

Cluster-based distributed controller technology for failure-tolerant networking

June 5 2014

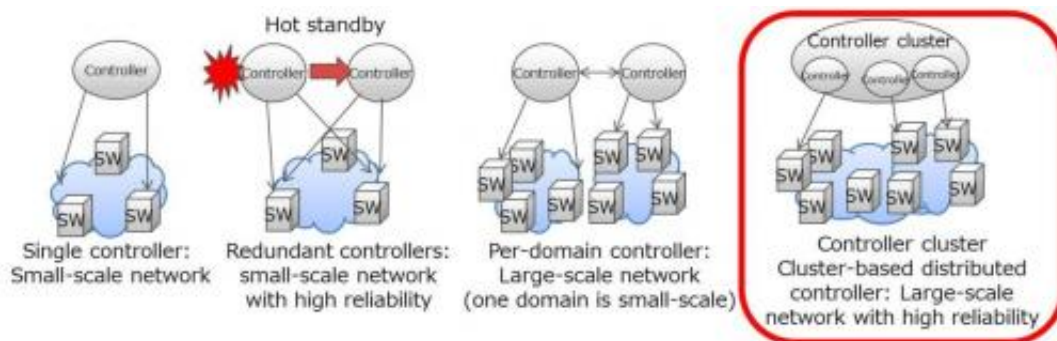


Figure 1: Controllers and network scale

Fujitsu Laboratories today announced that it has developed technology for cluster-based distributed controllers in large-scale networks that implements a wide-area software-defined networking (SDN) and that can automatically handle controller failures and load fluctuations. A cluster-based distributed controller runs on multiple physical controllers as a single logical controller to control multiple network switches.

Compared to conventional centralized controllers, cluster-based distributed controllers offer better scalability and improved failure tolerance. Until now, however, the problem was that they had difficulty handling sudden load Fujitsu Laboratories Ltd. today announced that it has developed technology for cluster-based distributed controllers in large-scale networks that implements a wide-area software-defined

networking (SDN) and that can automatically handle controller failures and load fluctuations. A cluster-based distributed controller runs on multiple physical controllers as a single logical controller to control multiple network switches.

Compared to conventional centralized controllers, cluster-based distributed controllers offer better scalability and improved failure tolerance. Until now, however, the problem was that they had difficulty handling sudden load fluctuations and coordinated control when there was a controller failure. Now, Fujitsu Laboratories has developed a distributed controller module for the coordinated control of multiple controllers, a load-balancing technology that transfers a switch being managed by one controller to another in a matter of seconds when a controller is under increasing load or has a failure, and an uninterrupted recovery technology.

These technologies enable SDNs to work reliably when traffic rises beyond initially expected levels, or when multiple controllers have failures. By deploying an SDN with these technologies to a wide-area network, infrastructure can recover quickly from disasters or other network failures while maintaining steady network operations. These technologies are being presented at Interop Tokyo 2014, opening June 11 at Makuhari Messe in Chiba, Japan.

Background

Existing SDNs such as OpenFlow(1) are designed for centralized control, which means that operating wide-area networks, configured with switches transferring large volumes of communication packets, as SDNs results in highly concentrated loads in the controller when the number of users increases. This will be an obstacle to the smooth provision of service, and if the controller itself fails, the switch that it had been managing can no longer be controlled. Fujitsu Laboratories solved these

problems by treating multiple physical controllers as a single logical controller that can handle centralized control of thousands of switches. This is accomplished through a proprietary cluster-based distributed controller technology (Figures 1, 2). This technology consists of a module for control applications that is an add-on to existing controller applications, and a distributed controller module that connects multiple distributed controllers as components of an OpenFlow controller so that, depending on loads, application and controller components can be added along with server resources.

Issues

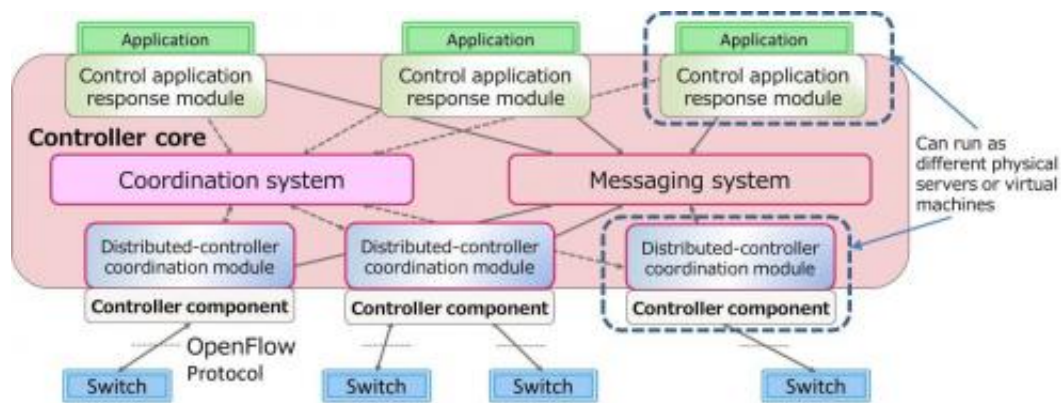


Figure 2: Cluster-based distributed controller overview

Cluster-based distributed controllers are different from centralized controllers in that multiple distributed controller modules need to be run in a coordinated way so that they do not compete with each other. Another challenge is ensuring continuity of control. Processes need to keep running even if a module fails, but difficulties are encountered with automatic switchovers when some controller components are heavily loaded or fail, and processing by the switches managing the controllers

slows down or control becomes unsustainable.

About the Technology

Fujitsu Laboratories has developed a load-balancing technology that automatically redistributes control loads in a cluster-based distributed controller, and a recovery technology that automatically reassigns controllers without interruption when one fails.

Load-Balancing Technology

Fujitsu Laboratories has developed a load-checking function as a new addition to the distributed-controller coordination module (Figure 3). This collects load information from each controller component (such as CPU utilization rate and number of switches) (step 1), and the coordination system periodically checks load information using one distributed-controller coordination module chosen as the "leader" based on module control number or other criterion (step 2) to detect load imbalances. If load rebalancing is judged to be needed according to the load-balancing logic, which switches to be reassigned are decided based on switch-reassignment logic, to balance the load according to a policy for CPU utilization rates and number of switches (step 3). As a result, the correspondence between the changed switches and the controllers is registered in the coordination system (step 4), and the load is balanced by reassigning the switches in accordance with the updated information from the distributed-controller (step 5).

Uninterrupted Recovery Technology

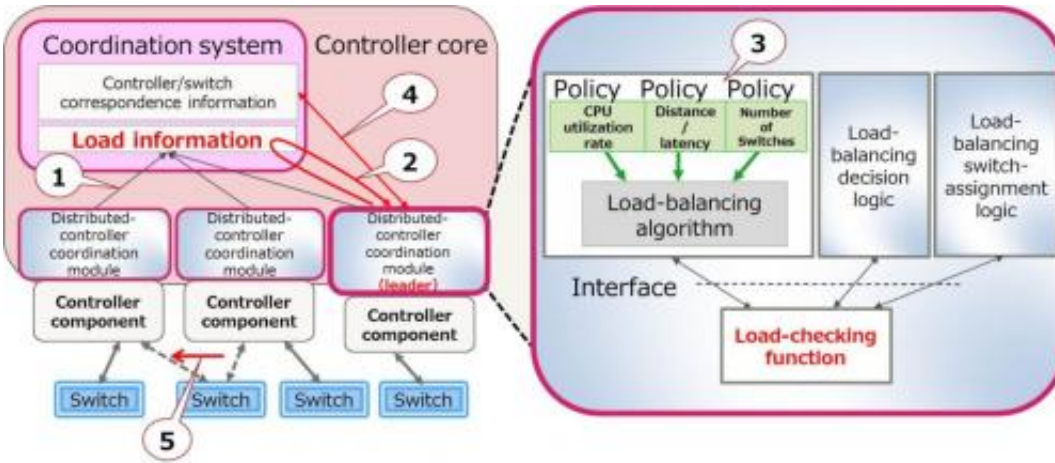


Figure 3: Load-balancing technology overview

Fujitsu Laboratories has developed a new failure-checking function for the distributed-controller coordination module (Figure 4). The distributed-controller coordination module chosen as leader detects a failure in a controller component (steps 1, 2) and determines a new controller component to manage the switches connected to the failed controller (step 3). This changes the controller/switch correspondence information to redistribute loads automatically based on controller-component load information (CPU utilization rates and number of switches) (step 4). The distributed-controller coordination modules that have not failed link to the information update and activate it to reassign the controllers managing switches (step 5) so that operations continue without any interruption in service. Because the controllers that are the reassignment destinations are decided using load-balancing technology, no controller should experience a sudden load spike that would cause it to shut down. Furthermore, even if the leader module itself suffers from a failure, the coordination system will detect a session interruption and select a new leader, and that leader module will determine controllers to manage switches again.

Results

Using the cluster-based distributed controller makes it possible to handle sudden load fluctuations and to maintain continuity of network services even when controllers fail, enabling stable, highly reliable operations of wide-area networks. For example, in the case of conventional controllers, when they are duplicated in the hot standby mode, i.e. active and on standby, for a ten-domain network, the total required number of the controllers is 20, or specifically two per domain. By contrast, using cluster-based distributed controllers, just one standby controller is added to the regularly running ten controllers, so that only 11 controllers are needed, enabling a reduction in the number of controllers by nearly half.

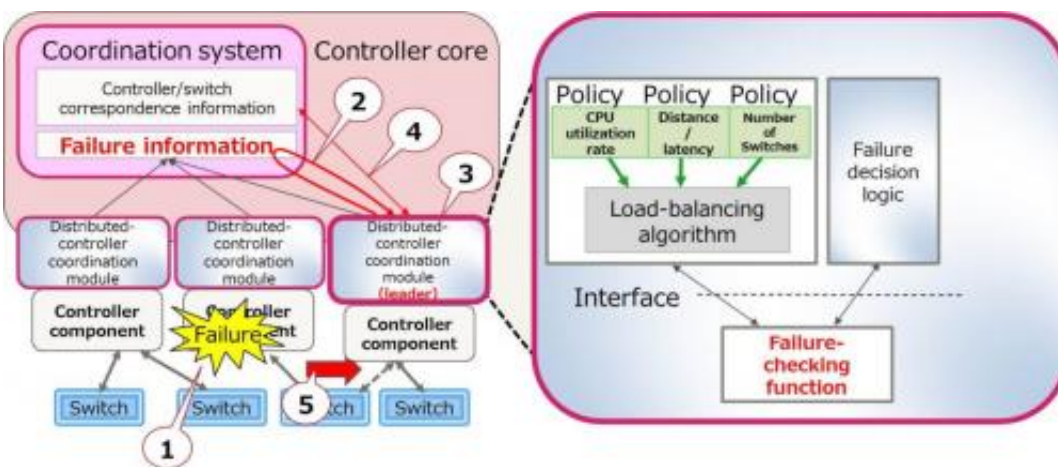


Figure 4: Uninterrupted recovery technology overview

Future Plans

This technology could be used in the networks of telecommunications carriers and other network infrastructure to achieve highly reliable, stable operations with lower deployment costs and lower operating costs.

Fujitsu Laboratories is continuing with research and development on control technology for cluster-based distributed controllers with the goal of a practical implementation in fiscal 2015.

Provided by Fujitsu

Citation: Cluster-based distributed controller technology for failure-tolerant networking (2014, June 5) retrieved 26 April 2024 from <https://phys.org/news/2014-06-cluster-based-technology-failure-tolerant-networking.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.