Data processing architecture can reconfigure content within IoT data processing stream
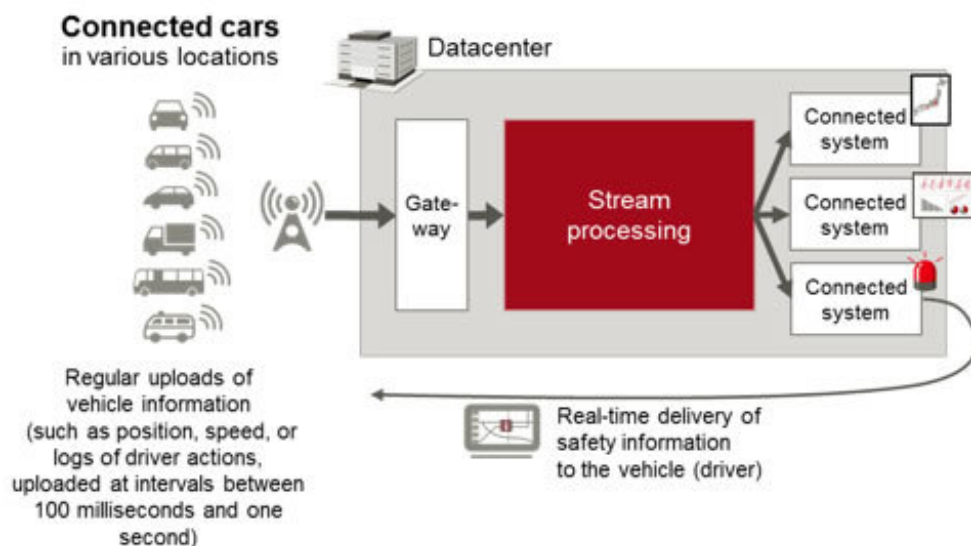
March 8 2018

Figure 1: Real-time services via the collection, analysis, and utilization of data from connected cars. Credit: Fujitsu

Fujitsu Laboratories Ltd. today announced the development of the Dynamically Reconfigurable Asynchronous Consistent EveNt-processing Architecture (Dracena), a stream processing architecture that can add or change content while processing large volumes of IoT data, without stopping. With recent advances in IoT technologies, it is expected that many real-time services will be created to utilize the large volumes of data flowing into the cloud from various devices across

factories, homes, and social infrastructure. In the progression towards autonomous driving with connected cars, researchers are considering the analysis of the vast amounts of information, such as speed and location, generated from vehicles, which can then be presented to drivers, in the form of warnings, for example.

Stream processing technology, which is effective in the high-speed processing of these sorts of huge volumes of data, has issues in that, because processing must be temporarily stopped when changing or adding processing content according to additions or improvements to services, the provision of services can be delayed. Now, Fujitsu has developed a new stream processing architecture that automatically switches to a newly provided data processing program when a parallelized data processing job has been completed, by separating stream processing into data reception processing and actual data processing so that data reception processing and current data processing are not stopped (patent pending).

As a result, in a simulation of the reception of a few dozen bytes of data per second from one million vehicles, Fujitsu has confirmed that this architecture is able to continue processing streaming data while adding or changing processing programs, with an average delay increase volumes of five milliseconds or less. Fujitsu Laboratories is looking to commercialize this technology during fiscal 2018 on the Mobility IoT Platform, offered by Fujitsu Limited, and extend it to other industry areas. Details of this technology were presented at DEIM2018 (the Forum on Data Engineering and Information Management), a conference being held in Awara, Fukui Prefecture, Japan, from March 4.

Development Background

With the recent development of IoT technologies, data has begun to be gathered from all sorts of objects and collected in datacenters, and it is

expected that by analyzing and utilizing this, a variety of new services will be created. In the case of connected cars, for example, it is thought that by collecting, analyzing, and utilizing data from automobiles in real time, it will be possible to relieve congestion, assist drivers, and improve the safety of autonomous driving (figure 1).
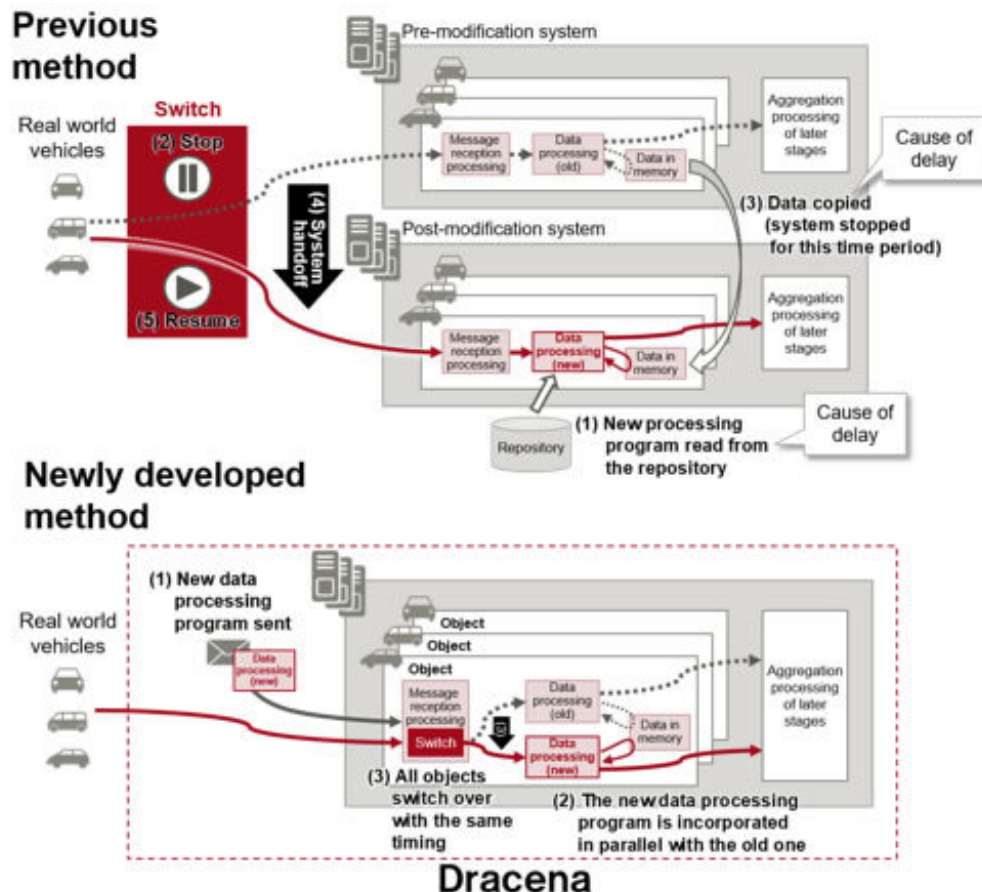

Figure 2: Differences between the existing technology and Dracena's non-disruptive update technology. Credit: Fujitsu

In order to rapidly process data, such as speed and location, that are generated on a second-to-second basis by huge numbers of cars in

motion, the most effective method is to construct a system that uses stream processing to process data in parallel, such as on a car-by-car basis. In order to add to or change the processing program according to service additions and improvements, the current method involves preparing two systems of the same scale in advance, using one for operations, making changes to the other one, and then quickly swapping them out. This method required both systems to be temporarily stopped, however, while the data, such as the speed or position of a car, held in the memory of the system in use, was copied over to the revised system. This made it difficult to produce services that required truly continuous operations, such as the real-time transmission of warnings to connected cars. In addition, because new processing programs were obtained from the database, known as a repository, congestion resulted with the numerous queries from large volumes of processing units, delaying overall processing.

Details of the Newly Developed Technology

Now, Fujitsu Laboratories has developed Dracena, an architecture that can modify the processing programs of a system while it is operating, without halting operations. With this technology, when changing or adding data processing contents, this architecture distributes the new data processing program as a message, in the same way data is distributed, to each individual processing unit, called an object, such as the processing unit for each car. This eliminates the impact on overall processing speed due to the concentration of queries on the repository. Moreover, by separating intra-object message reception processing and data processing in this architecture, the system is able to add in the new data processing program without stopping the message reception processing or the existing data processing, and then have all objects change over to the new data processing program with the same timing. This has enabled Fujitsu Laboratories to create a stream processing architecture in which the data processing program can be added to or

changed without stopping, in order to continue parallelized processing without holding back the flow of huge volumes of data for copying (figure 2).

The results of a simulated evaluation confirmed that, in a use case in which a few dozen bytes of data are transmitted once each second from one million vehicles, this architecture was capable of continuously providing services when adding a sudden-braking detection service in a situation where the system was already providing a service to detect excessive driving times, with an average delay increase volume of five milliseconds or less. This architecture will enable the rapid provision of real-time services that require uninterrupted operation and that can respond to problems occurring in society, including providing driving assistance for connected cars, supporting energy-saving usage of appliances, providing in-home health and safety monitoring, and providing travel guidance for tourists using smartphones. Moreover, this architecture enables users to adopt a build method in which they first build a base system aimed at simple analysis and utilization, and then gradually add new services. Using this technology in the case of automobiles, for example, it would be possible to begin with a system that reads signs of drunk driving based on steering wheel operation data, and then add new services layer by layer, such as combining this with map data to detect crosswinds at tunnel exits, or combining it with image data to detect the presence of illegally parked cars, which can be expected to improve the efficiency of service development.

Fujitsu aims to commercialize this technology during fiscal 2018 as a constituent element of the Mobility IoT Platform offered by Fujitsu Limited. In addition, Fujitsu is looking to extend this technology beyond the mobility field to business areas that require real-time services based on data that is continually generated at a high frequency, such as providing directions to people during events or in disaster situations.