DCI Product Concept Paper

Classification: APPROVED REFERENCE DOCUMENT

Confidentiality: PUBLIC

Version 1

October 2023

[DCI PCP]

Legal

THIS DOCUMENT HAS BEEN DESIGNATED BY THE INNOVATIVE OPTICAL AND WIRELESS NETWORK GLOBAL FORUM, INC. (“IOWN GLOBAL FORUM”) AS A APPROVED REFERENCE DOCUMENT AS SUCH TERM IS USED IN THE IOWN GLOBAL FORUM INTELLECTUAL PROPERTY RIGHTS POLICY (THIS “REFERENCE DOCUMENT”).

THIS REFERENCE DOCUMENT IS PROVIDED ”AS IS” WITH NO WARRANTIES WHATSOEVER, WHETHER EXPRESS, IMPLIED, STATUTORY, OR OTHERWISE, INCLUDING WITHOUT LIMITATION ANY WARRANTY OF MERCHANTABILITY, NONINFRINGEMENT OF THIRD PARTY RIGHTS, TITLE, VALIDITY OF RIGHTS IN, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, REFERENCE DOCUMENT, SAMPLE, OR LAW. WITHOUT LIMITATION, IOWN GLOBAL FORUM DISCLAIMS ALL LIABILITY, INCLUDING WITHOUT LIMITATION LIABILITY FOR INFRINGEMENT OF ANY PROPRIETARY RIGHTS AND PRODUCTS LIABILITY, RELATING TO USE OF THE INFORMATION IN THIS REFERENCE DOCUMENT AND TO ANY USE OF THIS REFERENCE DOCUMENT IN CONNECTION WITH THE DEVELOPMENT OF ANY PRODUCT OR SERVICE, AND IOWN GLOBAL FORUM DISCLAIMS ALL LIABILITY FOR COST OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, LOST PROFITS, LOSS OF USE, LOSS OF DATA OR ANY INCIDENTAL, CONSEQUENTIAL, DIRECT, INDIRECT, PUNITIVE, EXEMPLARY, OR SPECIAL DAMAGES, WHETHER UNDER CONTRACT, TORT, WARRANTY OR OTHERWISE, ARISING IN ANY WAY OUT OF USE OR RELIANCE UPON THIS REFERENCE DOCUMENT OR ANY INFORMATION HEREIN.

EXCEPT AS EXPRESSLY SET FORTH IN THE PARAGRAPH DIRECTLY BELOW, NO LICENSE IS GRANTED HEREIN, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS OF THE IOWN GLOBAL FORUM, ANY IOWN GLOBAL FORUM MEMBER OR ANY AFFILIATE OF ANY IOWN GLOBAL FORUM MEMBER. EXCEPT AS EXPRESSLY SET FORTH IN THE PARAGRAPH DIRECTLY BELOW, ALL RIGHTS IN THIS REFERENCE DOCUMENT ARE RESERVED.

A limited, non-exclusive, non-transferable, non-assignable, non-sublicensable license is hereby granted by IOWN Global Forum to you to copy, reproduce, and use this Reference Document for internal use only. You must retain this page and all proprietary rights notices in all copies you make of this Reference Document under this license grant.

THIS DOCUMENT IS AN APPROVED REFERENCE DOCUMENT AND IS SUBJECT TO THE REFERENCE DOCUMENT LICENSING COMMITMENTS OF THE MEMBERS OF THE IOWN GLOBAL FORUM PURSUANT TO THE IOWN GLOBAL FORUM INTELLECTUAL PROPERTY RIGHTS POLICY. A COPY OF THE IOWN GLOBAL FORUM INTELLECTUAL PROPERTY RIGHTS POLICY CAN BE OBTAINED BY COMPLETING THE FORM AT: www.iowngf.org/join-forum. USE OF THIS REFERENCE DOCUMENT IS SUBJECT TO THE LIMITED INTERNAL-USE ONLY LICENSE GRANTED ABOVE. IF YOU WOULD LIKE TO REQUEST A COPYRIGHT LICENSE THAT IS DIFFERENT FROM THE ONE GRANTED ABOVE (SUCH AS, BUT NOT LIMITED TO, A LICENSE TO TRANSLATE THIS REFERENCE DOCUMENT INTO ANOTHER LANGUAGE), PLEASE CONTACT US BY COMPLETING THE FORM AT: https://iowngf.org/contact-us/

Copyright © 2023 Innovative Optical Wireless Network Global Forum, Inc. All rights reserved. Except for the limited internal-use only license set forth above, copying or other forms of reproduction and/or distribution of this Reference Document are strictly prohibited.

The IOWN GLOBAL FORUM mark and IOWN GLOBAL FORUM & Design logo are trademarks of Innovative Optical and Wireless Network Global Forum, Inc. in the United States and other countries. Unauthorized use is strictly prohibited. IOWN is a registered and unregistered trademark of Nippon Telegraph and Telephone Corporation in the United States, Japan, and other countries. Other names and brands appearing in this document may be claimed as the property of others.

Contents

List of Figures

List of Tables

1. Introduction

Today's communication and computing infrastructures are reaching their limits. This prevents the realization of emerging use cases. The IOWN GF has been investigating emerging use cases that have demanding requirements for communication and computation: gap analyses show that current infrastructures cannot support these use cases due to inherent structural shortcomings.

The key issue in achieving the required performance and service quality is that current computation infrastructures are based on the premise that data is stored and processed locally in each server. The reason for this approach is that in general, in most other processing paradigms, the communication performance of the interconnects between servers is too slow and too unstable. As a result, the construction of these *hardware-centric* infrastructures enforces specific ways for how data must be handled, effectively treating all data in a similar manner.

In contrast, IOWN GF use cases demand that infrastructures need to have different ways to transfer, process, and store data based on the type of data. Recently, concepts that focus on the needs of the data first instead of the needs of the hardware are called *data-centric*.

As a solution to enable future use cases, the IOWN GF takes this line of data-centric thinking further: transforming today's segmented information processing landscape into a single unified data-centric infrastructure (DCI), from which users can request networking and computing capabilities to deploy applications with tailored communication stacks. DCI embraces data-centric concepts throughout all architectural levels, from the setup of individual processing elements up to global-scale optical path provisioning. To realize DCIs, key technologies that will enable this approach include flexible wide and local area optical networking, server hardware disaggregation, and geographically aware joint network and compute management.

This paper introduces the overall concept of IOWN GF DCI systems from a product perspective and outlines the benefits of using DCI for users and operators. While the details of the individual technical aspects are still underway at the IOWN GF, this text also gives an overview of the major technical means to implement DCI systems that are fit for actual business deployments, so that these systems can fully meet the requirements of the IOWN GF use cases.

This paper is structured as follows:

- First, the motivation for DCI in the context of IOWN GF use cases is discussed in Section 2.
- After that, the IOWN GF DCI solution is outlined in Section 3.
- Then, the main benefits for users and operators of such systems are highlighted in Section 4.
- Finally, the individual technologies needed for the realization of DCI systems are revisited in greater detail in Appendix A, an extensive example of a DCI system is shown in Appendix B, and requirements of IOWN GF use cases are summarized in Appendix C.

2. Current infrastructure challenges and the need for DCI

With ever-evolving technology, new use cases are emerging. The IOWN Global Forum is investigating and describing these use cases, including highly dynamic fronthaul setup for virtual radio access networks using IOWN for Mobile Networks [IMN], AI-Integrated Communication Interactive Live Music [ILM], Cyber-Physical-Systems industry management [RCRI], as well as disaggregated data centers and related on-demand resource-sharing platforms to realize data spaces that provide confidentiality, integrity, and compliance.

To improve cost efficiency, most of these use cases greatly benefit from dynamic establishment of high-quality communication between compute resources in different data centers and/or other locations of interest, that is, across multiple organizations. Further, the use cases listed earlier often have strict latency requirements for the communications between and within the heterogeneous computing units themselves, such as CPUs, GPUs, FPGAs, DPUs/IPUs, and NVMs. Use case requirements are further discussed in Appendix C.

Popular orchestration frameworks exist that could also be used to realize geographically distributed computing and edge computing. However, common hurdles remain toward supporting the demands of IOWN GF use cases for high usability, high efficiency, long geographic distances, and high performance all at the same time:

- **Protocol processing overhead:** Data streams from remote sites to local accelerators must pass through the host server protocol stacks and policy enforcement functions. These functions are typically implemented at least partially in software on the host CPU and are based on TCP and UDP protocols. Therefore, these stacks cause non-negligible overhead and have non-deterministic processing time. Such a CPU-centric packet processing flow may be acceptable for today’s connections over WANs. However, this protocol stack structure comprising CPU-based processing is driving up energy consumption. Furthermore, for communication in geographically distributed settings, this processing structure prevents the effective use of fully hardware-accelerated high-performance protocol stacks beyond CPU-based TCP and UDP, such as remote direct memory access (RDMA) protocol stacks.
- **Segregated WAN and compute:** In current infrastructures, the management of WAN resources and the management of compute resources are separate: WAN management frameworks allow users to create connections between locations such as data centers. However, the choice of servers that can be directly connected via the WAN is limited if such functionality is available at all. In turn, computing resource management frameworks offer rich functionality to configure and provision server hardware but can only configure high-quality communication between servers on a limited local scale. Therefore, in general, high-performance connections with guaranteed QoS can only be established either on a very high level between data center gateways, or on a very low level between resources in the same data center rack or in nearby racks, but not across these levels. This segregation impedes joint resource allocation and therefore high-performance communication across WANs.
- **Resource allocation based on pre-configured servers:** For IOWN GF use cases, using pre-assembled servers causes many additional issues: For users, these include cost increases for the servers, as typically, more than the required capabilities of the server are put in use. For operators, fewer customers are served than would be possible with more flexible resource allocation. Moreover, idle hardware still wastes energy, thus adding to the overall cost. For IOWN GF use cases, rigid resource allocation is particularly an issue for memory, storage, and accelerator devices.
- **Inflexible WAN Quality-of-Service (QoS):** Current communication infrastructure management frameworks for connections between servers across WANs either provide highly dynamic connection setup (with only classic packet-based best-effort transmission quality), or provide connections with transmission quality at the level of optical fiber channels including coordination between different organizational entities (with dynamic connection setup only possible between limited groups of predetermined servers), but not both at the same time.

- **Lack of user-friendly security mechanisms:** Security standards and implementations that apply to high-quality connectivity and sovereign computing already exist. However, with existing approaches, realizing complete data protection throughout the whole data lifecycle is highly complex.

Due to these obstacles, fulfilling functional requirements, such frameworks are ultimately bound by the performance limitations underlying today’s network infrastructures. As a result, today’s resource “clouds” are often confined to a single data center location.

In the end, due to the inability of today’s systems to fulfill requirements on strictly bounded latency and packet loss, performance stability, and guaranteed bandwidth across geographic locations, these IOWN GF use cases cannot be implemented in an economically feasible manner.

Simply upgrading current infrastructures in the same manner as in the past may keep yielding more throughput and more computational power. However, this will still not enable the high-performance, end-to-end communication capabilities that future use cases require. In addition, upgrading infrastructures in this way as in the past will further increase energy consumption, eventually making their power requirements untenable.

As a result, a new approach to networking and computing is needed that breaks the isolation of today’s cloud silos at multiple levels and fully integrates the networking and computing worlds.

The resources an application can access using current infrastructure are limited to, for example, those located within the same data center, in nearby servers, or to resources within the same server. Leveraging the IOWN GF’s use cases, such restrictions will be reduced, and resources will be sensibly accessible from a wider geographic area.

This will, in turn, fulfill users' needs to elastically reserve resources across edge and center with high quality-of-service guarantees for other emerging use cases. At the same time, such an approach that increases the overall infrastructure efficiency reduces both the costs for users and energy consumption. IOWN GF DCI systems are the solutions that will make the above possible.

3. Outline of the IOWN GF DCI solution

This section will outline how DCI systems are expected to overcome the hurdles as mentioned in Section 2. Readers interested in further details can find a more extensive overview in Appendix A.

DCI systems are expected to employ significant innovations. Taken together, these innovations will allow to overcome the hurdles highlighted in Section 2. These innovations include:

Processing and communication efficiency improvements will be realized by facilitating the use of accelerator technology to improve performance and reduce energy consumption:

- **High-performance protocols:** Software deployment will be integrated with high-performance protocols beyond TCP/UDP such as RDMA across geographic locations to fully utilize fiber-level connections.
- **Autonomy:** Accelerators, memories and storage, and other dependent or passive devices will be able to act autonomously and handle network transfers on their own to reduce communication overhead.

Joint network and compute management will be realized by providing resource abstractions and orchestration to break up the confines of cloud silos:

- **Combined network and compute orchestration:** Network and compute device configuration will be sufficiently abstracted for high-level orchestration to enable joint control of all involved devices.
- **Geographical placement optimization:** For application deployment, DCI orchestration will consider both network and compute requirements together according to the constraints provided by users.

Flexible resource usage will be promoted by resource-sharing software frameworks and hardware pools for logical server composition to reduce total required hardware, energy consumption, and cost:

- **Borderless sharing:** DCI systems will facilitate software frameworks for sharing remote pooled resources across networks among applications, servers, locations, and organizations to improve resource utilization.
- **Disaggregated resources:** DCI resource management will be integrated with hardware frameworks that compose on-demand logical servers according to user requirements from large hardware pools.

Flexible WAN-scale connections with QoS will be enabled by managing QoS for all devices and standardized inter-organization coordination:

- **All-QoS:** All buses and networks between processors and endpoints of dedicated optical paths will support QoS management to connect elements with full line speed even over long-range paths.
- **Multi-organization:** High-level orchestration systems will use standardized interfaces to coordinate both network and compute resource provisioning among legal entities or their departments.

Wide-scale security that can maintain data sovereignty: Security that can maintain data sovereignty throughout the data lifecycle, even for cross-organizational data distribution, will be facilitated by DCI systems and integrated with high-level orchestration to automate security configuration.

These individual means taken together remove the hurdles as introduced in Section 2. A more detailed treatment of each technique can be found in Appendix A.

4. Benefits for users and operators

As highlighted in Section 3, DCI systems increase overall system efficiency, enable joint optimization of network and compute domains, and enable flexible QoS-managed connections at WAN scale. With these advantages at hand, such a complete DCI system provides the following three main benefits:

- **Users can agilely deploy their geographically distributed applications with demanding performance constraints at nation-scale and beyond:** DCI systems offer high-performance interconnectivity on demand. Provisioning links between individual devices at arbitrary locations with 100+Gbps bandwidth and fiber-class latency together with computing hardware will only take seconds. Workloads are executed on the processor architecture(s) as designated by users. Accelerators become first-class resources, can be requested stand-alone, and can serve as direct endpoints for such connections, providing sufficient security.
- **Such advanced application deployment will be simple and natural for both users and operators:** Tenants can request such network and compute resources in tandem via a single unified interface. The DCI system will take care of coordination among multiple organizations, including hardware configuration parameter settings and monitoring configuration; therefore, users do not need to consider these complex factors. Operators can choose their optimal balance of business privacy and system efficiency when exchanging resource inventory information for cross-organization provisioning.
- **At the same time, the energy efficiency and carbon footprints of applications greatly benefit from the same functionality that boosts agility and usability:** Lowering the hurdle to use accelerators and shim protocol stacks such as RDMA improves base energy efficiency and suppresses idle resource provisioning. In a world without resource silos where geographically-distributed deployment is the norm, data and computation can be directed flexibly to where Green Energy is available.

Taken together, these three benefits enable users and operators to realize the future use cases as highlighted in Section 2: Flexible resource sharing and joint management for compute and network resources from server to WAN results in higher hardware utilization and higher energy efficiency. This improves cost-efficiency so that these use cases become not only technically possible, but also economically viable. The IOWN Global Forum is continuing work to make these benefits a reality.

Appendix A. DCI techniques

This appendix revisits the individual techniques that DCI systems rely on to improve efficiency, combine network and compute management, improve utilization, and provide WAN-scale QoS management in more detail.

Processing and communication efficiency improvements will be realized by combining high-performance communication and autonomous devices:

- **Long-range RDMA and other high-performance protocols:** With the QoS guarantees that DCI systems provide, RDMA finally becomes viable for connections to other locations across WAN networks. To cater to the future needs of developers, RDMA libraries will abstract different processor architectures, handle communication buffer management, and be straight-forward to use. For extreme use cases, even layer-0 optical communication between components and beyond data centers comes within reach. Restructuring protocol stacks in such a way will improve performance, scalability, and energy consumption.
- **Autonomous devices:** With the increased viability of RDMA, the need for devices such as accelerators, memories, and storage to act autonomously and to perform communication protocol processing on their own will emerge. This will be necessary to suppress overheads and to directly ingest data streams with extremely high bandwidth at low latency.

Joint network and compute management will be realized by providing resource abstractions and orchestration to tear down the confines of cloud silos:

- **Unified network and compute device configuration and abstraction:** DCI systems will comprise a wide variety of network and compute hardware to provide capabilities to users in an as-a-service fashion. For example, such hardware may include server system buses, network interface cards, intra-data center networks, data center gateways, and wide area networks, as well as GPUs, FPGAs, IPUs/DPUs, ASICs, and other domain-specific devices, respectively. A new generation of hardware control frameworks will abstract the configuration and start-up procedures of these devices from high-level orchestrators. This way, users and orchestrators do not need to care about low-level device configuration of network devices and can transparently provide initialization and start-up information for heterogeneous computing hardware. Logically, DCI systems will be structured accordingly: all computing and network resources except WANs are organized into clusters (DCI Clusters) that each have their own gateway (DCI Gateway) to access WANs. Each cluster is managed by a single controller (DCI Cluster Controller). These controllers abstract cluster hardware toward high-level orchestrators (DCI Infrastructure Orchestrators). WANs, too, is abstracted in a similar manner. All clusters and WANs in a DCI System are controlled by multiple peered high-level orchestrators.
- **Topology-aware and constraint-based orchestration:** DCI systems empower users to deploy geographically distributed applications with high-quality communication. The deployment requires orchestration systems that can provision both network and compute resources across edge and center locations at the same time. Furthermore, provisioning in the true clouds of the future gives more freedom for resource allocation to optimize utilization. Orchestrators will need to be able to automatically select suitable data centers and network paths based on the requirements of the user application. Orchestrators also allow geographic constraints, e.g., to designate customer sites or to fulfill regulatory requirements.

Flexible resource usage will be promoted by resource-sharing software frameworks and hardware pools supporting logical server composition to reduce total required hardware and reduce energy consumption:

- **Highly efficient resource sharing between application, servers, and locations via software:** DCI systems will provide the necessary means for software frameworks to request computing resources across multiple geographical locations and set up connections between these locations. Future software frameworks are expected to use these connections to implement the sharing of resources such as CPUs, GPUs, memories, and accelerators among applications, servers, and locations to improve resource utilization in a highly agile and dynamic manner and across networks.

- **Hardware pools, large logical system buses, and composable servers:** Hardware in DCI systems is expected to be installed and managed in hardware pools. From these pools, users can configure logical system buses to compose logical servers from one or multiple components, essentially allowing users to transparently build their own logical servers on-the-fly according to the needs of existing software. Being able to provision logical servers as small as a single accelerator card prevents partial overprovisioning of server hardware and lowers the hurdle to use accelerator hardware. Furthermore, installing hardware such as CPUs, GPUs, DPUs/IPUs/NICs, memory, and storage in pools allows to group identical hardware together physically, thus simplifying enclosures, maintenance, and hardware upgrades.

Flexible WAN-scale connections with QoS will be enabled by managing QoS for all devices and standardized inter-organization coordination:

- **QoS-aware interconnects and networks:** DCI system implementations will use QoS-aware high-performance network products for connectivity. For long-range WAN communication outside of data centers, the IOWN GF is developing the Open All-Photonics Network (Open APN). Inside data centers and customer premises, implementers are expected to use QoS-aware short-range networks and bus interconnects that are sufficiently performant to forward on Open APN connections at full rate to individual endpoints. Advanced gateways will either forward optical signals directly, adapt between long-range and short-range transmission methods, or forward to multiple hosts on the packet level to make true end-to-end, high-quality QoS connectivity possible.
- **Standardized interfaces for cross-organization provisioning:** Creating complex geographically distributed applications that span cities or countries may involve provisioning and configuring resources through multiple organizations. These organizations may include the owners of co-located servers, multiple data center operators, multiple WAN operators, and others. For operation efficiency, all phases of deployment must be fully automated, including the analysis of the user request, the negotiation and agreement about technical and business parameters, the provisioning and configuration of resources, and the final application deployment. Cross-organization coordination will be performed via standardized orchestration interfaces. Via these interfaces organizations will only need to provide peers with an abstract view of their resource inventory with a configurable degree of privacy.

Security at WAN scale for communication and computing will be available to users by default techniques provided in a transparent fashion. To achieve this support, DCI systems will facilitate the implementation of software and hardware frameworks that provide functionality to set up security for high-quality connections across WANs as well as secured computing spaces for sovereign computing. Apart from generic support, DCI systems will provide the means to realize security tailored to specific needs on a per-application basis.

The IOWN GF is continuing activities to outline how the techniques introduced above can be put together to realize next-generation infrastructures. These infrastructures will then be able to support new classes of use cases demanding geographically distributed high-performance processing and communication.

Appendix B. Example of a DCI system

Previous sections described the overall structure and components of DCI systems in an abstract manner. This appendix outlines a still simplified but more concrete example of a DCI system in Figure B-1.

The main components of the example DCI system that are shown here are multiple DCI Clusters and Open APN networks.

The Open APN networks connect the DCI Gateways inside the DCI Clusters to each other. DCI infrastructure orchestration (not shown) may coordinate multiple Open APN networks. This enables the establishment of fully QoS-managed connections or even direct optical paths between DCI Cluster resources across multiple Open APN networks.

The DCI Clusters comprise the computing resources. DCI Clusters may be structured differently internally depending on their purpose. Different types of possible DCI Clusters are illustrated:

- Large general-purpose DCI Clusters housing a variety of hardware components in pools for dynamic logical server composition are depicted top-left in the figure. One or multiple such clusters could replace today's data centers. The DCI Cluster illustrated here houses multiple hardware pools, possibly optimized for different workloads. Such hardware pools could, for example, be realized by means such as dedicated system bus fabrics or other system bus extension hardware.
- Smaller edge-site DCI Clusters tailored for VRAN and mobile edge computing (MEC) are shown at the bottom-left. Such clusters are expected to contain use case specific hardware, such as FPGA or GPU cards for eCPRI processing. Additionally, to make optimal use of scarce hardware under limited space and power availability, these resources are shown installed in hardware pools to compose servers depending on the load situation. DCIaaS QoS-managed connections could be used to attach diskless nodes of the edge site to storage in a central data center, for example via RDMA or non-volatile-memory-over-fabrics (NVMe-oF) protocols, illustrated here for a diskless mobile AI application server (dotted line).
- An on-premises DCI Cluster is shown at the bottom-right of the figure. In this example scenario, this cluster specializes in providing storage capacity that can be connected to resources in public data centers temporarily for processing. For example, on-premises storage could be connected with QoS-managed DCIaaS connections (dotted line) to GPU-based AI analysis nodes in public data centers for data analytics, benefitting from the remote high-efficiency AI inferencing hardware but keeping data storage local. Depending on the needs and requirements of the user, such clusters could also be remotely managed.
- Smart Building DCI Clusters are drawn at the top-right. Here, local aggregation servers receive video streams from cameras via a dedicated local network inside the smart building. The aggregation servers bundle these video streams and transfer these streams via QoS-managed DCIaaS connections (dotted line) to network accelerators in a remote data center. The network accelerators then forward the data to GPUs for massive AI analysis.

The example outlined here touches only a few selected scenarios for DCIaaS systems and shows only a limited number of use cases for illustration purposes. Real systems are expected to be much larger in scale, contain more hardware diversity both for computing and networking, and handle mixed workloads of far more use cases at the same time. The IOWN GF will work out more details about overall DCI systems by further updating the DCI Functional Architecture document and creating further PoC Reference documents.

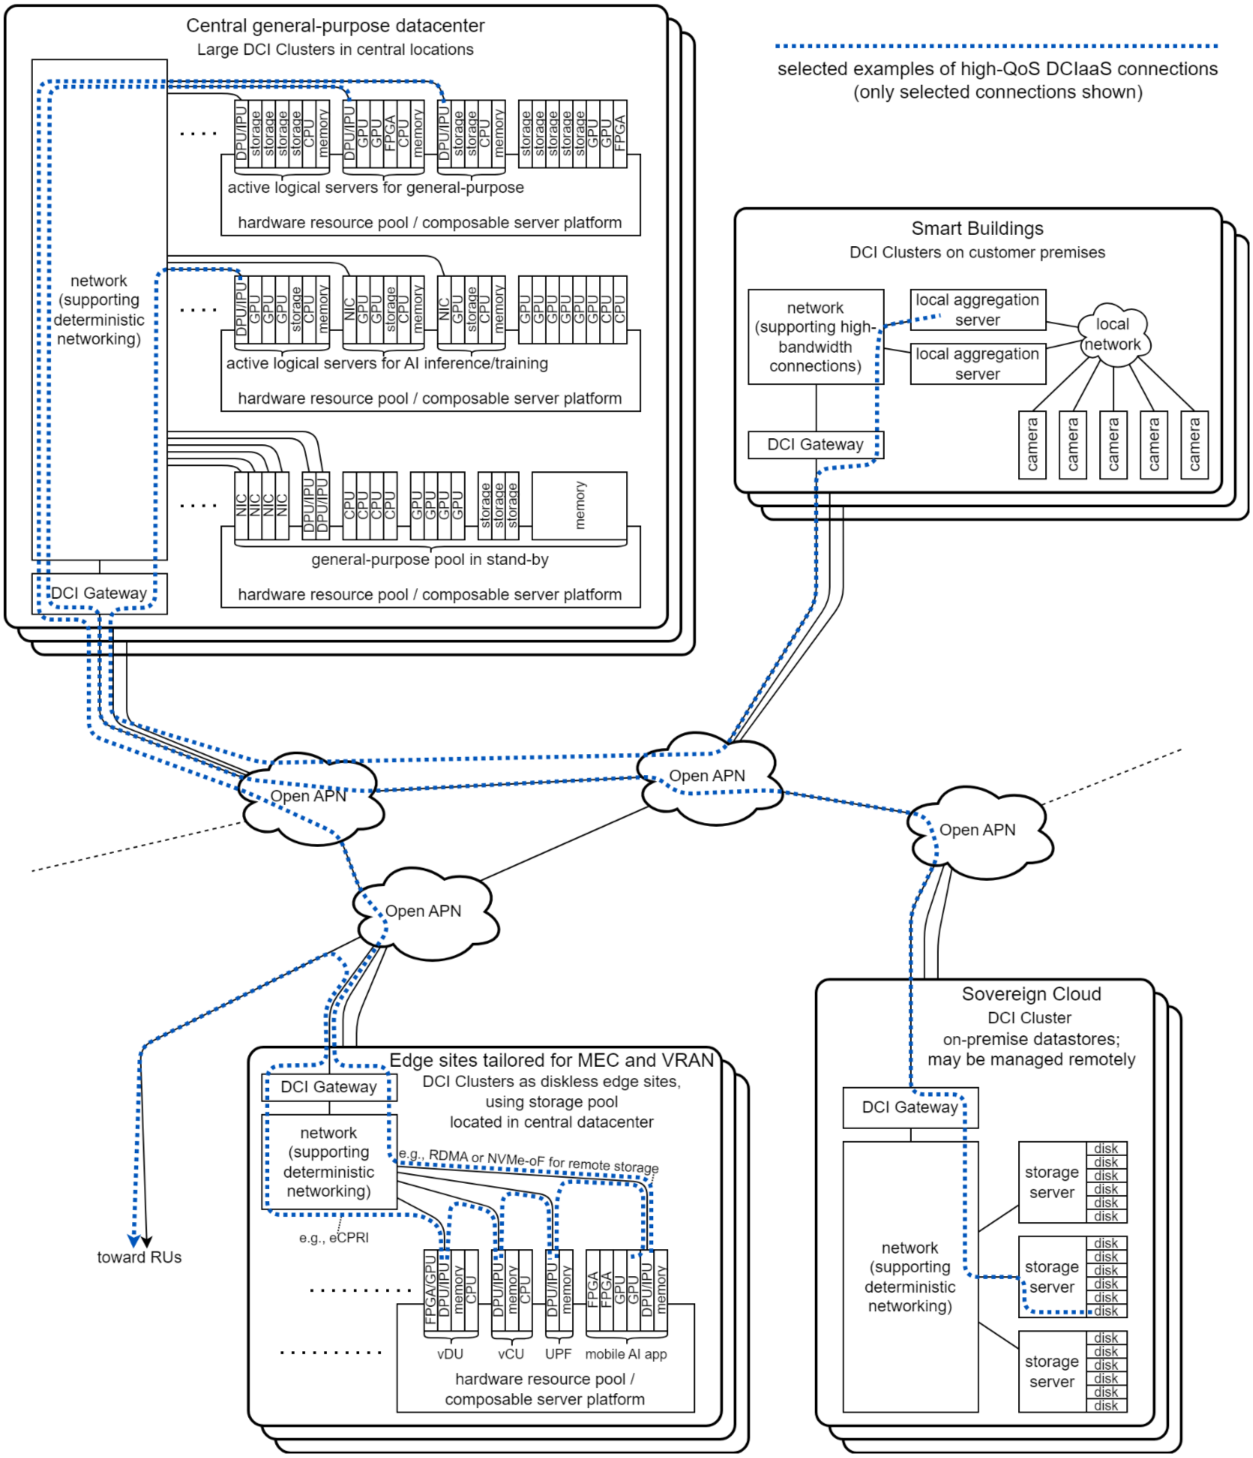

Figure B-1: Example of a DCI system.

Appendix C. Connectivity requirements of IOWN GF use cases

This appendix summarizes the connectivity requirements of IOWN GF use cases. An overview is presented in Table C-1: For major use cases, the purposes and requirements of high-QoS connections between geographically distant sites are outlined. In addition, endpoints of such connections are classified whether their location is constrained by the use-case itself, or if DCI systems have some degree of freedom to select the location of these endpoints. Based on this information, finally the main geographical constraint type for the high-QoS connections of each use case is given.

As a result, Table C-1 illustrates that IOWN GF use cases require high-QoS connections across geographically distant locations and sites. As mentioned in Section 2, for many use cases, at least for one end of such connections, the geographical location is not completely fixed by the use case itself, and the DCI system has some degree of freedom or the duty to choose such locations according to user performance, resource, and regulatory requirements. For example, in the vRAN use case, radio units (RUs, attached to mobile base station antennas) need to be connected to distributed units (DUs, signal processing servers). Here, the geographical location of the radio antennas (RUs) is fixed and cannot be changed by the DCI system. On the other hand, the location of the signal processing servers (DUs) is not inherently fixed by the use case; the user may delegate the decision of where to instantiate the signal processing nodes to the DCI system, with the DCI system being responsible to set up a connection with sufficiently QoS to carry mobile fronthaul data between the RU and DU endpoints. Therefore, in Table C-1, the main geographical constraint type for vRAN “RU ↔ DU” connections is stated as “fixed ↔ variable”.

This kind of freedom poses new requirements and challenges to orchestration technology. However, having such freedom provides a possibility to optimize the efficiency of a DCI system as a whole.

Table C-1: Connectivity requirements of IOWN GF use cases.

Use case	Purpose and requirements of high-QoS connections between locations (over WAN)	High-QoS connection endpoints: Geographically fixed	High-QoS connection endpoints: Geographically variable (likely benefit from joint network and compute orchestration)	Main geographical constraint type
5G Mobile Fronthaul / vRAN ***[IMN](Sec. 5.2.4.1.1, p. 58, App. A, p. 80, 82).***	• Transmission between radio unit (RU) and distributed unit (DU) via (e)CPRI protocol ○ Bandwidth of 25Gbps ~ 75 Gbps per RU ○ Latency ≦0.25ms.	• Antenna (RU) location.	• vDU, condition: located within latency range of the RU.	• fixed ↔ variable
AIC ILM - Live Music [ILM] ***[ILM](Sec. 2.2.1.1, p. 18), [IM](Sec. 2.3.2.2, p. 22).***	• Raw 4K low-latency video stream from rendering GPUs stream to HMD at 120fps. ○ Bandwidth of 69.12 Gbps (in case of 120fps). ○ Latency ≪10ms. • Raw volumetric video stream from Artist's capture studio to Edge DCs at 30fps or 60fps. ○ Bandwidth of 56.35 Gbps (in case of 30 fps).	• Head-mounted display (HMD) at user premises. • Artist's capture studio which is close to central DC. • Edge DCs.	• GPU, condition: located within latency range of the HMD.	• fixed ↔ variable
CPS AM - Area Management for Security ***[AM](Sec. 5.4.2.1, p. 47).***	• Aggregated media streams sent to AI analysis data center. ○ Bandwidth of 60 Gbps per edge location.	• Cameras, LiDARs, and other media sources at POIs.	• AI analysis data center: location mostly constrained by compliance requirements.	• fixed ↔ variable
Sovereign Cloud platform	• Interactions between local database store to remote CPUs or accelerators.	• Database storage devices on user premise(s).	• Data center locations of temporarily borrowed processing elements, conditions: depending on user applications.	• fixed ↔ variable
Distributed data center	• Streams between components, e.g., between compute and storage in different micro-data centers.	• Fully dependent on use case.	• Locations of both compute and storage (and also all other resources) is free within the user latency bounds.	• variable ↔ variable

Table C-1: Connectivity requirements of IOWN GF use cases. (continued)

Use case	Purpose and requirements of high-QoS connections between locations (over WAN)	High-QoS connection endpoints: Geographically fixed	High-QoS connection endpoints: Geographically variable (likely benefit from joint network and compute orchestration)	Main geographical constraint type
CPS IM - Remote Controlled Robotic Inspection [RCRI] ***[RCRI](Sec. 5.4.2, p. 26)***	• Low-latency high-reliability 8K 60fps data streams. ○ Bandwidth of 72Gbps (raw), 2.5Gbps (JPEG XS). ○ Latency ≦100ms.	• Robots at the industrial plant site. • Operator site.	• Not discussed up to now. In the future, footage may be streamed to separate storage for archival with lower performance requirements.	• fixed ↔ fixed (for direct control)
General requirements of databases and IOWN Data Hub	• Extremely high rate transmission for synchronization between reliability domains and incoming transactions from customer and edge locations. ○ Connection latency may be lower-bounded by the required distance for natural disaster resilience and/or upper-bounded by the requirements of the individual use cases.	• Aggregation point for IoT data for ingestion, such as CPS AM ingestion node. • Backend store location when attaching database frontends or analysis nodes to existing storage. • Source and destination locations of storage devices when moving existing databases between sites.	• Database frontends dynamically attached to existing storage. • Compute notes temporarily attached to existing storage. • Database replica (at provisioning time). • Database main copy (at provisioning time).	Depending on use case: • fixed ↔ variable • fixed ↔ fixed • variable ↔ variable
Digital Twin use cases **(Use cases considered include Green Twin, Human Digital Twin, Remote Control Operation, etc. [DTF])**	• Extremely high data volume requirements between endpoints to manage the high amount of data expected for the various use cases; ○ e.g., in the order of TB/s for the building managed in the green twin and disaster notification use cases [DTF]. • Connection latency can be fundamental for real time digital twin update in critical use cases. • Maintaining QoS while allowing cross-domain data flows.	• . Camera, Sensors, LiDAR, network and IoT devices in the infrastructure (e.g., building, vehicle, humans). • Aggregation points for IoT data, when applicable..	• Data pre-processing close to the source. • Generating simulation and ML models at powerful nodes. • Running simulation and ML models for mission critical missions within latency boundary.	Depending on use case: • fixed ↔ variable • fixed ↔ fixed • variable ↔ variable

References

- [AM] https://iowngf.org/wp-content/uploads/formidable/21/IOWN-GF-RD-RIM-for-AM-Use-Case-1.0.pdf
- [DTF] https://iowngf.org/wp-content/uploads/formidable/21/Digital_Twin_Framework_Analysis_Report.pdf
- [IMN] https://iowngf.org/wp-content/uploads/formidable/21/IOWN-GF-RD-Technical-Outlook-for-Mobile-Networks-1.0-1.pdf
- [ILM] https://iowngf.org/wp-content/uploads/formidable/21/IOWN-GF-TS-RIMforInteractiveLiveMusicUseCase-1.0.pdf
- [RCRI] https://iowngf.org/wp-content/uploads/formidable/21/IOWN-GF-RD-RIM_RCRI-1.0.pdf

History

Revision	Release Date	Summary of Changes
1	October 2023	Initial Release
1.1	October 2023	Editorial Change (fixed typo)