b The general purpose CPU part of a compute node consists of multiple CPU cores connected by a hardware shared memory system. Each CPU core must be based on the 64-bit extended IA32 architecture.

c The computation accelerator part must be constructed by a set of multiple computation cores connected by a hardware shared memory system. This part must be connected to the general purpose CPU part to perform high speed computing coordinating with it.

d Theoretical peak performance (double precision floating point arithmetic) of GP must be 2PF or more, and total memory bandwidth of GP must be 150TByte／sec or more. Each node must have two general purpose CPUs. Each node must be connected to the internal network with 100Gbps or higher and to the storage network 100Gbps or higher.

e Theoretical peak performance of double precision floating point arithmetic of ACC must be 4.5PF or more, and total memory bandwidth of ACC must be 600TByte／sec or more. Each node must have one or two general purpose CPUs and four computation accelerators in maximum. Each node must be connected to the internal network with 50Gbps or higher for each accelerator and to the storage network 50Gbps or higher for each general purpose CPU.

f Theoretical peak performance of single precision floating point arithmetic of DA must be 16PF or more and theoretical peak performance of half precision (16-bit) floating point arithmetic of DA must be 120PF or more. Each node must have one or two general purpose CPUs and eight computation accelerators in maximum. Each node must be connected to the internal network with 25Gbps or higher for each accelerator and to the storage network 50Gbps or higher for each general purpose CPU.

C Requirements for Storage

a Storage of the system consist of high-speed internal storage, large-capacity internal storage, shared storage, and cold storage. Each storage must be highly reliable and fault-tolerant.

b High-speed internal storage must be constructed by a set of high-speed internal storage nodes including NVMe SSD with a total capacity of 8PByte or more. It must allow the access from the compute nodes with equal to or more than 2TByte／sec bandwidth using NVMe over Fabrics. In addition, based on the isolation setting which can be configured only from the management node, it allows to restrict the access to specific storage nodes or drives, etc., for each compute node.

c Large-capacity internal storage must be constructed by a set of large-capacity internal storage nodes with a total capacity of 15PByte or more. It must allow the access from the compute nodes with equal to or more than 100GByte／sec bandwidth using iSCSI and so on. In addition, based on the isolation setting which can be configured only from the management node, it allows to restrict the access to specific storage nodes or drives, etc., for each compute node.

d A shared storage must be an object storage compatible with Amazon S3 and a total capacity of 30PByte or more. It must be connected to the external router with equals to or more than 100GByte／sec.

e A cold storage must be configured by magnetic tapes or optical disks. It must have a total capacity of 15PByte or more by the storage media, and allows to add storage media with equals to or more than 50PByte without expansion of the cold storage system.

D Requirements for Internal Network

a Compute Nodes (GP, ACC, DA) must be connected to each other and to the external routers using Spine-Leaf type IP network. The interfaces of compute nodes are Ethernet. The switches of the Infrastructure System must be configured by only the management nodes. The switches can separate the network attached to the compute nodes and avoid inter-communication among the compute nodes.

b Any network can attach to any compute nodes.

c RDMA communications must be available among the compute nodes.

E Requirements for Storage Network : The storage network must have a full bi-section bandwidth and it must be configured by only the management nodes. The network between compute node and storage node must be configured as separated network not to communicate each other.

F Requirements for External Router

a The one pair of external router must be composed by two or more routers for redundancy.

b On the routers, IPv4, IPv6, and MPLS routing must be available and all routing control methods commonly used in backbone routing including BGP, OSPF, and OSPFv3 must be supported. The routers can transfer data at wire rate with the minimum Ethernet frame size.

c The routers must connect to the External Network using two or more sets of 100Gbps Ethernet.

d The routers must connect to each Spine switches using two or more sets of the Ethernet links with 200Gbps or more bandwidths.

e The routers must connect to the other two supercomputer systems in Kashiwa II campus of The University of Tokyo, “AI Bridging Cloud (ABCI)” and “Integrated Supercomputer System for Simulation, Data, and Learning”, with 100Gbps or more bandwidths, and must configure the communication paths to these two systems.

f The routers have the function of VLAN, VLAN tag translation, and link aggregation.

G Requirements for Portal Server and Management Server

a Portal Servers must be provided for operating the web-based portals which are available for users of Infrastructure System to require and control the computing resources assigned to themselves. Portal Servers must be composed by two or more servers for redundancy.

b Management Servers must be provided for managing the resources of entire Infrastructure System as follows : Virtualized Resource Management Server, Operation Management Server, Authentication Server, Security Log Storage Server. Each Management Server must be composed by two or more servers for redundancy.

H “Infrastructure System” must satisfy the following software specifications :

a All compute nodes must be capable of running Virtual Machines (VM) and／or containers.

b The portal server and the management server must be Linux OSes.