Dataffinic Computing: Data-Centric Architecture to Support Digital Trust

● Masahisa Tamura ● Eiji Yoshida ● Kohji Yamada

Recently, there is a movement gaining momentum to create business reforms and innovations by storing and utilizing large volumes of data generated on various sites. In addition to the structured data handled in conventional databases, advancements are being made in the use of video and other unstructured. There are three requirements for efficient utilization of unstructured data: data processing performance, cost-performance ratio, and data management. Fujitsu Laboratories is working on the R&D of "Dataffinic Computing," which uses distributed storage systems to achieve high-speed processing of large volumes of data, as a data-centric architecture that supports the utilization of massive volumes of data. This paper outlines data neighborhood processing technology, large-volume memory technology, and high-speed thin client technology, which are elemental technologies of Dataffinic Computing. It also presents an approach to video monitoring systems as an application example.

1. Introduction

Recently, there is a movement gaining momentum to create business reforms and innovations by utilizing massive data generated on various sites. Along with this movement, in addition to the structured data such as customer data and point of sale (POS) data conventionally handled in databases, advancements are being made in the use of rapidly increasing video, log, and other unstructured data. In utilizing unstructured data, the use of information obtained through data processing such as AI analysis allows for efficient analysis.

There are three requirements to utilize massive, ever-increasing volumes of data efficiently.

The first requirement is data processing performance. In conventional data center architecture, the data storage and processing units are separated, connected with a network. Therefore, as the amount of data to be processed increases, data movement to the processing unit sees a bottleneck occur in the processing performance. Scalable high-speed processing according to the increase in data is necessary to resolve this issue. The second requirement is the cost-performance ratio. Storage costs depend on storage performance and amount of data. If the large amount of data processing is required in the high-performance storage, the storage costs become enormous. Therefore, the use of storage devices with different levels of performance and different costs according to the application is key. The third requirement is data management. When utilizing data, the data itself becomes an important differentiating factor from other companies. Therefore, it is necessary to carry out data processing efficiently without allowing data leaks occurring.

In order to satisfy these requirements, Fujitsu Laboratories has been working on R&D for Dataffinic Computing using distributed storage as a data-centric architecture to support the utilization of massive data.

This paper outlines data neighborhood processing technology, large-volume memory technology, and high-speed thin client technology to realize a data processing platform for safe, secure, and efficient utilization of massive data. It also presents an approach to video monitoring systems as an application example of Dataffinic Computing.

2. Dataffinic Computing

The term Dataffinic Computing, which has been being researched and developed by Fujitsu Laboratories, combines "data" and "affinity." To meet the requirements for the utilization of massive data described in the previous section, Dataffinic Computing uses data neighborhood processing technology for high-speed and scalable processing of stored data, large-volume memory technology to achieve a high cost-performance ratio, and high-speed thin client technology for efficient remote data processing. The result is a data processing platform that allows for safe, secure, and efficient utilization of massive data (**Figure 1**). This section describes these technologies.

2.1 Data neighborhood processing technology

This subsection presents data neighborhood processing technology for processing massive data at high speeds over distributed storage.

For unstructured data to be utilized, data processing such as AI analysis is required. As the volume of data increases, however, a bottleneck in the data movement to the processing unit occurs. On the other hand, while data can be processed at high speeds by processing the data in storage without moving it, the degradation of the storage performance due to the impact of data processing poses a problem. The difficulty in achieving these functions at the same time is a challenge when utilizing unstructured data.

In order to solve this issue, Fujitsu Laboratories has been studying data neighborhood processing, which achieves both high-speed data processing and stable storage operations by processing data while controlling resources to be allocated to distributed storage.[1] An execution platform that allows for data processing is built on distributed storage nodes to achieve data processing on the nodes where the data exist.

To execute data neighborhood processing, the storage node that stores the data to be processed is first calculated from the data management information of the distributed storage. Next, a data processing application is deployed on the node to execute processing, and the result is stored in a disk (**Figure 2**). In general, the volume of data of the data processing application and results is small when compared with that of the data to be processed, and data processing can be achieved while the costs of data movement are kept to a minimum.

In addition, data processing takes place within the storage node. Therefore, no inter-node communication occurs even if a storage node is added to increase capacity. This allows the processing performance to be scaled according to the increase in the data volume. Furthermore, the storage node constantly monitors the amount of resources such as the CPU and memory used by the original storage processing and individual data neighborhood processing. If the resources allocated to the storage processing are insufficient, some of the

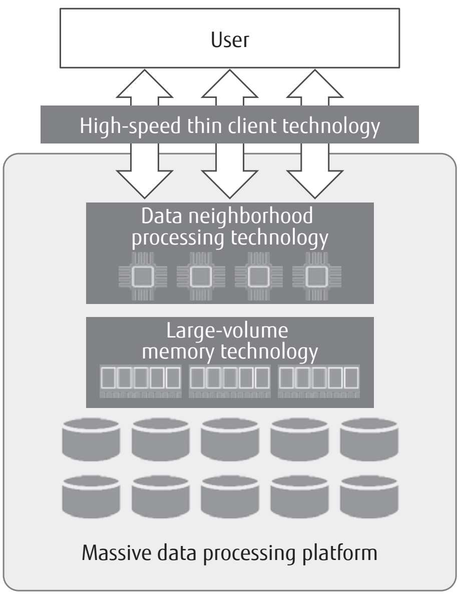

Figure 1
Configuration of Dataffinic Computing.

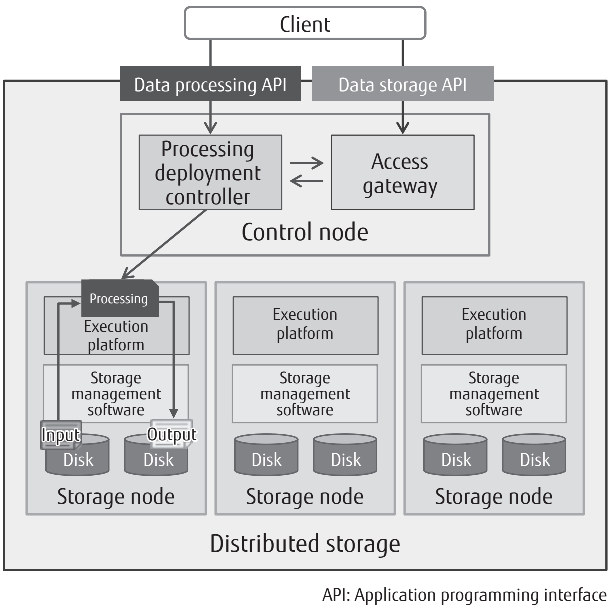

Figure 2
Configuration of data neighborhood processing system.

data neighborhood processing can be stopped to free up resources, thereby realizing stable operation of the original storage processing.

2.2 Large-volume memory technology

This subsection describes a large-volume memory technology that realizes high performance storage.

As the volume of data to be analyzed increases, higher speeds are required not only for data processing in the CPU and graphics processing unit (GPU) but also for data input/output. In-memory processing is effective to deal with this issue. In-memory processing is a method of executing processing by storing all data in the main memory on the server, which reduces the disk input/output overhead, and realizes an incomparably higher-speed processing. However, dynamic random-access memory (DRAM) conventionally used for main memory is costly, which presents a problem when preparing large-volume memory required for processing large volumes of data.

In order to solve this problem, Fujitsu Laboratories has been studying a method of using non-volatile memory as the main memory to realize in-memory processing with a high cost-performance ratio.[2] Non-volatile memory is characterized by low unit volume costs and large volumes as compared with DRAM. However, its speed is lower, meaning a technology to reducing speed degradation is important for its utilization. Accordingly, we are studying a technology for actively arranging in DRAM the data with high access frequency according to the processing content as well as a data access technology that enhances performance taking into account the characteristics of the non-volatile memory.

Use of these technologies allows the memory volume to be increased tenfold without speed degradation. This enables processing that handles larger volumes of data to be made faster. In addition, application of the large-volume memory technology to nodes of distributed storage allows them to be used as large-volume memory for application processing at the time of data processing and as high-speed disks for storage access. This allows for the efficient integration of storage and processing.

2.3 High-speed thin client technology

This subsection describes high-speed thin client technology, which enables data in a cloud to be operated remotely.

In today's data-centric society, the development of AI technology and the improvement of computing performance have given momentum to the movement to acquire new knowledge and create new business and innovation by utilizing large volumes of stored data. For this reason, the stored business data itself becomes an important differentiating factor from other companies. Therefore, to prevent leakages of business data and ensure their integrity against disasters, many companies have introduced thin client environments. These aggregate business data on a server in the cloud to enable processing simply by a client accessing it.

In a thin client environment, only client operation information at hand is transferred to the server. The data are processed on the server in the cloud and the desktop image of the server is transferred to the client. In video playback and CAD design operations, however, the update area of the desktop image is large and the update frequency is high as well. This posed problems such as increased CPU load causing a lower transfer frame rate of the desktop image and difficulty in ensuring sufficient operability due to increased traffic.

In order to solve these problems, Fujitsu Laboratories has developed the high-speed display technology Remote Virtual Environment Computing (RVEC) intended for virtual desktop environments.[3], [4] RVEC has the following features.

- The CPU load and the traffic are reduced as well by analyzing the update frequency for each area of the desktop screen and compressing it with the optimum coding system for the individual area.
- Good operability is maintained even in environments with frequent packet losses and network delays by switching between the Transmission Control Protocol (TCP) and User Datagram Protocol (UDP)-based transmission modes.
- The operation response is improved by estimating the network bandwidth available between the server and client and controlling the transmission frame rate according to the bandwidth.

These features allow the data stored in the cloud to be processed and visualized without the need to transfer them to the client. As a result, the time required for transferring massive data can be decreased and information leakages due to data leaks can be reduced.

3. Application example of Dataffinic Computing

Types of massive data for which the application of Dataffinic Computing effectively functions include image data, log data, sensor data, and genome data. As an example, this section presents the application of Dataffinic Computing to a video monitoring system that handles large volumes of video data (**Figure 3**).

Road administrators of national and local governments have installed a large number of monitoring cameras on roads to help prevent accidents and ensure safety. These are intended for use in video checks and real-time detection of accidents and other incidents, as well as aggregation and storing of the recorded monitoring videos in the data center (cloud). One method of utilization of the large volumes of video data stored is video data analysis, which is used to search for missing persons such as children and the elderly in past videos and checking detailed situations involved in traffic accidents and criminal investigations.

Video data analysis requires large volumes of video data stored in a storage system to be transferred to a server for processing. It requires higher speed transfers of video data than for storing videos, which causes the amount of traffic between the storage system and the server to increase. Accordingly, transfers of video data cause a bottleneck in video data analysis, possibly lowering the speed of analysis processing. In addition, transfers of video data in the cloud to the client for checking and analytics not only reduce user operability due to the transfer of video data but also raise the possibility of video data leakages, resulting in privacy and information leakage problems.

Furthermore, these analysis, checking, and analytics processes for video data stored in the cloud may be performed simultaneously from multiple locations under the road administrator. Since these video data processes require large volume memory, simultaneous execution of multiple data processes may cause the memory capacity of the server to run short, resulting in lower processing response. In the worst case scenario, the system may go down.

The following describes the effect of applying Dataffinic Computing to a video monitoring system that has these issues.

Application of the data neighborhood processing technology to video data analysis allows video data analysis to be processed in the storage system without the need for transferring the video data in the storage system to the server. Therefore, the transfer of the video data does not cause a bottleneck, improving the processing speed of the video data analysis. In order to quantitatively verify this effect, we measured the time required for the process of detecting people and cars from 50 GB video data. As a result, it was confirmed that a data analysis speed 10 times faster than that of conventional methods is possible. Data neighborhood processing allows more video data than before to be

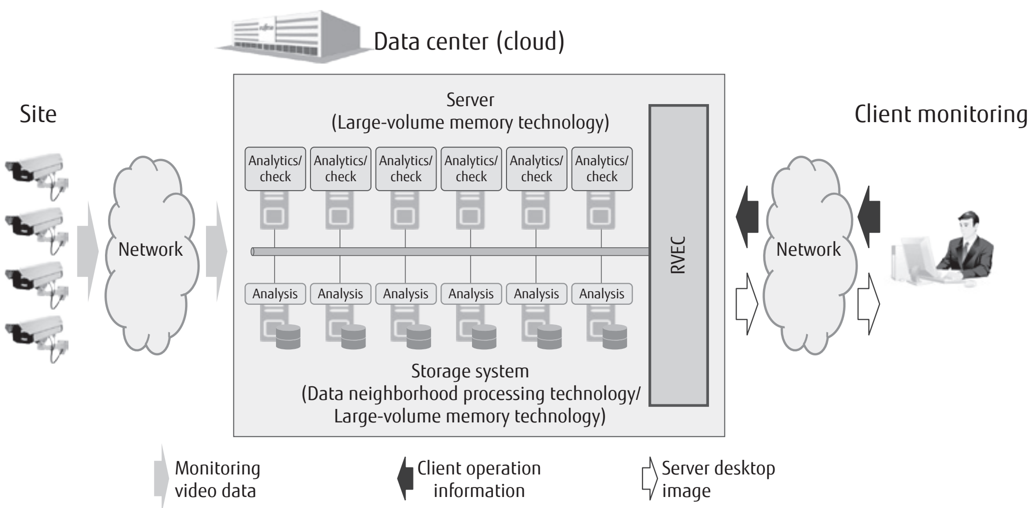

Figure 3
Video monitoring system.

analyzed at high speeds, which makes it possible to quickly identify traffic accidents or detect missing people.

In addition, application of RVEC to video checking and analytics allows video playback and analytics software in the cloud to be operated by the client, enabling the client to check and analyze the video. As a result, even when monitoring work is outsourced, videos can be checked and analyzed without transferring the video data stored in the cloud to the client, preventing leakages of video data. When video data analysis, checking, and analytics processes are simultaneously carried out, application of the large-volume memory technology to video monitoring systems allows memory with 10 times larger capacity than conventional memory to be mounted. This ensures sufficient operability and stable operations without server memory shortages.

As described above, by applying Dataffinic Computing, i.e., the data neighborhood processing technology, RVEC and large-volume memory technology to video monitoring systems, the processing performance, security, and stability of the system can be improved. This in turn leads to the realization of a trustworthy business environment.

4. Conclusion

This paper presented Dataffinic Computing using distributed storage, on which Fujitsu Laboratories is working in relation to the data-centric architecture to support the utilization of massive data.

In the future, we intend to promote the development of technology for more efficient storage and management of massive data and conduct field trials in cooperation with customers. Doing so, we aim to promptly provide technology that can truly support customers' businesses.

References

1) Fujitsu Laboratories: Fujitsu Develops Platform Technology to Support High Speed Processing of Massive Data in Distributed Storage.
https://www.fujitsu.com/global/about/resources/news/press-releases/2018/0920-01.html

2) Fujitsu Laboratories: Fujitsu Conducts Memory Expansion Technology Field Trial, Achieves System Performance Equivalent to 10 Servers.
https://www.fujitsu.com/global/about/resources/news/press-releases/2018/0919-02.html

3) M. Matsuda et al.: "High-Speed Thin Client Technology for Mobile Environment: Mobile RVEC." FUJITSU Sci. Tech. J. Vol. 49, No. 2, p. 184–189 (2013).
https://www.fujitsu.com/global/documents/about/resources/publications/fstj/archives/vol49-2/paper07.pdf

4) K. Yamada et al.: Screen Transfer System Realizing High Image Quality and Low Bandwidth via GPU Built-in Hard Codec. 23rd Annual Conference of the Virtual Reality Society of Japan, September 2018. (in Japanese)
http://conference.vrsj.org/ac2018/program2018/pdf/11C-2.pdf

Masahisa Tamura
Fujitsu Laboratories Ltd.
Mr. Tamura is currently engaged in the research and development of data system technology.

Eiji Yoshida
Fujitsu Laboratories Ltd.
Dr. Yoshida is currently engaged in the research and development of data system technology.

Kohji Yamada
Fujitsu Laboratories Ltd.
Mr. Yamada is currently engaged in the research and development of video coding technology and video coding applied technology.