

FalconEye[†]: Data Center Status Extraction via Vision Transformation Techniques

Ki-Woong Park, Woomin Hwang, and Kyu Ho Park
 Computer Engineering Research Laboratory
 Korea Advanced Institute of Science and Technology
 {woongbak, wmwang, kpark}@core.kaist.ac.kr

Abstract

In this study, a vision monitoring system that is applicable to the maintenance of data centers was developed and applied to data center status extraction. The vision monitoring system, which is intended to complement system monitoring tools, such as IPMI and Nagios, has the additional benefit of enabling continuous monitoring of the external status of data centers. This system, which is based on vision transformation techniques, involves three main steps: camera calibration, where the characterized physical point is determined in the data center by the setting of the system parameters; the development of vision transformation subroutines, which are aimed at transforming relatively large images into vertically expanded image streams; and the development of image analysis subroutines for the purpose of investigating images that are transformed for the data center status extraction. This *FalconEye* system of this work is implemented with the aid of a cloud computing platform, called *iCube Cloud*¹.

1. Introduction

Data center monitoring has evolved from a simple alert system into a critical element that ensures maximum availability of data centers [5]. In an ideal world, software and hardware would be deployed in a data center and simply work without any problems. However, reality poses a considerable challenge and involves resource-intensive tasks because of the need for around-the-clock monitoring and management to safeguard the increasingly complex environment of data centers. Data centers or server administrators must keep a watchful eye on the data and make sure the various systems remain stable and available [3]. For a small server group system with only a handful of servers, an administrator can just periodically do a manual check of the status. However, for data centers, where there are thousands of servers to monitor at any given time, the monitoring and alerting process pose a considerable challenge [12].

¹ <http://www.icubecloud.com>: Cloud service platform & testbed of this work




Figure 1. *iCube Cloud* Data Center Map and Image

Although sophisticated monitoring and alerting capabilities are common in physical equipment, such as the uninterruptible power supply (UPS) and fire suppression systems, the vision monitoring of data centers is a critical element for maintaining maximum availability. In the construction phase of a data center (which is known as the system layer development), the external status of each server platform is indicated by blinking Light-emitting diodes (LEDs). The status of this type of system must be viewed holistically and watched proactively for the sake of system administration and construction. Because the targets of such monitoring include the storage, network, and power status of the data center, considerable information is collected for such tasks as system stabilization, maintenance, and debugging.

As a remedy to these problems, this paper presents a vision-based monitoring system called *FalconEye*. *FalconEye* detects the current status of each server by transforming the captured image. The transformation involves a series of three operations: firstly, sub-images that indicate the system status are extracted from the captured image; in the next step, the cropped images are stripped horizontally; and, finally, the horizontal images are expanded into vertical streams for a certain duty cycle (of 200 ms). The three steps facilitate the analysis of the image streams for the data center status extraction.

The remainder of the paper is organized as follows: In Section 2, we discuss background issues and relevant works. In Section 3, we present the overall system design and components of the proposed monitoring system. In Section 4, we evaluate the performance of the proposed monitoring system. Finally, in Section 5, we present our conclusions and suggested future works.

2. Background and Related Works

A data center monitoring system that ensures maximum availability of data centers has been actively studied and developed in the research area of grid or cloud computing. The technology for monitoring the status of a data center goes back to the days of centralized mainframes and includes such practices as walking around with thermometers and relying on a data center administrator to investigate the spatial environment of the data center [5]. Although conventional system monitoring tools provide a well-defined interface for checking the system health and status, many limitations exist with the traditional monitoring solutions in use today including compatibility, operational limitations and cost. The limitations due to compatibility generally stem from additional hardware requirements or a heterogeneous interface. For instance, IPMI (Intelligent Platform Management Interface) [1] requires a certain compatible interface equipped with baseboard management controller for interaction between system management software and platform hardware [9]. In terms of the interface, the use of a heterogeneous interface such as OPC [6], SNMP [2], Modbus [8] and BACnet [10] for data center monitoring may be one of the limitations. The use of a widely varied interface greatly complicates the development of a solution that enables different devices to be monitored with disparate protocols because such a solution would require additional costs and hardware, such as a protocol converter.

Another limitation with traditional monitoring systems is the operational restriction. *Nagios* [7], for example, provides a system and network monitoring interface at the application layer. That could be unsuitable solution of data center monitoring for cloud computing environment. Because in that approach virtual machines are dynamically

turned on and off, it will severely hamper the usability of the monitoring solution.

Data center monitoring solutions must do more than simply alert data center management of immediate problems. They must actually enable real-time data to be captured and analyzed so that management can make intelligent decisions about the data center infrastructure. In this paper, we present a vision monitoring system can be used in the maintenance of data centers. The vision monitoring system, which is intended to complement the above-mentioned system monitoring tools, has the additional benefit of enabling continual monitoring of the external status of the data center without the complication of compatibility issues or additional hardware requirements.

3. Internal Analysis of *FalconEye*

3.1. Study Site, *iCubeCloud*

Figure 1 shows the deployment of *FalconEye* as it monitors the *iCube Cloud* Data Center (located on the campus of the Korea Advanced Institute of Science and Technology (KAIST)) [4]. Four networked cameras are used to periodically collect images of the data center. The *FalconEye* monitoring system is deployed at a certain location to extract information on the status of the data center. The computing environment is used for the research and development of a peta-scale cloud computing platform. At this site, several companies and three research institutes (Seoul National University, the University of Paris VI, and KAIST) use the platform for research activities and for commercialization development, particularly with regard to semantic search engines, searchable media broadcasting services, and even the platform itself. The platform is also used in some courses by a number of students who need an individual hosting machine [11]. Thus, the aggregated computing platform accommodates massive participation, collaboration, and content creation by unpredictable end users. The realization of full cloud computing services requires around-the-clock monitoring and management. Before this study, we were involved in the monitoring and logging of the data center status for the purpose of ensuring a stable cloud service. The process, which requires both human and IPMI [1] participation three times a week poses a considerable challenge and resource-intensive activity, especially from the point of physical monitoring by humans.

3.2. *FalconEye* System Architecture

Figure 2 shows the overall architecture of the *FalconEye* monitoring system. The three major components of our architecture are listed as follows:

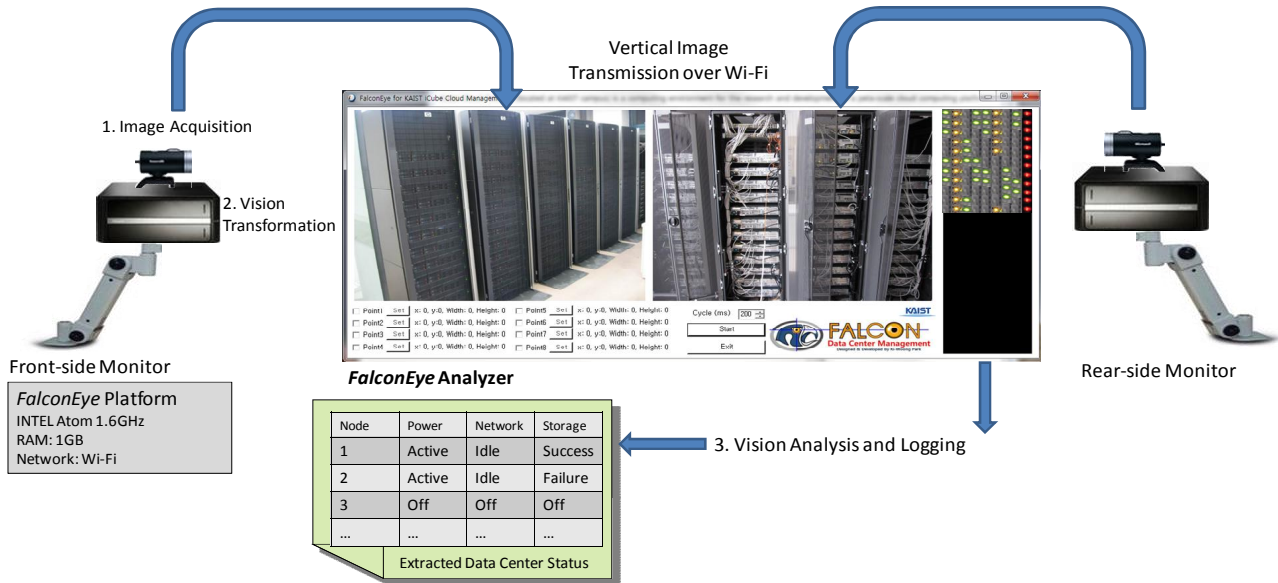


Figure 2. FalconEye Overall System Architecture

- **Front/Rear-side Monitor:** This component periodically collects images periodically (every *200 ms*) from a camera set up on a bare-bone platform as shown in Figure 2. It then transforms the relatively large images into vertically expanded image streams and transmits the transformed images to the *FalconEye* analyzer over a Wi-Fi interface.
- **FalconEye Analyzer:** This component analyzes the received transformed images and generates log information for the data center status extraction. If an abnormal status is detected, *FalconEye* automatically sends a notification e-mail with the generated log.
- **Camera Calibration:** A characterized physical point such as a network or a power LED of each server in the data center is determined by the parametric settings of the system. The parameters determine the configurations of the camera's internal parameters (such as the effective focal length and the principal point position). In this work, we configure the internal camera parameters manually by setting the position and size of the characterized physical point.

3.3. FalconEye Operating Procedure

The proposed system provides a vision-based monitoring system that can be used in the maintenance of data centers. It detects the current status of each server by transforming the captured image. The transformation involves a series of operations: firstly it extracts the sub-images that indicate the system status from the captured image; it then produces a stream of images, where the horizontal axis is a colored image for each server that is being monitored and the vertical stream is accumulated over time; and, finally, it facilitates the analysis of the image streams for the data center status extraction. The realization of this system, which is based on vision transformation techniques, involves four main steps:

- **Image Acquisition:** Images are collected from the cameras of the bare-bone platforms; the cameras are positioned in certain places and directed obliquely to the front side and rear side, respectively with a *1.6* mega pixel resolution. The bare-bone platforms are connected to the *FalconEye* analyzer by the Wi-Fi interface, and the analyzer captures images of the data center every *200 ms*. The collected images are simultaneously transformed in the next step.
- **Vision Transformation:** This process transforms the captured images (with a resolution of *1.6* mega pixels) into vertically expanded image streams. The transformation is based on the correspondence between several characterized physical points and the current time. The transformed images are then transmitted to the *FalconEye* analyzer over a Wi-Fi interface.
- **Image Analysis:** The images are analyzed for the purpose of investigating the transformed images for the data center status extraction. The data center status can be extracted from the transformed images every *200 ms*. For every received transformed image, the data center status that corresponds to the time of




Figure 3. Transformed Image & Status Monitoring

the image acquisition is determined by logging the information of the data center status extraction. If an abnormal status is detected, *FalconEye* automatically sends a notification e-mail with the generated log.

4. *FalconEye* Evaluation

Figure 3 shows a transformed image and a detected server status of the data center. By means of vision transformation, the *FalconEye* analyzer can detect the sequence of the LED status (on the temporal axis), the LED color (in terms of the RGB value), and the corresponding server location (on the spatial axis). In the results of Figure 3, each column is mapped to the main LED of a certain server node; the first row is mapped to the first 200 ms of the total monitoring time. From the color aspect, the blinking green LED shows that the corresponding node is on the *ACTIVE* state, whereas the orange LED shows the *IDLE* state, and the red LED shows that the corresponding node is on the *Power-OFF* state.

Figure 4 shows how the image analysis overhead mutates as the number of physical points increases for the data center status extraction. The number of physical

points ranges from 50 to 2000. For an evaluation of the performance characteristics, the *FalconEye* Image Collector (CPU: *Intel Atom 1.6 GHz*, RAM: 1 GB) is connected with the *FalconEye* image analyzer (CPU: *Intel Core2Duo 2.2 GHz*, RAM: 2 GB) over an 802.11g Wi-Fi network. For the native approach, the analysis overhead increases drastically as the number of physical point increases, mainly as a result of the computation-communication overhead between the camera and the image analyzer because of the need to transmit and analyze the overall image for every image capture. In the case of *FalconEye*, there is a much smaller variation in the image analysis overhead because *FalconEye* mitigates the computation-communication overhead between the camera and the image analyzer, which is enhanced with the vision transformation technique.

Figure 5 shows the accuracy deviation between the extracted data center status and the real data center status as the image capture cycle is varied. The accuracy decreases drastically after 210 ms because the status of the LEDs (that is, whether they are blinking or ON position) cannot be detected accurately with a relatively longer image capture cycle (greater than 210ms).




Figure 4. Comparison of Image Analysis Overhead




Figure 5. Accuracy Analysis with Varying Capture Cycle

5. Conclusion and Further Work

Our aim was to provide a full-fledged data center status extraction system tailored for a cloud computing data center. To accomplish this aim, we describe the development of a video-monitoring system enhanced by a vertical image transformation based on a low-cost, bare-bone platform and camera equipment. The results were promising because the system can extract the data center status from a vertically expanded image stream of characterized zones for data center monitoring.

From the respect of economic issue, *FalconEye* makes it possible to construct an economic monitoring system than other traditional solutions. One *FalconEye* image collector can cover 80 server nodes. One bare-bone machine can accommodate maximum four *FalconEye* image collectors. In addition, *FalconEye* does not require any hardware and software module equipped with each server node.

In the next step of this study, we plan to explore a more accurate image analysis mechanism so that we can minimize or eliminate false positives because it is important to watch over systems in the data center and make sure that everything is functioning properly.

6. Acknowledgments

The authors wish to thank their anonymous referees for all of their invaluable comments and suggestions. They would also like to thank Chulmin Kim and NexR Co. Ltd for their help in conducting some of the experiments in this study. The work presented in this paper was supported by MKE (Ministry of Knowledge Economy, Republic of Korea), Project No. A1100-0901-1710.

7. References

- [1] Intel Press, "Intelligent Platform Management Interface (IPMI) Specification", IPMI Specification Document Ver 2.0, 2004
- [2] RFC 1065, "Structure and Identification of Management Information for TCP/IP-based internets", <http://tools.ietf.org/html/rfc1065>
- [3] Ristenpart T., Tromer E., Shacham H, and Savage S., "Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds", 16th ACM Conference on Computer and Communication Security (CCS) 2009.
- [4] NexR Co. Ltd. "Cloud service platform & testbed", <http://www.icubecloud.com>
- [5] C.Cowan and C.Gaskins, "Monitoring Physical Threats in the Data Center", WhitePaper #102 APC Press, 2010
- [6] OPC Foundation, "OPC UA For IEC 61131-3 Companion Specification", <http://www.opcfoundation.org/>
- [7] Nagios Core Development Community, Ethan Galstad, "Nagios Core Version 3.0 Documentation", <http://www.nagios.org>, 2010
- [8] Modbus Org, "Modbus Specifications and Implementation Guides", <http://www.modbus.org/specs.php>
- [9] Aperture Press, "Current Trends in Data Center Monitoring", White Paper, Aperture Technologies, Inc. 2008
- [10] BACnet Org, "BACnet, Building Connectivity", White Paper, Real Time Automation 2008
- [11] CCI:U (Cloud-Computing Initiative for Universities), "Introduction to CCI:U", <http://www.cciu.or.kr/>
- [12] Greenberg, A., Hamilton, J., Maltz, D. A., and Patel, P., "The cost of a cloud: research problems in data center networks", SIGCOMM Comput. Commun. Rev. 39, 1 (Dec. 2008), 68-73.