Services Infrastructure for Financial Industry Use Case

Classification: APPROVED REFERENCE DOCUMENT

Confidentiality: PUBLIC

Version 2

April 2025

[FSI Use Case]

Legal [draft]

THIS DOCUMENT HAS BEEN DESIGNATED BY THE INNOVATIVE OPTICAL AND WIRELESS NETWORK GLOBAL FORUM, INC. ("IOWN GLOBAL FORUM") AS AN 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/

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

1. Introduction

1.1. Opportunities for the Financial Industry

Perhaps more than any other industry, the financial services sector is deeply engaged in its ongoing digital transformation. This journey has not been without some obstacles, as the performance limitations of current technologies, increased regulation, and reliability requirements have conspired to impact business agility and service resiliency while presenting new opportunities to improve infrastructure and advance new service models.

Many IT solutions vendors offer cloud computing platforms that either serve as a platform for or can be integrated with existing IT systems. However, the very nature of financial institutions demands hybrid cloud solutions capable of connecting cloud infrastructure among multiple data centers with stable bandwidth and low latency. Further complicating the matter, existing hybrid cloud solutions are significantly complex, expensive, and challenging to implement.

The following digital initiatives are just some of the opportunities that technologies such as the Financial Services Infrastructure solution from the Innovative Optical and Wireless Network (IOWN) Global Forum will enable organizations to address.

- **The Ongoing Shift to Digital Banking Services:**
  As of 2024, it's estimated that around 3.6 billion people (Juniper Research, 2020) globally are using digital banking services. In the United States alone, the number of digital banking users is projected to reach nearly 216.8 million by 2025 (Statista, 2023). Additionally, about 73% of customers (Deloitte, 2018) use online banking platforms at least once a month, indicating a significant adoption of digital banking services among consumers. As the world becomes more dependent on digital financial services, financial institutions must increase their computing power to deliver rich user experiences (UX) and improved resilience to serve the fundamental needs of society and the economy. Financial Services Infrastructure with IOWN technology increases the agility of digital services by providing greater customer touch points, performance, and resiliency.
- **Cross-Industry Collaboration/Competition:**
  Today's financial institutions provide diverse services, from microfinance between individuals and small companies to sharing risks with the most prominent businesses and governments. As a result, collaborating with and sharing data across multiple industries is essential for new services. For example, new financial services such as Buy Now Pay Later (BNPL) are made possible through collaboration between business companies, Fintech companies, and financial institutions. These new services will be more direct and provide faster service to new customers. They are accelerated by B2C companies, also known as bank-as-a-service providers, and embedded and automated via technology, including AI. Financial Services Infrastructure with IOWN technology will be essential in linking disparate industries and institutions, fostering new markets and business segments.
- **Leveraging Data to Provide a Personalized Customer Experience:**
  A superior customer experience provides critical competitive differentiation and is quickly becoming a standard expectation for financial services customers. A promising example is to offer a customer an investment product that complements the customer's asset portfolio. This service example requires the collection and analysis of a variety of social and business trend data that is updated daily. It is then expected to detect and collect events such as creation, update, and deletion in each trend data source (these are likely to be small in data size). Subsequently, a large number of transactions will be generated. Financial Services Infrastructure with IOWN Technology provides high-performance interconnections to rapidly collect data from connected data stores and build personalized services that consistently provide immediate responses to the user.
- **Navigating and Automating Regulatory Compliance:**
  Keeping up with the ever-changing landscape of financial regulations, such as the Digital Operational Resilience Act (DORA) in Europe, the California Consumer Privacy Act (CCPA) in California, and various global anti-money laundering standards, is a constant challenge. IT departments must ensure that their systems and procedures

are compliant. Hence, financial institutions must leverage services provided by cloud providers who have a far larger customer base and have invested heavily to track compliance. That said, some services should also be contained on-premises due to regulations. Once realized, Financial Services Infrastructure with IOWN Technology can bridge public cloud and on-premises services, enabling financial institutions to use services provided by cloud providers and provide secure and resilient services to end-users in compliance with various regulations.

- **Maintaining Legacy Systems While Delivering Next-Gen Services:**
  Financial institutions must maintain legacy systems because they often cannot discontinue service to their existing customer base. As a result, they are sometimes forced to maintain some systems for generations. In other cases, some financial services rely on legacy systems (e.g., mainframes) that are difficult to deploy off-premises and constrain other peripheral systems to the same site. Financial Services Infrastructure with IOWN enables the flexible placement of peripheral systems at different sites.

1.2. Purpose of the Services Infrastructure for Financial Industry Use Case Project

The Services Infrastructure for the Financial Industry Use Case project aims to deliver performance and reliability far beyond existing solutions, enabling financial entities to compete successfully in this crowded and highly regulated market. The IOWN Global Forum, an alliance of the world's leading technology innovators, has dedicated its collective efforts to developing new technologies that enable the financial sector to address these challenges with agile and resilient services.

To fulfill its purpose, the project proposes that a financial service interconnect its own data centers and industrial, private, and public clouds to build and operate its own zone (like availability zones), allowing application deployment and workload migration. The Services Infrastructure leverages highly advanced IOWN Global Forum technologies, such as Open All-Photonic Networks (Open APNs), which primarily rely on optical data transmission (using laser light rather than electricity) as the basis for data transmission. This approach offers far higher energy efficiency, performance, resiliency, and more dynamic and flexible network management than existing technologies.

Eventually, the project's value is reduced operational expenditures, which provides a more cost-effective approach with a lower Total Cost of Ownership (TCO) and greater Return on Investment (ROI) compared to legacy data infrastructure solutions.

1.3. Objective and Scope of the Services Infrastructure for Financial Industry Use Case Project

The Services Infrastructure for Financial Industry Use Case project aims to define a reference design for computing infrastructure across multiple data centers, incorporating the advanced capabilities needed by financial institutions for agile and resilient operations.

The scope of this project is to:

1. Describe the Services Infrastructure for the Financial Industry Use Case and its key requirements.
2. Define the Technology Evaluation Criteria, which include reference cases and critical benchmarks.
3. Develop the Reference Implementation Model, which provides a practical implementation of IOWN technologies as a reference model for realizing the use case.
4. Define the Proof of Concept (PoC) Reference, which provides guidelines for conducting PoCs for the use case to evaluate the Reference Implementation Model with the defined Technology Evaluation Criteria.
5. Develop and evaluate the PoC based on the Reference Implementation Model and PoC Reference.

This document covers Steps 1 and 2 of this activity, which aims to engage early adopters in the financial industry.

The primary purpose of this document is to outline how IOWN Global Forum technologies will help financial services organizations develop game-changing services with agility while maintaining very high service resiliency with Financial Services Infrastructure with IOWN-enhanced disaster recovery capabilities.

The following sections of this document describe use cases and define their key features and performance objectives. This project will also define technology evaluation criteria to meet these features and objectives, including reference cases, key benchmarks, and evaluation guidelines.

2. Use Cases

Section 1 describes several key points: operational resiliency, agility, lower TCO, and regulatory compliance. Financial institutions can leverage IOWN technology to address these issues. This technology effectively supports financial use cases with hybrid cloud environments across multiple data centers. Considering the target objectives described in section 1.3, we distilled two significant and promising use cases that are introduced below:

- Intra-regional application deployment and workload migration for operational resiliency and agility.
- Inter-regional back-ups and workload migration for resiliency.

Beyond resilience, future reports will address how to use the technology for these use cases to also improve sustainability for green data centers.

2.1. Intra-regional Application Deployment and Migration to Improve Operational Resiliency and Agility

2.1.1. Description

This use case aims to improve the flexible usage of computing resources, e.g., resource pooling and dynamic resource allocation, for operational resiliency, agility, and cost optimization. If applications can be migrated without service outages and computing resources can be used more flexibly, it will improve scalability and facilitate the provision of new services such as Banking-as-a-Service, which relies on external service requests and makes it challenging to predict transaction volumes. In terms of operational resiliency, low-priority services can be moved to another data center so that high-priority services can respond to topical events. As another example, all services could be deliberately evacuated to another location immediately upon a local disaster such as a fire.

In this use case, a financial services institution builds and operates one logical computing zone by interconnecting private data centers, leased spaces in co-location data centers, operators' offices, and private/public clouds in one region (Note).

Note: We define “a region” as a geographical area within a 50km radius. This comes from the geographical measures in Japan, but different values can be applied to other countries. See Appendix A for the definition of region.

2.1.2. Key Requirements

The computing resources in one logical computing zone should achieve the following key requirements (KRs) and features:

- Virtual Machines (VM) Migration across data centers
  - ▫ KR1.1 --- Downtime for VM online migration: less than 1 second (Note 1)
    - ▪ One data center to another data center or public cloud
      - ✧ When migrating an application from one data center to another or to a public cloud, the downtime needed during VM migration is short enough not to impact the application's functionality (Note 2).
      - ✧ In online migration, a sender server copies the virtual data of the original VMs (including VM images and memory) and sends it to another data center. Then, the differential data of the original VM during the data transfer is sent to the destination while suspending the original VM. After the VM migration is completed, the migrated VM is resumed. Also, access to the VM will follow the migrated VM. The downtime is the time needed to ensure data and configuration consistency between when the original VM is suspended and when it restarts at another data center.

- ✧ Downtime includes the processing time to ensure data and configuration consistency so that the application can continue successfully.

Note 1: However, financial institutions expect downtime for VM migration to be closer to 100ms, because the shorter the downtime, the more services can be operated flexibly.

Note 2: VM image conversion is required when migrating to a public cloud, resulting in longer VM downtime. Even in such cases, downtime should be minimized. In addition, data that cannot be replicated to the public cloud will access storage in the data center from which it is migrated. The performance degradation of data access must be acceptable to the application.

2.2. Inter-regional Back-Ups and Migration to Improve Resiliency

2.2.1. Description

In this use case, a financial services institution achieves disaster recovery and inter-regional workload migration, e.g., for energy demand and supply balancing, by using infrastructures in multiple regions with replication between sites close to real-time. IOWN technology can simplify systems and enhance their effectiveness. For example, it enables the placement of the secondary system directly at the backup data center, allowing us to bypass the redundant intermediate system typically found in the main data center. This approach not only streamlines the system architecture but also significantly improves data security and system efficiency.

The distance between two sites should be geographically meaningful for each geographical location and should be defined as two kinds of distance: minimum and maximum. The minimum distance is the distance that must be maintained for disaster countermeasures in the inter-region use case. On the other hand, the maximum distance is almost the same as the maximum distance defined for each geographical location and does not necessarily have to be achieved in this use case.

Note: See Appendix A for the size of regions.

We prioritized a disaster recovery scenario and defined RPO/RTO (Recovery Point Objective/Recovery Time Objective) from the perspective of business continuity.

There are priorities in system groups because business continuity is one of the key demands of today’s financial organizations. A financial institution has the highest priority for business-critical data and services, such as accounting systems and its ledger data (Tier 1 Systems, Top Priority Group, < 1% of all systems). It also places major systems as a moderate priority to be resumed in reasonable duration (Tier 2 Systems, Second Priority Group, < 10% of all systems). Data synchronization and application migration are mandatory to realize this system's autonomous recovery. The remaining systems other than Tier 1 and Tier 2, the general group, are referred to as Tier 3.

In traditional systems, it is challenging to replicate data synchronously to a remote data center, either causing application performance degradation or asynchronous replication that either tolerates a certain level of data inconsistency or assumes operational coping. Essentially, Tier 1 systems require synchronous replication to ensure continuity of financial transactions.

2.2.2. Key Requirements

For financial service operation, all systems, Tiers 1-3, should achieve zero RPO totally (Note 1). However, this requirement cannot be readily translated into the database’s replication performance because of the following reasons:

Due to the nature of database systems, there is a trade-off between data protection capability and update throughput. Synchronous replication can achieve data protection at the DB layer. However, it degrades update throughput. In particular, performance degradation worsens as the latency between the replicated servers increases.

On the other hand, asynchronous replication can suppress performance degradation. However, in case of disaster, there would be updates that had been accepted but not yet been replicated to other sites.

Given the above, we should find a solution that can achieve good update throughput while tolerating a certain amount of recovery time, as summarized below:

- KR2.1 --- RPO
  - Tier 1,2,3 Systems:
    - Zero (Note 1).
    - User data on system use for financial services cannot be lost.
- KR2.2 --- RTO (Note 2)
  - Tier 1 Systems:
    - Less than half an hour.
    - Additional time is acceptable if a human operator needs to restore unprocessed transaction data or check the consistency of data.
  - Tier 2 Systems:
    - Less than two hours.
    - The backup period is not always real-time, and hourly or daily is ordinally timing for full backup.
  - Tier 3 Systems:
    - Less than 24 hours.
    - The backup period is not always real-time, and daily or special maintenance time (e.g., holidays) is also ordinarily a time for full backup.

Note 1: RPO is not intended to be achieved solely by the DBMS but also includes the application layer and the manual operations of the financial company.

Note 2: We have confirmed that this RTO is sufficient by interviewing people with experience in financial institutions.

3. Technology Evaluation Criteria

First, we define reference cases to accurately evaluate these reference designs under the specific conditions that meet the key requirements defined in Section 2. Then, we select key benchmarks for each use case based on the technology evaluation definition in Figure 3-1 to evaluate reference designs under the conditions of the defined reference cases. Because the development of instrumentation methods for green data center utilization is in its very early stages, technology evaluation criteria of green data centers are slated for future study.

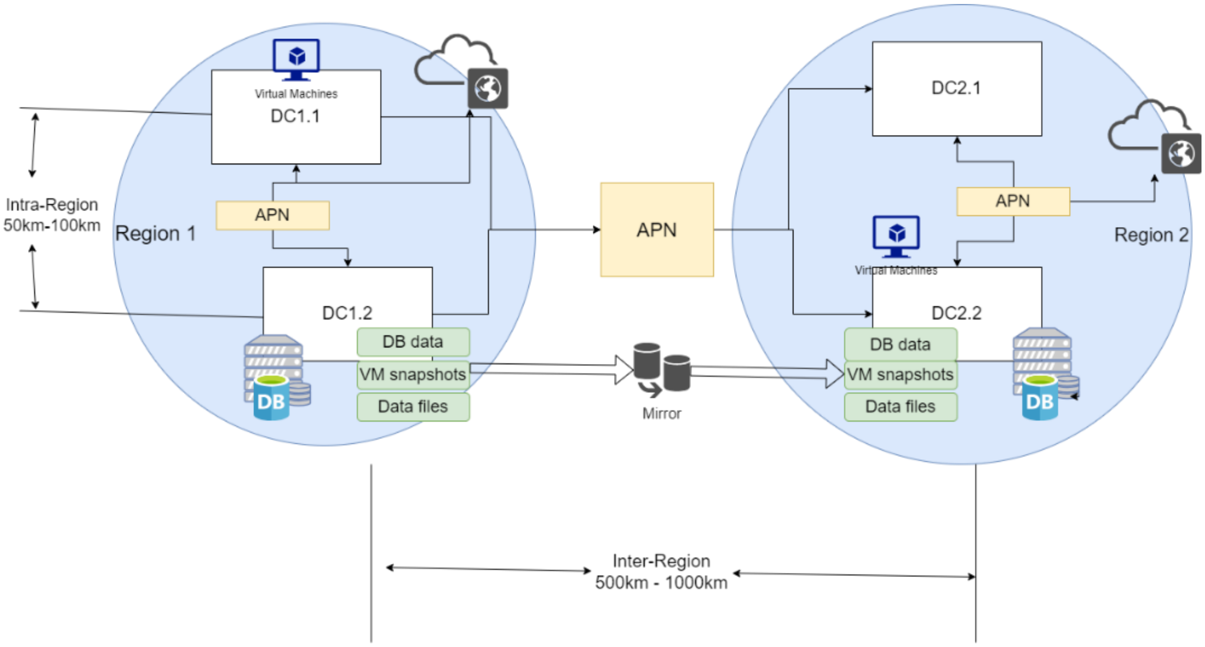

Figure 3-1: Technology Evaluation Definitions

3.1. Intra-regional Application Deployment and Migration to Improve Operational Resiliency and Agility

This section describes the technical evaluation criteria for the use case 2.1.

3.1.1. Reference Case

- Today, each data center site has 100-1000 physical commercial-off-the-shelf (COTS) servers running many financial applications on guest OSs in 1000 - 10000 virtual machines.
- The normal backend traffic volume between two data center cluster sites is less than 100 Gbps, and it will increase over 100 Gbps (up to 200 Gbps) during data replication, etc.
- Database journal data is to be used for live replication, and data files are replicated among data centers placed intra-regionally within an area whose distance is between 50km and 100km (0.25 msec to 0.5 msec in light speed).
- There is a constant delay between the network interfaces of servers across data centers and Open APN networks. The bandwidth ranges from 10-40 Gbps, and one-way latency is 0.5 msec, with a general optical fiber propagation delay of 100 km.
- The bandwidth of the VM host server I/F can be up to 25 Gbps.
- The VMs to be migrated are assumed to have a data size of 100 GB and a memory size of 16 GB and are to be hot-migrated or cold-migrated.

3.1.2. Key Benchmarks

3.1.2.1. Scenario 1: DC1.1 to DC 1.2 VM Migration

Description

- For VMs at the origin server running in DC 1.1, hot migration to DC 1.2 is requested. First, the memory contents of the VM running in DC 1.1 are copied to a VM host located in DC 1.2 in the background.
- The VM running in DC 1.1 is switched to a suspended state. Then, the parts of the VM memory that were dirtied by the VM after the previous copy to DC 1.2 data are copied again, resulting in identical memory images in DC 1.1 and DC 1.2. After that, the migrated VM is restarted on DC 1.2 based on the VM data that was copied from the origin data center DC 1.1.
- The network configuration is updated so the new VM in DC 1.2 replaces the original VM in DC 1.1.

Metrics

- Preparation time to migrate VM data from the VM migration start instruction before VM migration downtime occurs.
- Downtime in VM connection switchover in VM migration from DC 1.1 to DC 1.2.

3.1.2.2. Scenario 2: VM Import from DC1.2 to Public Cloud

Description

- Assuming a system test when adding a new financial service, VM data from the standby system site (DC 1.2) will be replicated to the public cloud and started in the public cloud with VMs.
- Data not replicated to the public cloud will be accessed in storage at DC 1.2.

Metrics

- The time from preparation start to when the application is ready to launch.
- Performance degradation in latency and throughput of test app benchmark scores between DC 1.2 and public cloud.

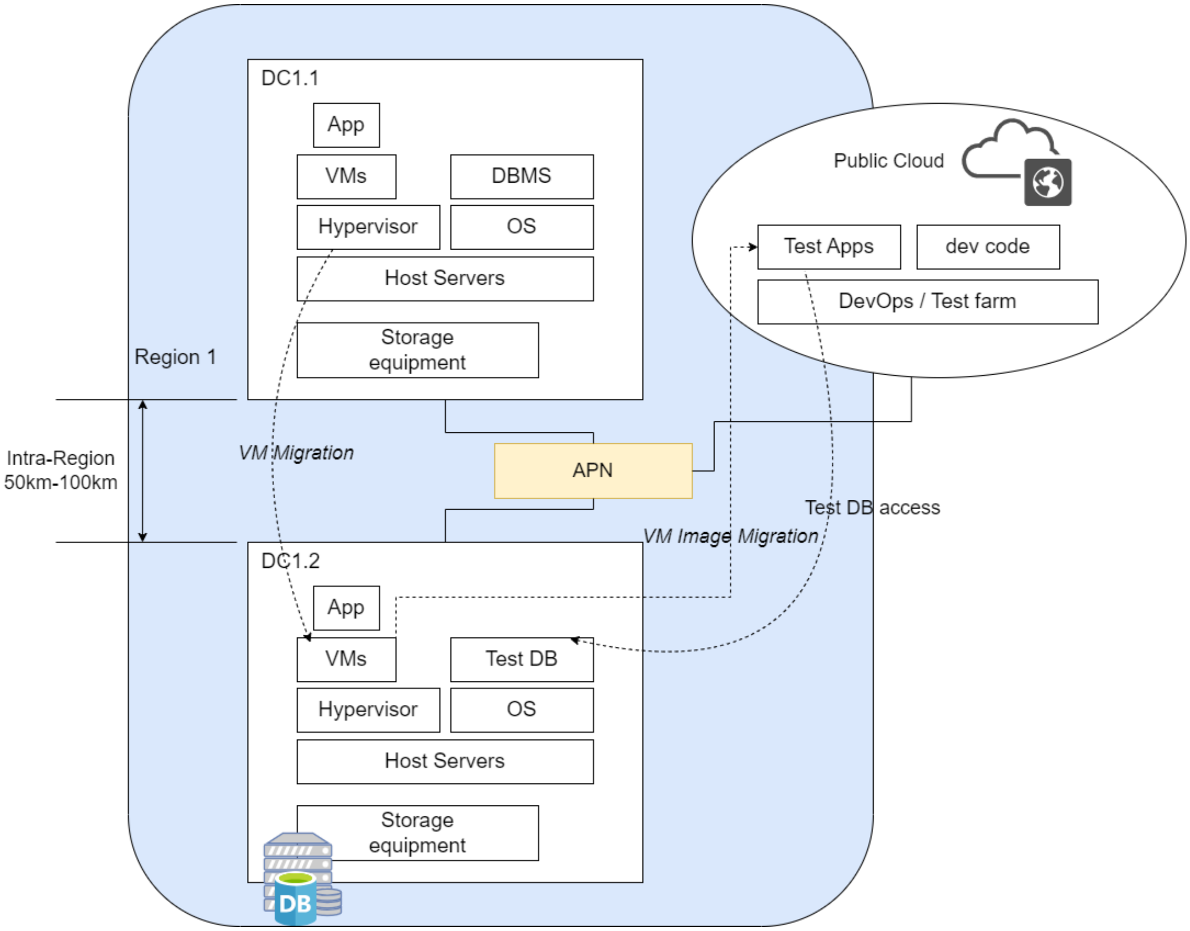

Figure 3-2: Data Flows in the Intra-Region Technology Evaluation Scenario

3.2. Inter-regional Back-Ups and Migration to Improve Resiliency

This section describes the technical evaluation criteria for the use case 2.2.

3.2.1. Reference Case

- Both minimum and maximum distances are defined for each geographical location, e.g., 250 - 2,500 km in the case of Japan.
- Each data center site has 100-1,000 physical COTS servers running many financial applications on guest OSs placed in 1,000 - 10,000 virtual machines.
- The network lines between the two data center cluster sites are redundant.
- Normal backend traffic volume between two data center cluster sites is less than 10-100 Gbps, depending on the Tier 1, 2, and 3 ratios. It will increase to over 100 Gbps (up to 200 Gbps at peak) during data backup, etc.
- File transfer of Database Management System (DBMS) data (log journals, commit log files, etc.), VM disk images, and other file data (log files) is also required.

3.2.2. Key Benchmarks

3.2.2.1. Scenario: Inter-Region VM Migration

Description

- As described in KR2.1, the DB of Tier 1 systems is replicated from DC1 to DC2, the data of Tier 2 systems is backed up hourly, and the data of Tier 3 systems is backed up daily.
- When a disaster occurs, and the shutdown of services running in DC1 is expected to be imminent, data from VMs running in DC1 are moved to DC2, which is more than 250 km away in the case of Japan(Figure 3-3).
- System migration is performed in the order of Tier 1, Tier 2, and Tier 3.

Metrics

- Downtime of VM migration from DC1 to DC2 (for all Tiers).
- DB Performance, in Transaction / Sec (for Tier 1).
  - Performance degradation should be minimized in TPS and round-trip time when migrating Tier 1 systems from DC1 to DC2.
- Backup speed (for all Tiers).
  - To accurately measure data backup speed, it is essential to consider the processing time needed to save (flush) data from the application layer, where the application software operates, to permanent storage. Most real-world Database Management Systems (DBMS) use advanced mechanisms to quickly read, transfer, and write data, such as parallel data processing (handling multiple data streams at once, as explained in Appendix B). Therefore, the backup time is determined by the arrival time of the most delayed data. This means that multiple factors beyond the raw available bandwidth need to be considered when designing for specific performance goals.
  - Backup speed is not equivalent to throughput because backup will be done with snapshot milestones, which are the latest, consistent locations of all recent data. It means the most delayed data arrival is essential. (see Appendix B).
- RPO
  - Evaluate system RPO. It is better to be smaller. RPO zero will be realized with the IT system and operational procedures. If some portion of data is guaranteed to be RPO zero, it simplifies an operational recovery procedure.

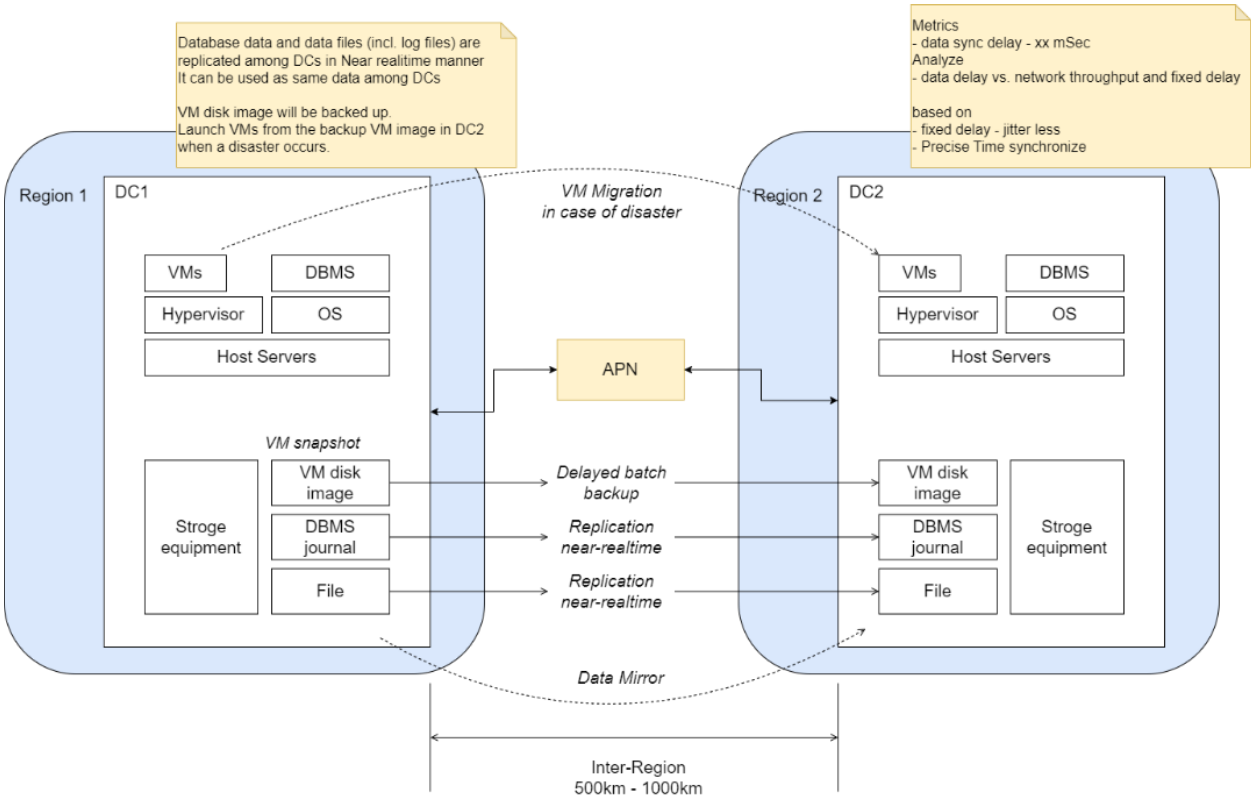
Database data and data files (incl. log files) are replicated among DCs in Near realtime manner
It can be used as same data among DCs
VM disk image will be backed up.
Launch VMs from the backup VM image in DC2 when a disaster occurs.
Metrics
- data sync delay - xx mSec
Analyze
- data delay vs. network throughput and fixed delay
based on
- fixed delay - jitter less
- Precise Time synchronize
Region 1
DC1
VMs
DBMS
Hypervisor
OS
Host Servers
VM snapshot
Stroge equipment
VM disk image
DBMS journal
File
VM Migration in case of disaster
APN
Delayed batch backup
Replication near-realtime
Replication near-realtime
Data Mirror
Region 2
DC2
VMs
DBMS
Hypervisor
OS
Host Servers
VM disk image
DBMS journal
File
Stroge equipment
Inter-Region 500km - 1000km

Figure 3-3: Data Flows in the Inter-Region Technology Evaluation Scenario

4. Conclusion

The Services Infrastructure for Financial Industry Use Case will improve the quality of continuous financial services under any conditions, including disaster situations such as earthquakes, floodings, power shortages, storms, natural disasters in general, misconfigurations or human error, “other infrastructure disruptions”, etc. This level of reliability is critical to financial institutions for supporting Business Continuity Planning.

As explained in Section 1, this document covers the first step of the Forum’s activity to engage early adopters in financial institutions. Further documentation, which will contain more technical details than this document, will be issued to evaluate the feasibility of this vision through the development of a Reference Implementation Model and Proof-of-Concept demonstrations.

References

[JUNIPER2020]: Juniper Research, March 2020.

[STATISTA2023]: Statista, “U.S.: digital banking users 2025”, May 2023.

[DELOITTE2018]: Deloitte insight, December 2018.

[DORA]: Digital Operational Resilience Act (DORA).

[CCPA]: California Consumer Privacy Act (CCPA).

[JMA]: Japan Meteorological Agency, “The 2011 off Pacific Coast of Tohoku Earthquake Distribution of JMA Seismic Intensity”.

Appendix A. Geographically Meaningful Size of Regions and Distance Between Regions in the Case of Japan

We define “a region” as a geographical area within a 50km radius, i.e., 100km in maximum distance. We classify the applicable cases into two categories: a regional system and a national-wide system, as shown in Figure A-1.

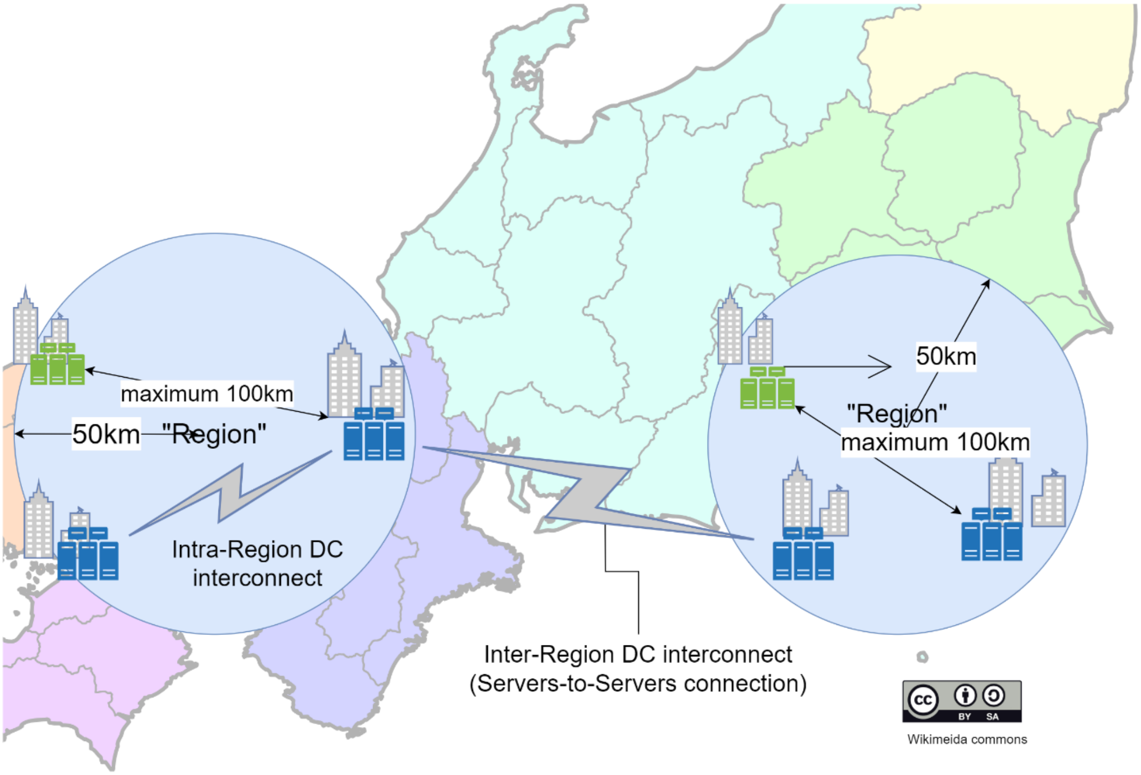

Figure A-1: Geographical Cases for Financial Service Infrastructure with IOWN

The distance between two sites in the nationwide system should be enough for them to be located in different seismic areas, as Japan has had significant earthquake activity over the last decade. This illustrates that a 500km distance is an appropriate separate range for potentially disaster-affected areas (Figure A-2) [JMA].

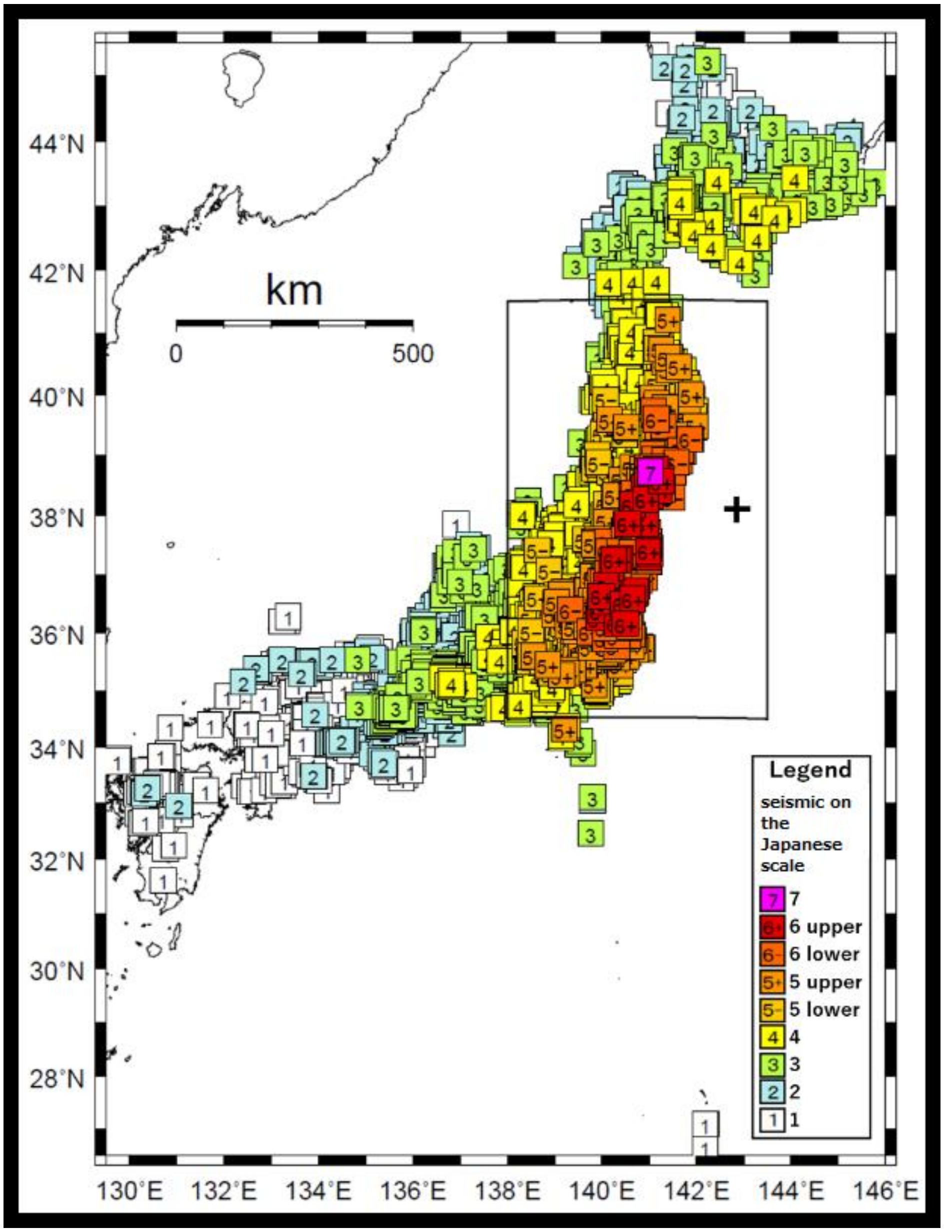

Figure A-2: Affected Area by Mega Earthquake, 2011

(Translated from a publication licensed under CC-BY from Japan Meteorological Agency)

Appendix B. Note for Measurement of Data Transfer Delay

This appendix demonstrates how an All-Photonic Network (APN) improves use cases through detailed behavior analysis. An APN's very high bandwidth allows the sender and receiver to transfer data with great speed, but we should consider the total data flow and validation period.

The applications and software layers required for APN use cases, such as a DBMS and hypervisor, write data in their timing and then “sync” in the file system on demand from a data transfer application or “commit” in the database system. The backup process can be triggered to transfer and start reading written data. The data will be in a high-speed medium, like main memory, as a cache and immediately begin transferring. The receiver can receive data chunks and write this data to a storage medium concurrently.

After all the data arrives, the receiver can process the “validation” of the received data and flush it into a persistent medium, such as SSD storage. This provides time to finish the backup process. “Sync” or “commit” is important because the APN application requires “consistency” among data, which is a “checkpoint.”

APNs can reduce the delay between the application checkpoint and the time when the remote site flushes/syncs to it. The backup completion time will be predictable and consistent, allowing the start and end times to be easily integrated into the schedule without needing significant buffer time. A daily checkpoint backup system only guarantees data from the day before use, which is yesterday’s backup data. Hourly checkpoints provide a two-hour lag. Jitter-less fixed delay provides the ultimate small duration for a checkpoint; the only limitation is on the software layer.

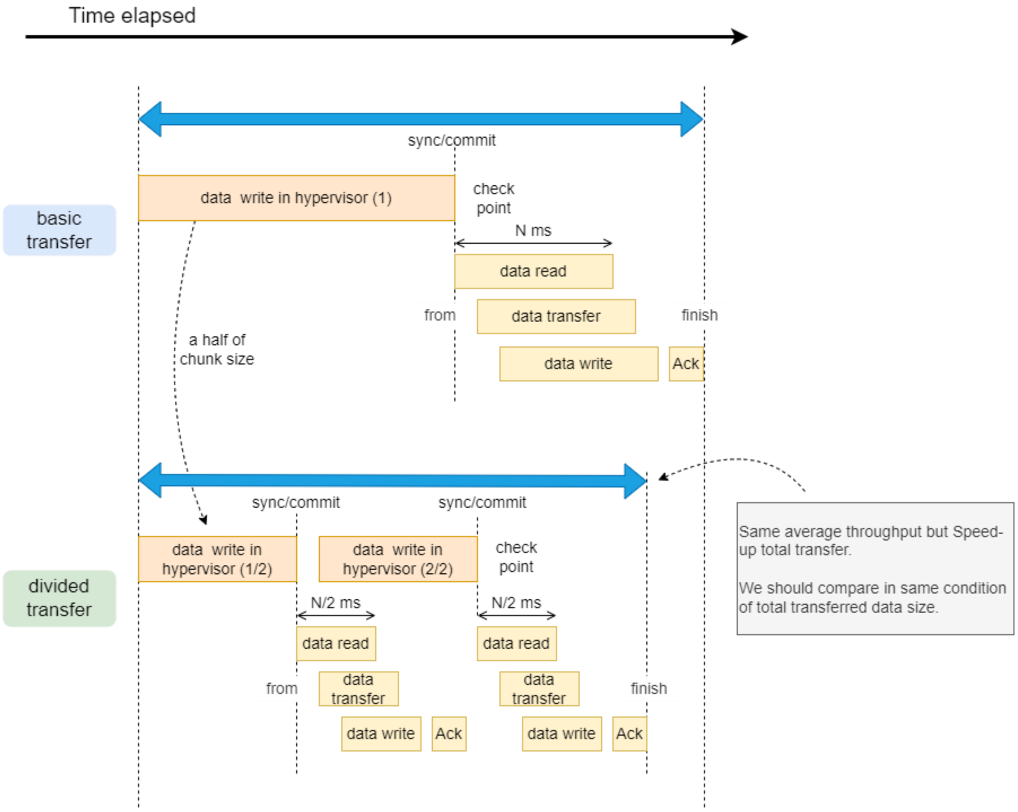

Same average throughput but Speed-up total transfer.

We should compare in same condition of total transferred data size.

Figure B-1: Data Replication of Regional System Benchmark Conditions

History

Revision	Release Date	Summary of Changes
1	July 2024	Initial Version
2	April 2025	• Change the distance of the Inter-Region scenario to be varied. • Fixed some typos.