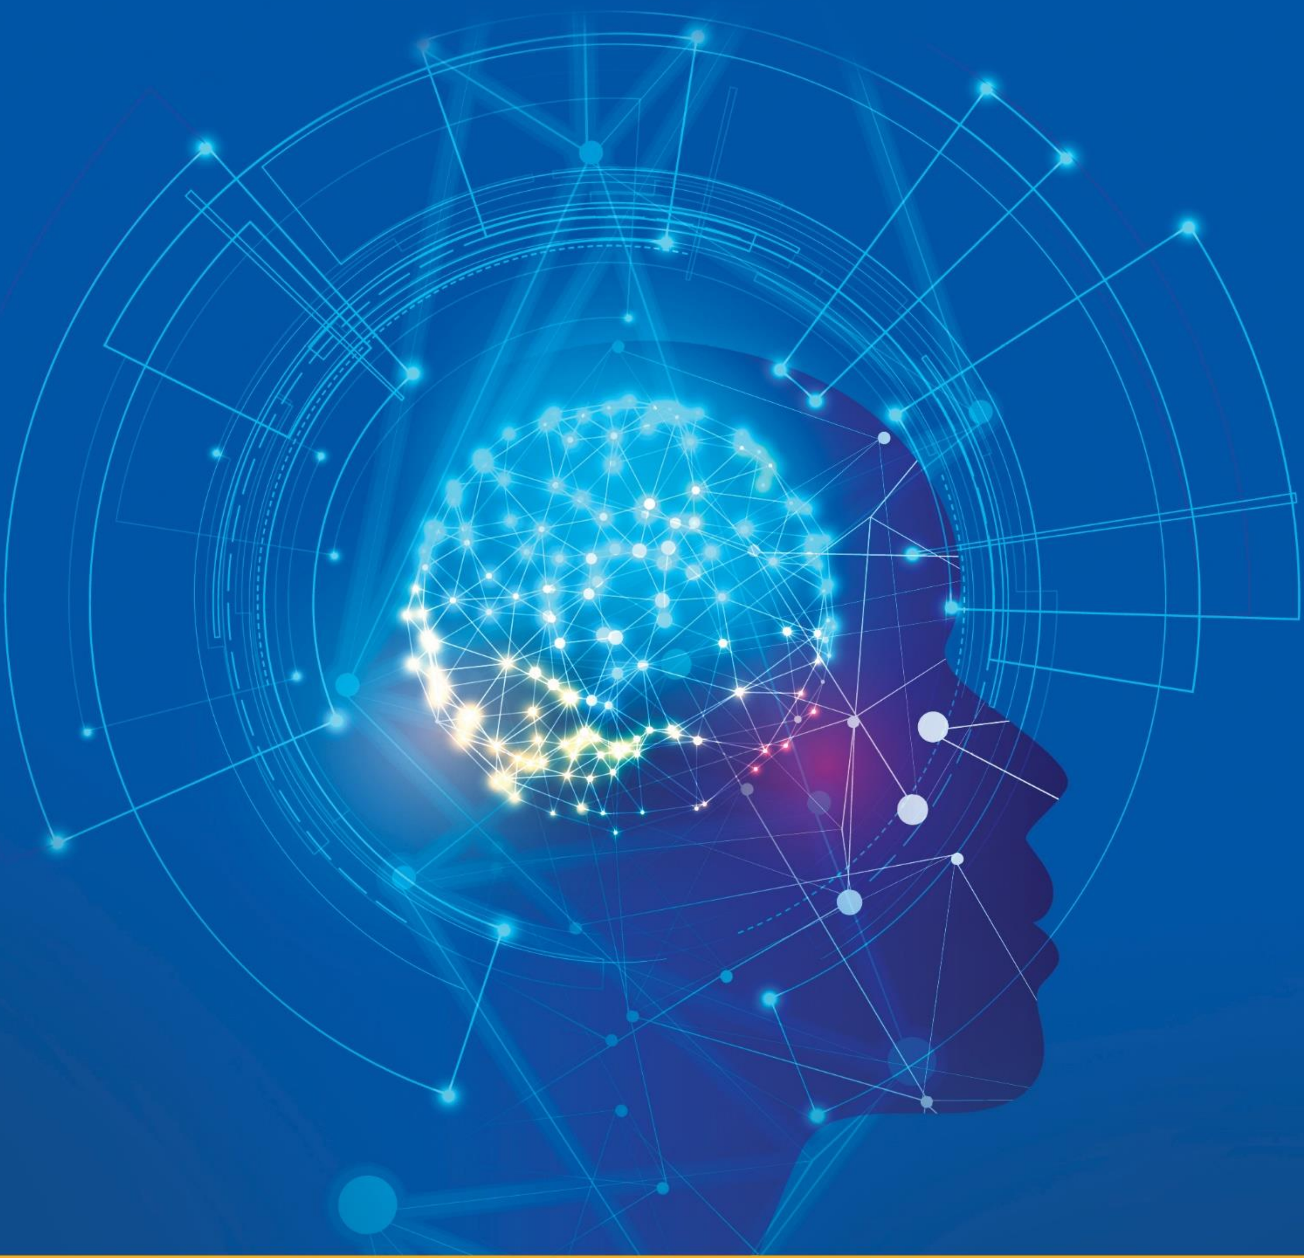

JANUARY 2020

This page has been intentionally left blank

Acknowledgments

In writing this Approach Paper, Arnab Kumar, Tanay Mahindru and Ira Saxena of NITI Aayog have made valuable contributions.

The contributions of several of our technology partners are also thankfully acknowledged.

Anna Roy
Senior Adviser
NITI Aayog

Contents

Introduction

Series of 'Approach Papers' detailing key recommendations of NSAI

NITI Aayog released the National Strategy for Artificial Intelligence (NSAI) discussion paper in June 2018, in pursuance of the mandate entrusted to it by the Hon'ble Finance Minister in the Budget Speech of 2018 – 2019. NSAI highlighted the potential of Artificial Intelligence (AI) in boosting India's annual growth rate by 1.3 percentage points by 2035 and identified priority sectors for the deployment of AI with Government's support (Healthcare, Agriculture, Education, Smart Cities and Mobility). NSAI also emphasized on four broad recommendations in supporting and nurturing an AI ecosystem in India: (a) promotion of research; (b) skilling and reskilling of the workforce; (c) facilitating adoption of AI solutions; and (d) the development of guidelines for 'responsible AI'.

While the discussion paper strived to clearly delineate the Government's role in the promotion of AI and identify potential initiatives, 'in-depth' analysis of the key recommendations was also subsequently pursued to develop implementation blueprints. Termed 'Approach Papers', these documents strive to present detailed plans for implementation of selected recommendations of the strategy.

It is thus with extreme pleasure, I present the first in a series of 'Approach Papers', titled 'AIRAWAT: An AI Specific Cloud Compute Infrastructure'.

Though novel in its scope, AIRAWAT is well in line with India's recent approach to innovation in fields of emerging and digital technology fields. This has been an approach of facilitation of innovation, rather than implementation, where we have seen large government funding for the creation of digital infrastructure aimed at enabling research and innovation, like the creation of the Unified Payments Interface (UPI), an underlying infrastructure for payments. UPI has grown tremendously over just 4 years as multiple products and innovators have leveraged it's capabilities and is widely credited for India's digital payments revolution.

As a computing facility designed specifically to execute tasks relevant to Machine Learning (ML) / Deep Learning (DL) applications, it is our hope that AIRAWAT will have a similar effect of bolstering AI research and application in India. This paper, while highlighting the urgent need for such a facility for Indian researchers and innovators, benchmarks other similar facilities being developed across the world and proposes the architecture, governance structure and mechanism of selecting various stakeholders involved in the implementation of AIRAWAT.

Anna Roy

Senior Adviser, NITI Aayog

Background

Recommendation from National Strategy for Artificial Intelligence

AI has evolved as a transformative new technology, capable of delivering large incremental value to a wide range of sectors. For India, AI presents the potential to address some of the biggest challenges that are currently faced: access, affordability and availability of quality healthcare, education, agronomics, mobility solutions etc.

While not a new technology, the recent advances in innovation, specifically in three areas: (a) computing power, (b) data storage and (c) volume of digitized data, have led to AI-based applications taking center stage in presenting a radical new approach to solving business and governance use cases.

Government's focus on digitalization and impressive strides under Digital India have enabled generation of large quantum of digital data. However, access to specialised compute and storage facilities would be crucial to achieve the economic potential of AI. AI has the potential to raise India's annual growth rate by 1.3 percentage points and add USD957 billion to India's economy in 2035, as per the studies done by Accenture[1].

Recognising the potential of AI to help achieve the goal of India becoming a USD5 trillion economy, NITI Aayog was tasked to establish the National Program on AI, with the aim to guide the research and development of new and emerging technologies. In pursuance of the above mandate, NITI Aayog released the "National Strategy for Artificial Intelligence" (NSAI) on 4th June, 2018.

NSAI analyses the current landscape of AI research and adoption in India, and identifies the impediments that handicap our progress. Among the key ones include:

(a) Lack of scale for experimental validation;

(b) Lack of facilities to support large scale experimental test beds; and

(c) High cost and low availability of computing infrastructure required for development, training and deployment of AI based services

To address these handicaps, one of the key recommendations is to set up an AI-specific cloud infrastructure to facilitate research and solution development in using high performance and high throughput AI-specific supercomputing technologies, nicknamed AIRAWAT, i.e. ***AI Research, Analytics and knoWledge Assimilation*** platform.

[1] Rewire for Growth: Accelerating India's Economic Growth with Artificial Intelligence, Accenture

Need for AI Specific Cloud Compute Infrastructure

Why High Performance Computing is not good enough?

AI solutions, more specifically ML (including DL) solutions, require processing huge number of calculations quickly, thus necessitating increased processing power. Developing ML solutions can be seen as a two-step process: (a) training and (b) inference. "*The training phase is essentially an optimization problem in a multi-dimensional parameter space, and involves building a model that can be used to provide a wider generalization in the inference process. In DL, a model usually consists of a multilayer network with many free parameters (weights) whose values are set during the training process. Once trained, the model now needs to be deployed on real-world data in the inference mode. For many applications, this inference step needs a trained model that is fixed for consistency, reproducibility, liability, performance or regulatory reasons*"[2]. The demands and intensity of training and inference would thus determine the need for advanced processing capabilities.

The recent advances in ML have been driven by the advent of specialised computing infrastructure. The computing infrastructure landscape for AI is continuously evolving and leadership in AI will entail adopting and adapting to latest AI-specific computing innovations. The biggest breakthrough in ML was realised in early 2000s when Graphics Processing Units (GPUs), initially developed for gaming and 3D graphics, were trained to process ML solutions. Since then, technology giants have invested heavily on AI hardware infrastructure, notable ones being nVIDIA's DGX, Google's Tensor Processing Units (TPUs), Microsoft's Azure, IBM's Watson and Intel's Nervana. The evolution that started from moving AI computations, which can be understood as mostly linear algebra operations, from CPUs (Central Processing Units) to GPUs, have now moved to specialised chips, TPUs etc., designed specifically for parallelized linear algebra computations.

AI computing infrastructure is distinct from High Performance Computing (HPC) infrastructure and the difference needs to be well understood for purposes of future infrastructure planning. HPCs, with its origins in particle physics simulations, have dominated the hardware development for several decades. "*HPCs are designed by aggregating clusters of computers designed specifically for delivering higher performance (as compared to a typical desktop computer or workstation) in order to solve large problems in science, engineering, or business*[3]". From a storage perspective, AI infrastructure involves very large datasets and storage transactions that are read-dominated at the beginning of each epoch (an epoch is defined as one complete pass-through of the dataset, inclusive of multiple iterations of model parameter updates). This differs from typical HPC applications which

[2] Future Computing Hardware for AI, IBM Research
[3] InsideHPC

are write-intensive. ML / DL training is usually static, involving large groups of random reads, accessed repeatedly, since the same data is used for training over and over. The following chart captures representative difference between an HPC and a GPU-enabled AI compute infrastructure:

Figure 1: *Difference between HPCs and GPU-enabled AI Compute Infrastructure*

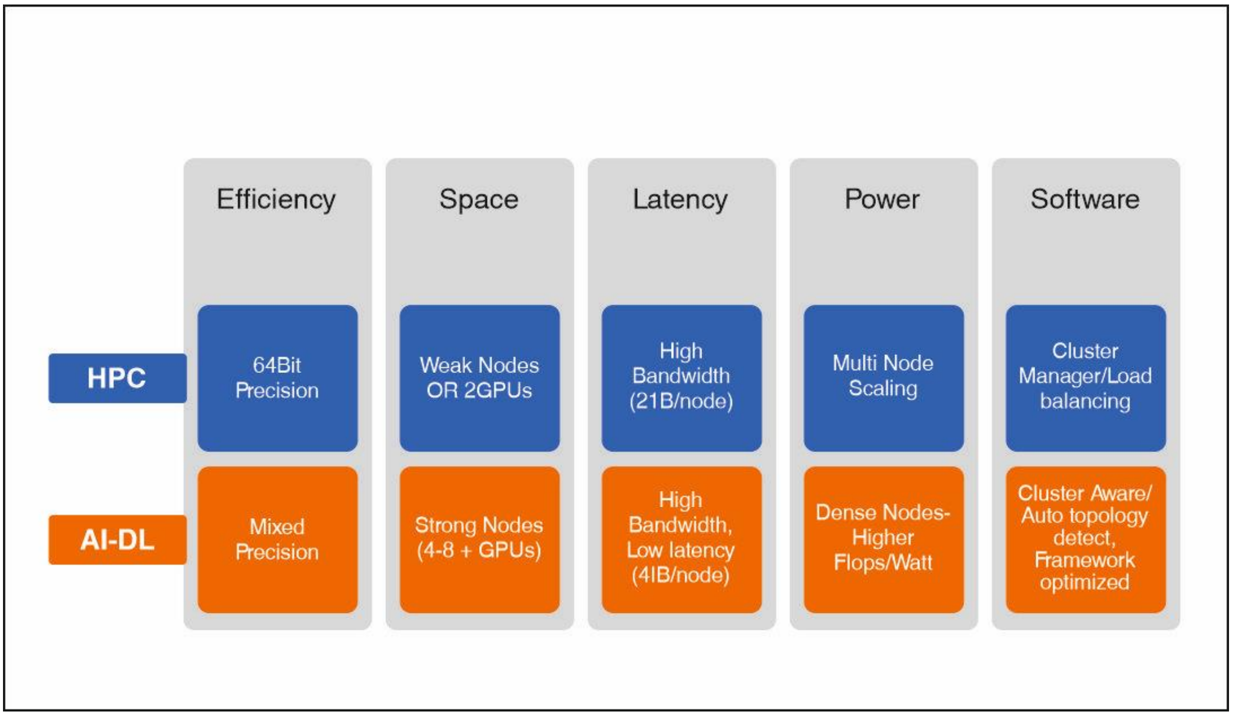

Source: nVIDIA, Wipro

The iterative nature of optimization of ML / DL algorithms necessitates the availability of a large amount of specialized computing resources for their continuous testing. The lack of availability of these resources is often cited as a major hurdle to the creation of a vibrant ecosystem for research in AI in India[4]. It is envisaged that if made available, the specialized compute resource would not only significantly improve the outlook of research in the field in India, but also increase India's competitiveness in international conferences and journal publications. The building of an indigenous compute facility, rather than increasing reliance on third party solutions (AWS, Azure, etc.) would also allay concerns of data privacy, while simultaneously increasing capacity to create and deploy similar facilities in India in the future.

[4] Landscape of AI / ML Research in India, Itihaasa Research and Digital, 2018

Benchmarking India's AI Computing Infrastructure

Not in the same league as US and China

Countries and institutions have invested heavily in AI-specific large computing infrastructures. Examples include the Summit supercomputer at Oak Ridge National Laboratory (developed for US Department of Energy) and AI Bridging Cloud Infrastructure (ABCI) commissioned by National Institute of Advanced Industrial Science and Technology (AIST) in Japan. These systems have been designed keeping in mind AI specific workloads, and thus could be perceived as gold standards in designing AI computing infrastructure.

As noted above, these platforms involve high speed machines capable of doing faster calculations, consuming energy much lower than traditional supercomputers and ability to efficiently store and process petabytes of data.

Amongst the biggest handicaps that India faces includes lack of AI infrastructure. Ranking and detailing for the 500 most powerful non-distributed computer systems in the world is done by TOP500. The TOP500 rankings[5], released twice every year, are based on a LINPACK benchmark, which is a measure of a system's floating-point computing power, or how fast a computer solves a dense system of linear equations. The list of top 500 supercomputers (all of which now have more than one petaflop of capabilities each), as benchmarked by TOP500, is dominated by China with 228 such facilities followed by the USA (117 systems) and Japan (29 systems).

***Table 1**: Top 10 Supercomputers in the World (November 2019)*

Rank	Name	Manufacturer	Country	Year	Segment	Rmax [TFlop/s][6]	Rpeak [TFlop/s][7,8]
1	Summit	IBM	USA	2018	Research	148,600	200,795
2	Sierra	IBM / nVIDIA / Mellanox	USA	2018	Research	94,640	125,712
3	Sunway TaihuLight	NRCPC	China	2016	Research	93,015	125,436
4	Tianhe-2A	NUDT	China	2018	Research	61,445	100,679
5	Frontera	Dell EMC	USA	2019	Academic	23,516	38,746

[5] https://www.top500.org/

[6] A system's Rmax score describes its maximal achieved performance

[7] A system's Rpeak score describes its theoretical peak performance

[8] Mflop/s is a rate of execution, millions of floating point operations per second. Whenever this term is used it will refer to 64 bit floating point operations and the operations will be either addition or multiplication. Gflop/s refers to billions of floating point operations per second and Tflop/s refers to trillions of floating point operations per second.

6	Piz Daint	Cray	Switzerland	2017	Research	21,230	27,154
7	Trinity	Cray	USA	2017	Research	20,159	41,461
8	ABCI	Fujitsu	Japan	2018	Research	19,880	32,577
9	SuperMUC-NG	Lenovo	Germany	2018	Academic	19,477	26,874
10	Lassen	IBM / NVIDIA / Mellanox	USA	2018	Research	18,200	23,047

A total of 145 of all the supercomputers on the list (nearly 28 percent) feature elements acceleration or co-processing, with 133 of those systems using Nvidia GPUs – essentially graphics co-processors rearchitected and retooled as parallel processing engines[9].

Compared to China or the USA, India has been sliding down the charts in the TOP500 ranking, with only two such supercomputing systems in the list, down from five such systems in the top 500 just 18 months ago. The primary reason for India's downward trajectory has been the accelerated investments in building new supercomputers / upgrading existing systems globally, which has pushed the entry criteria for top 500 from 716 Tflop/s mark in June 2018 to 1,142 Tflop/s mark in November 2019.

Table 2: India's Top 5 Supercomputers

TOP500 Rank	Name	Manufacturer	Year	Segment	Rmax [TFlop/s]	Rpeak [TFlop/s]
57	Pratyush	Cray	2018	Research	3,764	4,006
100	Mihir	Cray	2018	Research	2,570	2,809
N/A	InC1	Lenovo	2018	Industry	1,123	1,413
N/A	SERC	Cray	2015	Academic	902	1,244
N/A	IITM	iDataPlex	2013	Research	719	791

Notwithstanding the fact that India's supercomputing facilities fare on the lower spectrum compared to the rest of the world, the current supercomputing facilities in India are not designed for AI applications. The existing infrastructure does not lend itself to be upgraded for AI workload, are designed for specific purposes, (e.g. Indian Institute of Tropical Meteorology supercomputer designed for weather modelling), and are running at full capacity. The existing traditional supercomputing infrastructure in India are also available only at a few places e.g. top-tier institutes and Government establishments, making them inaccessible to the larger ecosystem of start-ups and

[9] https://www.zdnet.com/article/the-rise-fall-and-rise-of-the-supercomputer-in-the-cloud-era/

other institutions. Existing programmes aimed at computing infrastructure upgradation, e.g. the National Supercomputing Mission (NSM), are predominantly HPC oriented with no focus on AI infrastructure.

The current approach to developing AI compute infrastructure in India has been a decentralised one i.e. building localised limited AI computing infrastructure, which limits its applicability, including:

(a) can only cater to small-scale R&D work;

(b) requires tremendous efforts and investment to collaborate, aggregate and administer; and

(c) costs manifolds as data center, support and administration costs are usually repetitive.

Even the most ambitious of our AI infrastructure efforts planned, the upcoming AI supercomputing facility at CEERI Delhi, will have a modest capability of 5 petaflops[10].

Other approaches to AI compute in India has been to depend on cloud based AI services from the likes of AWS and Microsoft Azure. While these are efficient solutions to address compute facilities, limitations include data sharing concerns, non-predictable and high bandwidth costs etc. Such an approach is suitable perhaps for pay-as-you-go and small instance requirements.

From a storage perspective, GoI's MeghRaj is also designed for cloud services, and not for AI workloads. The underlying architecture is CPU based, and can't be upgraded to add GPU nodes.

[10] https://www.ceeri.res.in/csir-ceeris-delhi-centre/

Introducing AIRAWAT

Key design considerations

Existing and recent efforts of the Government, viz. NSAI and the National Mission for Interdisciplinary Cyber Physical Systems (NM-ICPS), have emphasized the need for enhancing both the core and applied research capabilities in AI, through initiatives like setting up of COREs (Centers of Research Excellence), ICTAIs (International Centers Transformational AI) and Innovation Hubs. In addition several other initiatives are being taken by governments and private sector to increase the adoption of AI, both in governance and private enterprises. These initiatives would spur the demand and necessity for state-of-the-art and specialised AI computing infrastructure.

In order to meet this demand and tackle the challenges associated with lack of access to computing resources highlighted, it is proposed that an AI-specific compute infrastructure be established. Such an infrastructure will power the computing needs of COREs, ICTAIs and Innovation Hubs, as well as facilitate the work of broader spectrum of stakeholders in the AI research and application ecosystem (startups, researchers, students, government organizations, etc.).

The proposal to establish **India's own AI-first compute infrastructure** is aimed to facilitate and speed up research and solution development for solving India's societal challenges using high performance and high throughput AI-specific supercomputing technologies. The key design considerations for this infrastructure are:

1. *Institutional framework for implementation:* an interdisciplinary task force
2. *Structure of the facility*: whether it should be centralized (in a single location), decentralized (access from across multiple locations) or utilize existing infrastructure (through existing Cloud Service Providers or existing HPC infrastructure);
3. *Modes of access:* whether it should be made available similar to access mechanisms for a traditional HPC or through as a fully managed cloud service;
4. *Architecture of facility:* what would constitute the roader technical design considerations;

The proposed infrastructure is acronymed AIRAWAT, i.e. the "AI Research, Analytics and knowledge Assimilation platform") and the design suggested is in line with the recommendations of the NSAI.

Institutional framework for Implementation

NSM envisages developing a supercomputing grid of more than 70 high-performance computing facilities. AIRAWAT is expected to complement the infrastructure developed for NSM, with specific focus on AI computing.

Given the inter-disciplinary nature of AI that would involve multiple entities, NITI Aayog recommends setting up of an inter-ministerial body (Task Force), with cross-sectoral representation, to spearhead the implementation of AIRAWAT. The Task Force may include representation of both developer community and user domain experts of this infrastructure facility, in advisory capacity, to ensure that the design of the facility is robust and is truly reflective of the demands of the stakeholders and keeps innovating with the evolving nature of technology.

The proposed development of AIRAWAT would be in line with the approach the Government has taken of developing common public infrastructure and enabling the various stakeholders to leverage the public good to innovate and achieve the stated goals. This approach has led to India leapfrogging the world in the field of digital payments by building world's most advanced payment system, UPI. UPI, which was developed as a public good, in partnership with 12 banks, has now more than 143 banks live on it and has registered more than USD125bn in transactions since its inception in August 2016. UPI now constitutes more than 50% of all online payments, and has raced ahead of cards and other modes to become the most preferred payment method.

Developing AIRAWAT is expected to similarly invigorate the AI ecosystem in India, addressing the computing infrastructure needs of startups, academicians, researchers etc. As such, the AIRAWAT should be seen as an essential public good and funded by the Government. The necessary funding for AIRAWAT may be provided by supplementing funds under the NSM.

Structure of the Facility

As per the recommendations of NSAI, the key design philosophy for AIRAWAT shall be guided by the need to democratise access to AI computing infrastructure. As discussed earlier, the efforts in building AI compute capabilities hitherto have been rather limited in scale and scope, and have led to islands of modest excellence with capabilities of a few petaflops, providing limited access to the wider user base. AIRAWAT is expected to obviate the inefficiencies that result from decentralised small infrastructure, and provide efficiencies of scale and scope, by building one large, ambitious and common infrastructure that is accessible across India. A centralized facility is recommended for AIRAWAT in order to ensure increased accessibility and utilization as well as ability to support large scale and more diverse R&D projects as explained below. Table 3 below brings out the pros and cons of alternate structure for AIRAWAT.

Table 3: Comparison of different potential structures for AIRAWAT

Structure	**Comments**
Utilizing existing HPC infrastructure	• Current Installed Supercomputers are designed for HPC • Rigid to upgrade to AI workload: HPC overloaded • New initiatives e.g. NSM are also HPC focused
Creating a new decentralized facility	• Optimal only for small R&D • Collaboration / aggregation / workload distribution / administration challenges • Repetitive cost
Utilizing existing 'public cloud' infrastructure	• Data sharing concerns • Lack of clarity and policy on data security / privacy • Non-predictable and high bandwidth costs • Suitable for pay as you go and small instance requirements
Creating a new centralized facility	• No data sharing concerns • Reuse existing high bandwidth infra (e.g. National Knowledge Network) • Efficient utilization in multi-user and multi-tenant environment • Can support both small experiments as well as grand challenges / big data

Modes of Access to the Facility

HPCs are typically made available to users through a private network and require a degree of comfort in dealing with system level user interfaces for their access. HPC access mechanisms also lack the virtualization and management tools typically available with cloud services such as AWS, Azure or Google Cloud, and have limited features in scaling up resources effectively.

Increasingly, HPC providers both domestically and globally are thus transitioning to 'HPC as a Service' models which "focus on exposing HPC resources using elastic, on-demand cloud abstractions, aiming to combine the flexibility of cloud based models with the performance of HPC based systems".[11] This is evidenced in CDAC's current aim to transition to cloud-based access to

[11] Cloud Paradigms and Practices for Computational and Data-enabled Science Engineering, Parashar (2013)

HPC grids[12], and examples of larger supercomputing facilities being made available as cloud services (Japan's ABCI, for instance)[13].

It is proposed that AIRAWAT consider a similar approach of cloud-based access, given the often flexible requirements of AI compute tasks. With regard to broader considerations of access, AIRAWAT may be made accessible to users across the country through National Knowledge Network (NKN), for which NKN may need to be upgraded suitably.

Architecture of the Facility

From a technical specification perspective, the most important aspect is building an AI infrastructure that is scalable and flexible, and can cater to rapidly changing AI development landscape.

We are currently in the phase of narrow AI, defined by performance in a single domain with human or superhuman accuracy and speed for certain tasks, which have been broadly adopted in applications from facial recognition to natural language translation. We are just at the beginning of Broad AI, which encompasses multi-task, multi-domain, multi-model, distributed and explainable AI. Transfer learning and reasoning are central to expanding AI to small datasets. Reducing the time and power requirements of AI computing is fundamental to the development and adoption of Broad AI solutions, and thus will dictate the technical specifications of computing infrastructure being designed.

While the technical specifications for AIRAWAT will be evolved and designed through an open request for proposal process, it is recommended that the technical capabilities may be designed on the lines of the Summit and ABCI facilities. The broad specifications that may be considered for AIRAWAT architecture may include:

(a) Multi-tenant multi-user computing support

(b) Resource partitioning and provisioning, dynamic computing environment

(c) ML / DL software stack – training and inferencing development kit, frameworks, libraries, cloud management software.

(d) Support for varieties of AI workloads and ML / DL frameworks for user choices

(e) Energy-saving, high teraflops per watt per server rack space

(f) Low latency high bandwidth network

(g) Multi-layer storage system to ingest and process multi-petabytes of big data

(h) Compatibility with NKN (with upgrade to NKN, if needed)

---

[12] CDAC website
[13] ABCI website

Figure 2: *Proposed Architecture of AIRAWAT*

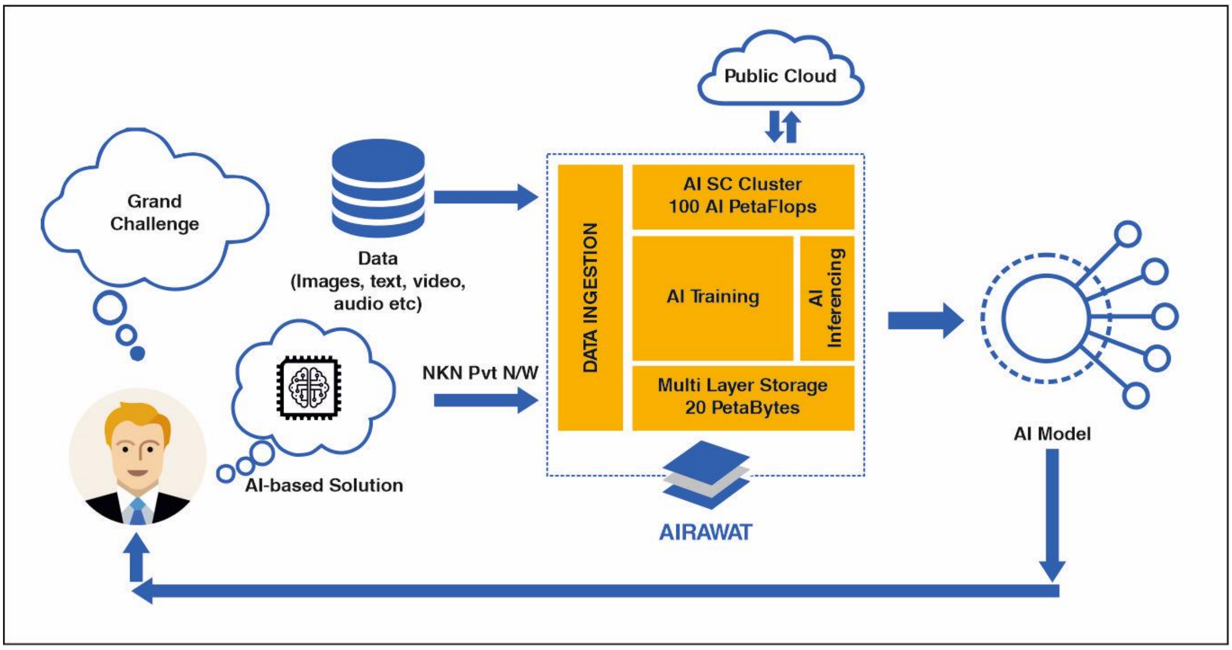

Source: NITI Aayog

The proposed architecture, with composite compute and storage infrastructure allows maintaining large data sets (thus eliminating the need for separate data centres and addressing data integrity concerns), and proximity of compute facility for efficient processing of data-intensive tasks viz. training of algorithms on large (both number and size) datasets.

The expected infrastructure, with capabilities of more than 100 peta flops (in the simplest sense, an AI flop is a measure of how fast a computer can perform deep neural network operations), would be more than the combined computing facility of top 20 supercomputers in India, and will put India on the global AI map, at par with the likes of Europe and Japan. Energy efficiency will be a key aspect of the facility, with the aim of putting AIRWAT in the list of top global green supercomputers. The facility would also enable storing of India's massive data sets from areas like healthcare, agriculture locally in a high throughput and efficient storage.

This new centralised AI infrastructure would alleviate any data sharing concerns (eliminating need to share data at multiple decentralised locations), is aimed at reusing existing high bandwidth infrastructure (e.g. NKN), is a better approach to utilization of computing resources in multi-user and multi-tenant environment, and has the scaling flexibility to include both small experiments as well as solving grand challenges / big data.

The use cases for AIRAWAT may vary from Big Data Analytics to specialised AI solutions across multiple domains viz. Healthcare (precision diagnostics, non-invasive diagnostics etc.), Agriculture (precision agriculture, crop infestations, advanced agronomic advisory etc.), weather forecasting, security and surveillance, financial inclusion and other services (fraud detection), infrastructural tools viz. NLP etc.

AIRAWAT: Key Governance Considerations

Ownership, Roles and Responsibilities

Development of AIRAWAT would need to address several issues relating to governance of the facility, including (a) ownership of the facility, (b) procedure for selection of the entity responsible for development and maintenance of the facility; and (c) roles and responsibilities for operation of the facility

Ownership of the Facility

Given the quantum of sunk investments required for building supercomputing facilities and their role in facilitating innovation and knowledge creation, development of AIRAWAT should be classified as public infrastructure requiring public funding. As noted above this would also be in line with the present policy of the Government towards building 'digital facilities' as public goods, such as the Unified Payment Interface (UPI).

Other factors in choosing between a privately owned infrastructure versus public ownership include concerns over data sharing, cost, and connectivity to existing public network infrastructure. These factors have been summarized in Figure 3 below.

Figure 3: *Ownership considerations for AIRAWAT*

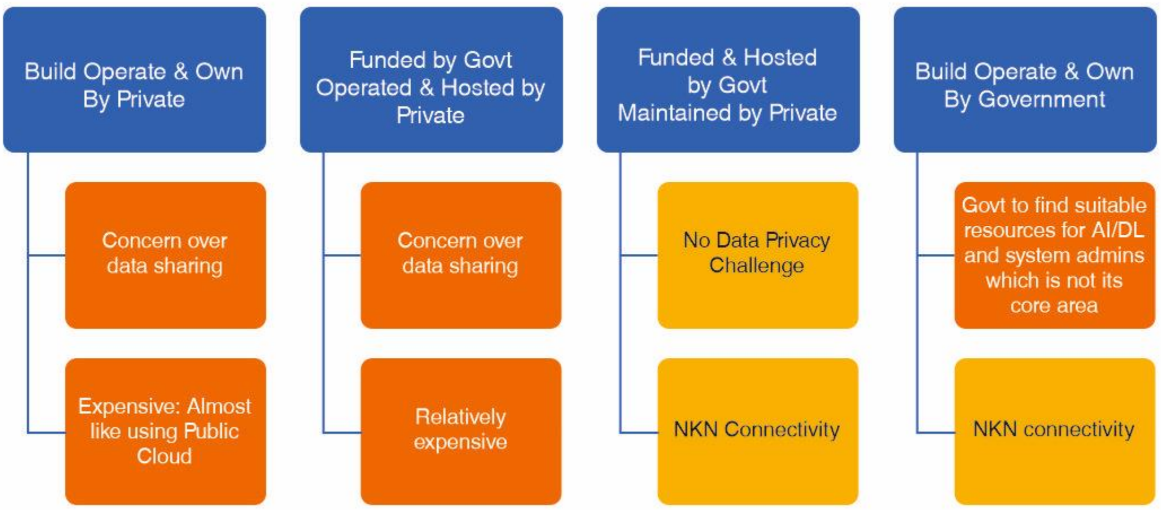

Source: NITI Aayog

In view of the above considerations, it is recommended that AIRAWAT adopt a mechanism where it is funded by the government and hosted at an academic institution (*Host Institute*).

Apart from state of the art, specialized hardware infrastructure, such as AI focused processors, which are currently only manufactured by select players across the world, including the most commonly used supercomputing processors by nVIDIA, Intel, and AMD, AIRAWAT would also need a robust software stack (including data management, workload optimization, automation tools) to maximize the efficiency of the processors. Thus, the development of this facility would primarily require engagement with a *System Integrator*, with demonstrated capability in setting up large computing facilities specialised for AI, who would design and implement the entire technology stack for the facility, and bring on board the suitable vendors for each layer of the stack. For example, the contract for building Japan's AI supercomputing facility, ABCI, was given to Fujitsu, a system integrator, which partnered with Intel, nVIDIA, and other vendors to develop the facility. The responsibility of upgrading and maintaining the facility on an ongoing basis will also be entrusted with the System Integrator with necessary skillset and experience. The operation and maintenance of the facility needs to be structured in a manner that creates sufficient incentive to leverage existing infrastructure

Procedure for selection of Host Institute and System Integrator of the facility

The Host Institute for AIRAWAT may be selected by a limited call for proposals from top-tier educational institutes, through a challenge method, based on demonstrated capability of hosting such an advanced computing facility and commitment to extend necessary support as may be required.

The System Integrator would be chosen through an open tendering process.

Roles and responsibilities for operation of AIRAWAT

A proposed organizational structure for governance of the facility that may be considered is given below:

Figure 4: Organizational structure for governance of AIRAWAT

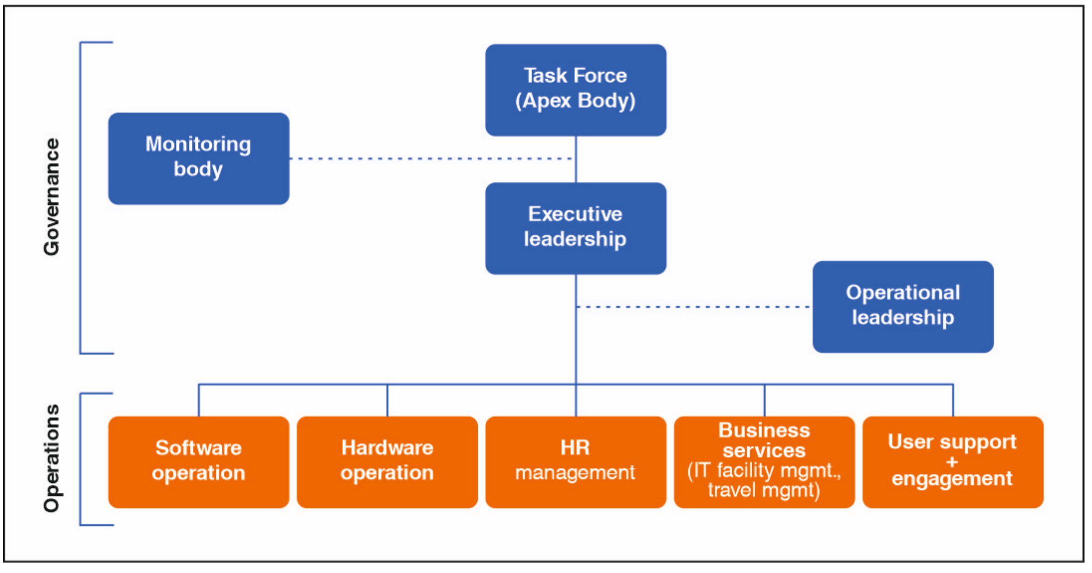

Source:
NITI Aayog

Table 4: Comparison of different potential structures for AIRAWAT

Body	**Function**	**Composition / Comments**
AIRAWAT Task Force (Apex Body)	• Overall governance of AIRAWAT • Approval of overall access and usage policy of AIRAWAT • Approval of pricing policy for access to AIRAWAT for non- research usage • Approval of large scale projects to be taken up by AIRAWAT facility • Review of financial performance of AIRAWAT on the basis of inputs of the Monitoring Body • Review of administrative performance of Host Institute on the basis of inputs of the Monitoring Body • Review of potential improvements to AIRAWAT functioning on the basis of report submitted by Advisory Body	• Inter-ministerial task force for implementation of AIRAWAT
AIRAWAT Monitoring Body	• Responsible for evaluation of AIRAWAT facility, Host Institute, and System Integrator, and submission of periodic report to the AI Task Force on their performance	• To be set up by the Task Force
AIRAWAT Executive Leadership	• Development of access and usage policy of AIRAWAT facility • Development of pricing policy of AIRAWAT facility	• Dean / Senior Faculty of the Host Institute, as specified in proposal for hosting of AIRAWAT, and approved by AI Task

	• Approval of small scale projects to be taken up using AIRAWAT facility • Participation in meetings of the AIRAWAT Advisory Body • Ensuring creation of linkages with CORE, ICTAI, Innovation Hubs and Moonshot Projects • Representation of AIRAWAT facility in international and domestic, academic and industry forums	Force

AIRAWAT
Operational Leadership
	• Day to day management of AIRAWAT operations, including management of System Integrator and Host Institute resources for below functions:
	  ○ Hardware operations
	  ○ Software operations
	  ○ User support and engagement
	  ○ Business services
	  ○ HR management
	• Representation of AIRAWAT facility in international and domestic, academic and industry forums
	• Organization of meetings of the AIRAWAT User Groups
		• Chief Operating Officer (COO) for AIRAWAT selected by the Host Institute and approved by AI Task Force
		• It is expected that additional resources be hired to support the COO and form the 'Office of the COO'. The Office of the COO will be funded by AIRAWAT's budget.

AIRAWAT User Groups	• To provide advice and feedback to AIRAWAT Steering Committee on the current and future state of AIRAWAT operations and	• All principal investigators and users on approved AIRAWAT user projects are AIRAWAT User Groups

	services ● Promote the effective use of the high performance computing facilities at AIRAWAT by sharing information about experiences in using the facility ● Serve in an advisory capacity to help determine the computational requirements and needs of the Government	members, and will remain so for 2 (two) years following the conclusion of their AIRAWAT Project

Administration of AIRAWAT facility will be a joint effort of the System Integrator and the Host Institute, with the following division of responsibilities:

1. System Integrator:
   a. responsible for the procurement and operation of hardware and software of the AIRAWAT facility
   b. maintenance and upgradation of AIRAWAT including server rack upgrades, software upgrades, facility cooling, etc.
2. Host Institute:
   a. responsible for the administration human resources of the AIRAWAT facility human resources, including HR management,
   b. delivery of business services such as IT management, travel administration, etc., as well as the delivery of user support and engagement services
   c. providing land and building for hosting AIRAWAT cater to a large part of India.

AIRAWAT: Financial Implications

The estimated financial outlay for building AIRAWAT will include the following components:

(a) Equipment (GPU / TPU supercomputers, storage, switches for internet connection)

(b) Facility setup and upgrade

(c) Recurring costs viz. maintenance, personnel, training workshops, contingency funds etc.

The equipment costs will have the following sub-components:

- AI specific processing units (could be GPUs, TPUs, as relevant)
- Other servers (data ingestion, cluster managers, inferencing, accelerators)
- Data Centres
- Software: for both hardware management and ML / DL
- Storage capabilities
- Network capabilities
- Service and support

The cost of these individual sub-items will best be discovered through an open bidding process.

Way Forward

Given the nascent stage of India's AI ecosystem, a dedicate cloud based computing infrastructure is needed to facilitate and speed up the AI research and solution development. AIRAWAT is envisioned to be a leading-edge AI computing technology platform, thus enabling the key players to bring an AI revolution in the country - students, researchers, startups, corporate and government organizations.

It is recommended that a Task Force, as discussed earlier, be set up immediately to oversee the development of AIRAWAT. The Task Force will need to seek funding for implementation and the timeline for setting up AIRAWAT is expected to be six months from the day financial approvals are received. The Task Force may call for proposals from system integrators through an open bid route, and the model request for proposal document prepared by NITI Aayog, with focus on outputs and outcomes, may be used for that. In parallel, the Task Force may seek interest from academic institutions ready to host AIRAWAT.

India has the potential to position itself among leaders on the global AI map, and AIRAWAT will be an important enabler in realising this aspiration.

Appendix I: Details of Summit and ABCI[14]

The Summit supercomputer at Oak Ridge National Laboratory, commissioned by the US Department of Energy, embodies multiple features of system-level purpose-built architecture for AI computation. The supercomputing facility, developed by IBM, was ranked the #1 most powerful supercomputer in the world in June 2018. The Summit architecture is designed not only for raw performance, but specifically tailored for AI workloads.

Key features of the Summit include 200 Petaflops of processing capability, 250 petabyte storage capacity and speed of 25 gigabytes per second between nodes. Summit employs multiple hardware and software approaches to address data transport, connectivity, and scalability. Summit's compute nodes each contain dual IBM POWER9 CPUs, six NVIDIA Volta GPUs, over half a terabyte of coherent memory (high bandwidth memory + DDR4) addressable by all CPUs and GPUs, plus 1.6TB per node of non-volatile RAM that can be used as a burst buffer or as extended memory. Second generation NVLink allows CPUs and GPUs to share data up to 4X faster than x86-based systems. Dual-rail Mellanox EDR InfiniBand interconnects, used for both storage and inter process communications traffic, deliver 200 Gb/s bandwidth between nodes.

AI Bridging Cloud Infrastructure (ABCI) supercomputer has been commissioned by the National Institute of Advanced Industrial Science and Technology (AIST) in Japan and is being integrated by Fujitsu. ABCI is aimed specifically to offer cloud access to compute and storage capacity for artificial intelligence and data analytics workloads.

[14] Reproduced mostly in original text from websites and various news releases and press clippings, including https://blog.mellanox.com/2017/11/what-does-it-mean-to-summit/ and https://www.nextplatform.com/2017/10/12/japans-abci-system-shows-subtleties-separating-ai-hpc/

Figure 5: *ABCI platform*

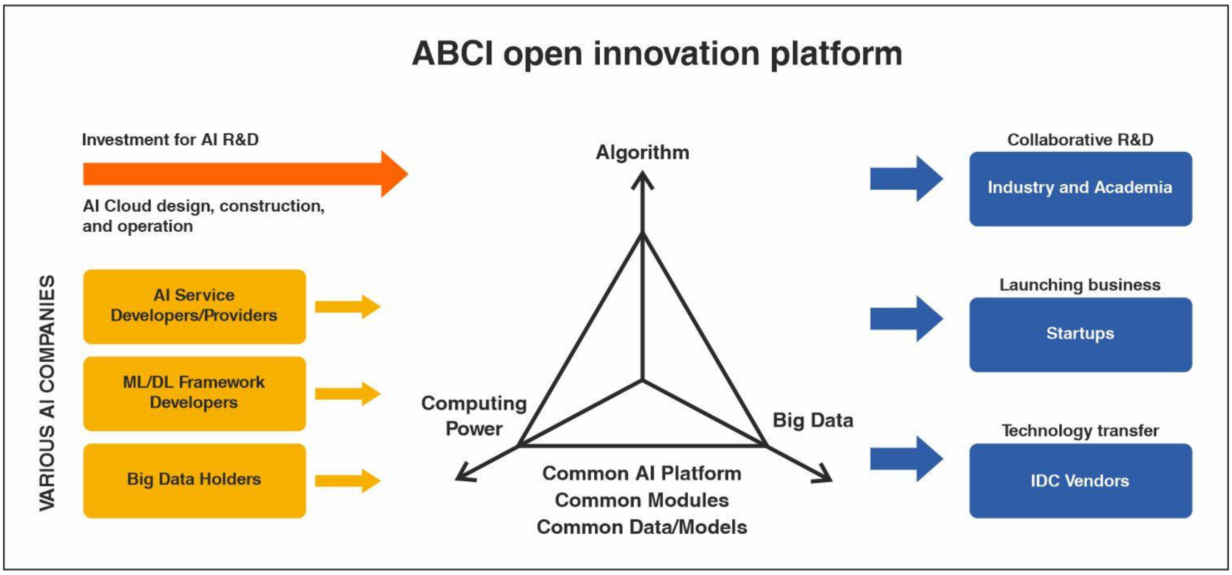

Source: TheNextPlatform.com

The ABCI system will be using the new "Volta" Tesla V100 GPU accelerators, which sport Tensor Core units that deliver 120 teraflops per chip for machine learning training and inference workloads. ABCI is aimed to deliver a machine with somewhere between 130 petaflops and 200 petaflops of AI processing power, which means half precision and single precision for the most part, with a power usage effectiveness (PUE) of somewhere under 1.1, which is a ratio of the energy consumed for the data center compared to the compute complex that does actual work. The system is expected to have about 20 PB of parallel file storage and, with the compute, storage, and switching combined, burn under 3 megawatts of juice.

The ABCI system will be comprised of 1,088 of Fujitsu's Primergy CX2570 server nodes, which are half-width server sleds that slide into the Primergy CX400 2U chassis. Each sled can accommodate two Intel "Skylake" Xeon SP processors, and in this case AIST is using a Xeon SP Gold variant, presumably with a large (but not extreme) number of cores. Each node is equipped with four of the Volta SMX2 GPU accelerators, so the entire machine has 2,176 CPU sockets and 4,352 GPU sockets. The use of the SXM2 variants of the Volta GPU accelerators requires liquid cooling because they run a little hotter, but the system has an air-cooled option for the Volta accelerators that hook into the system over the PCI-Express bus. The off-the-shelf models of the CX2570 server sleds also support the lower-grade Silver and Bronze Xeon SP processors as well as the high-end Platinum chips, so AIST is going in the middle of the road. There are Intel DC 4600 flash SSDs for local storage on the machine. It is not clear who won the deal for the GPFS file system for this machine, and if it came in at 20 PB as expected.

As per Fujitsu, the resulting ABCI system will have 37 petaflops of aggregate peak double precision floating point oomph, and will be rated at 550 petaflops, and 525 petaflops off that comes from using the 16-bit Tensor Core units that were created explicitly to speed up machine learning workloads. That is a lot more deep learning performance than was planned, obviously.

AIST has raised USD172 million to fund the prototype and full ABCI machines as well as build the new datacenter that will house this system. About USD10 million of that funding is for the datacenter. The initial datacenter setup has a maximum power draw of 3.25 megawatts, and it has 3.2 megawatts of cooling capacity, of which 3 megawatts come from a free cooling tower assembly and another 200 kilowatts comes from a chilling unit. The datacenter has a single concrete slab floor, which is cheap and easy, and will start out with 90 racks of capacity – that's 18 for storage and 72 for compute – with room for expansion.

Figure 6: ABCI Cooling Infrastructure

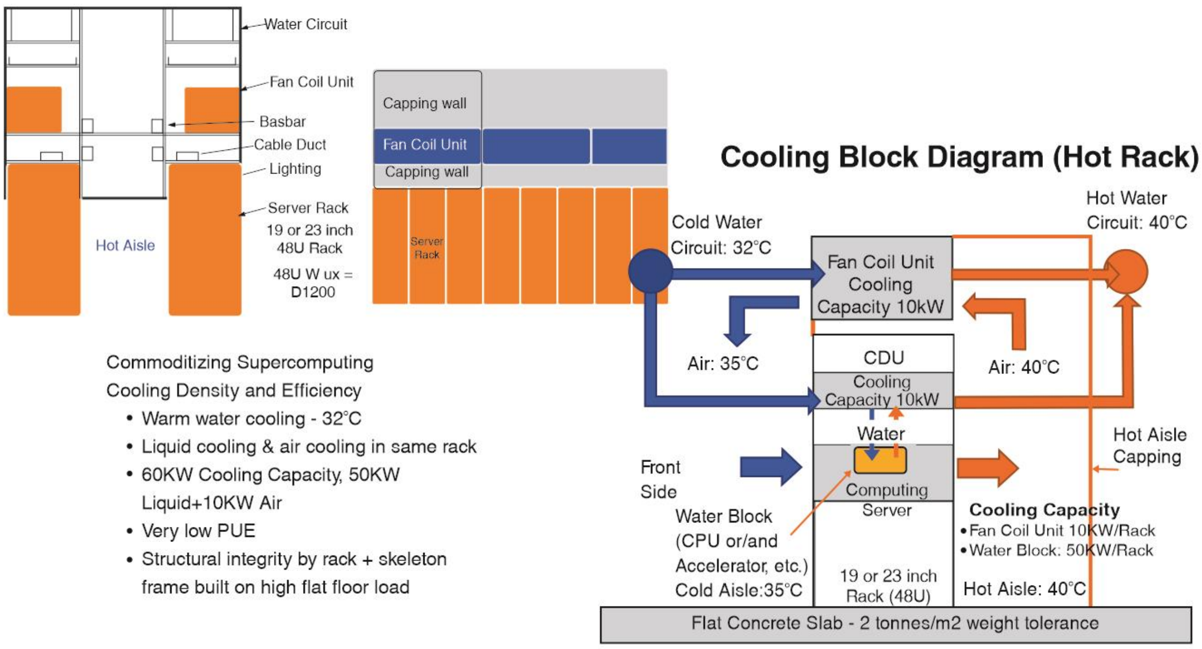

Source: TheNextPlatform.com

One of the key features of the ABCI design is the rack-level cooling, which includes 50 kilowatts of liquid cooling and 10 kilowatts of air cooling. The liquid cooling system uses 32 degree Celsius water and 35 degree Celsius air. The water cooling system has water blocks on the CPUs and GPUs and probably the main memory, and there is hot aisle capping to contain it and more efficiently remove its heat.

Figure 7: *ABCI architecture*

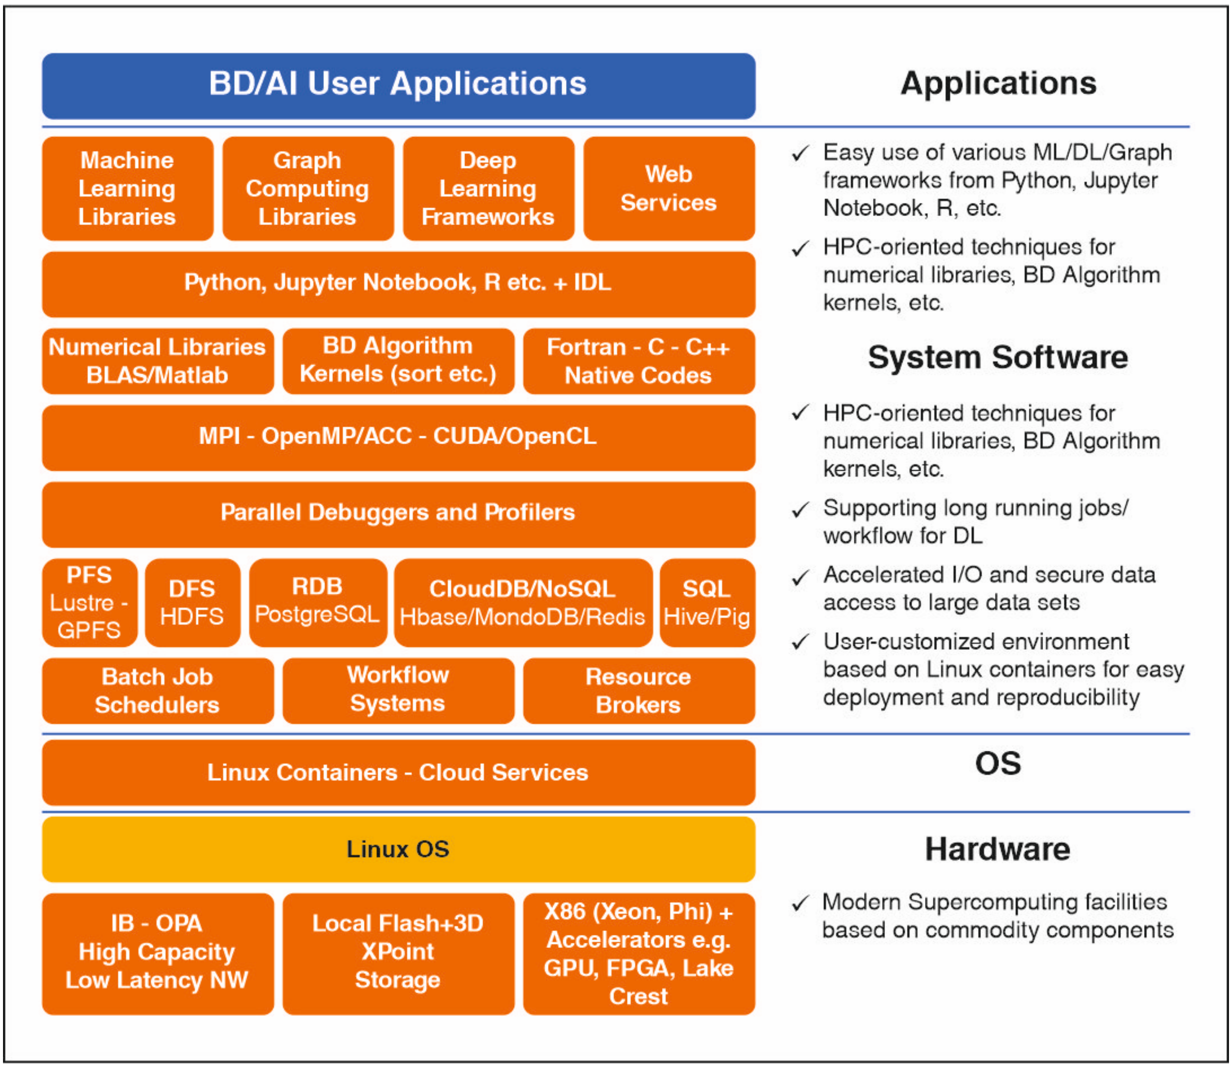

Source: TheNextPlatform.com

The HDFS file system that underlays Hadoop data analytics is a key component of the stack, as are a number of relational and NoSQL data stores. And while there is MPI for memory sharing and the usual OpenACC, OpenMP, OpenCL, and CUDA for various parallel programming techniques, and some familiar programming languages and math libraries, the machine learning, deep learning, and graph frameworks running atop the ABCI system make it different, and also drive a different network topology from the fat trees used in HPC simulations where all nodes sometimes have to talk to all other nodes.