





Towards Real-Time Monitoring of Data Centers Using Edge Computing

Brian Setz^(✉)  and Marco Aiello 

Department of Service Computing, University of Stuttgart, Stuttgart, Germany
{brian.setz,marco.aiello}@uni-stuttgart.de

Abstract. Introducing the Internet of Things paradigm to data centers enables real-time monitoring on a scale that has not been seen before. Real-time monitoring promises to reduce the data center's operational costs and increase energy savings. As data centers can house over a hundred thousand servers, the potential number of data points that can be collected every minute is in the order of hundreds of millions. In this work-in-progress paper, we stipulate about the impact that real-time monitoring of data centers has on the network infrastructure, and demonstrate that the impact is indeed significant enough to disrupt the data center's network. We therefore propose a preliminary solution based on edge computing that minimizes the load on the network when performing real-time monitoring of a data center.

Keywords: Data centers · Real-time monitoring · Edge computing

1 Introduction

Data centers form the backbone of the modern Internet, it is their computational resources that enable many of the services that are present on the World Wide Web today. Modern data centers are massive in size, covering areas of tens of thousands square meters, housing many thousands of individual server racks [3]. It is therefore not surprising that data centers are responsible for almost 3% of the energy consumption in the United States [9]. Monitoring data centers assists in improving the energy efficiency by discovering comatose or zombie servers. These comatose servers are performing no useful work, yet still consume energy. It is estimated that up to 30% of servers are comatose [10]. Monitoring is also critical in preventing outages, which can have a wide-spread global effect [14]. Preventing outages is also critical for upholding the Quality of Service as is specified in the Service Level Agreements. Furthermore, monitoring also aids the expansion planning process of data centers by predicting future cooling and space requirements as the data center grows.

The presented research is funded by the Netherlands Organisation for Scientific Research (NWO) in the framework of the Indo-Dutch Science Industry Collaboration programme with project NextGenSmart DC (629.002.102).

The emergence of the Internet of Things (IoT) paradigm enables monitoring of data centers at a scale that was not possible in the past. A wide variety of hardware and virtual sensors can be utilized to collect different types of data, which in turn can be used to evaluate dozens of sustainability and performance metrics [12]. As a result, the amount of data that can be collected in this environment is of massive proportions: a data center of 100 000 servers, each of which report 50 distinct metrics every second, would result in 300 000 000 data points every minute. Collecting data at such fine granularities enables the real time monitoring of the data center in its entirety. However, if this data would be collected at the high frequencies required for real-time monitoring, a different problem arises: the quantity of transmitted data would be sufficiently large to negatively impact the data center’s network infrastructure.

The question we pose in this work is: *how can we leverage a data center’s network infrastructure to efficiently monitor a data center in real time by utilizing the edge computing paradigm?* With the goal of answering this question, we first analyse the common network architectures found in data centers. Next, we look at the potential data sources that can be found in a data center in order to determine the size of the raw data and the required network throughput. This is followed by a preliminary design of an edge-based data collection platform that takes advantage of a data center’s network infrastructure to reduce the load on the network. Finally, we discuss the results we have obtained thus far, as well as the steps we have planned for our future research.

The remainder of this paper is organised as follows. Section 2 introduces the related work, followed by a description of data center network architectures in Sect. 3. An analysis of the data that can possibly be collected in a data center is made in Sect. 4. Next, in Sect. 5 the proposed edge-based architecture is introduced. Followed by the conclusion in Sect. 6. We note that, since this is a work-in-progress paper, there is no evaluation section.

2 Related Work

Real-time monitoring is a technique that has become more popular in several domains with the emergence of the Internet of Things. In smart grids, for example, real-time monitoring promises to assist in the prevention of severe safety accidents by automatically identifying threats. The authors of [6] identified that real-time monitoring of smart grids would cause an increase in data that would be too large to handle using the traditional cloud computing paradigm. In their solution they introduce edge computing as a key component of their real time monitoring solution, reducing the network load by more than 50%.

In our previous work, a data set of 2.5 billion data points was collected from a data center [11]. A total of 13 different data types were collected from more than 160 servers every 10 s. The type of data collected includes CPU temperature and utilization, network utilization, air temperatures, power consumption, and more. The data is used to train models that can estimate the status of a server. This work provides a glimpse into the potential amount of data that can be collected in a data center.

The authors of [8] developed a method for real-time monitoring of data centers using an IoT approach. The data is environmental data, such as the temperature and the humidity level. These values are collected every 10 s, a total of 1.4 million data points were collected. Their IoT platform is based on a simple web service which accepts data collected by the custom-built sensors. The monitoring takes place on the level of individual racks.

In the work of [7], an approach is proposed for monitoring a data center in real time using low-power wireless sensors. The collected data includes temperature, humidity, airflow, air pressure, water pressure, security status, vibrations, and the state of the fire systems. The need for collecting data from servers for monitoring purposes is also recognized. The authors envision that some type of IoT platform is required for the collection, processing, storage, and management of the data. The envisioned platform is not designed or implemented.

In [5], the authors describe the role that edge computing has in the Internet of Things. They propose a layered model in which millions of IoT devices connected to thousands of edge gateways, which in turn connect to hundreds of cloud data centers. The authors also recognise the need for data abstraction, which uses edge gateways to reduce the volume of the raw data before sending it to the data center. However, deciding the extent by which the data should be reduced is an open problem, according to the authors.

The related work shows that some effort has been made to introduce hardware sensors and virtual sensors to data centers. The type of data that has been collected thus far is limited, however. In this work, an architecture is described that allows the collection of a much wider variety of data. None of the related work consider the increased network load when introducing real time monitoring to data centers.

3 Data Center Network Infrastructure

Data centers are facilities containing large amounts of computational, storage, and networking resources. These resources are mounted in 19-inch racks, which are metal enclosures with standardized dimensions. The capacity of a rack is expressed in Rack Units (U), and determines the quantity of equipment it can house. The standard full height rack is 42U tall. Rack equipment such as servers and switches often occupy between 1U to 4U of space, with blade server enclosures consuming up to 10U of space. Efficiently connecting all the rack equipment to the network can be a challenge, and the design of the data center network affects the networking efficiency at which the connected equipment operates.

The most widely used network architecture in data centers is the *3-layer data center network architecture* [2], shown in Fig. 1. As the name suggests, this architecture consists of 3 distinct layers: a core layer at the top, an aggregation layer in the middle, and an access layer at the bottom. Equipment, such as servers, that require network access are connected to the access layer, usually with 1 or 10 Gigabit links. The access layer is commonly implemented as a network switch located at the top of a rack (ToR switch) or at the end of a row

of racks (EoR switch). The aggregation layer aggregates the different ToR and EoR switches, to enable network connectivity between racks. The links between ToR and EoR switches are commonly 10 or 40 Gigabit. The aggregation layer switches all connect to the core layer, these links can often be up to 100 Gigabit. The core layer is responsible for providing uplinks to the Internet.




Fig. 1. An example of a 3-layer data center network architecture.

There are also other network architectures currently in use in data center, such as Facebook’s data center fabric approach [4]. This approach is similar to approaches taken by Google and eBay. The notion of a server pod is introduced, which is essentially a standalone cluster consisting of racks and servers, containing up to 48 ToR switches and 4 special fabric switches. These fabric switches are responsible for interconnecting the servers in a single pod. To connect different pods, a network spine is introduced consisting of up to 48 spine switches per spine plane. This approach is highly scalable, as computational resources can be increased by introducing more pods, and the network capacity can be increased by introducing more spine planes.

Another approach is the Fat Tree data center network [1]. This approach is similar in design to the 3-layer approach, but provides guarantees regarding the available bandwidth for each server in a rack. This is done by carefully planning the numbers of switches in each layer, and increasing the number of links between individual switches the higher up the hierarchy they are. Any horizontal slice in the network graph has the same amount of bandwidth available.

Despite the significant differences between the available data center network architectures, they all contain an access layer with ToR and EoR switches in one form or another. As we show later in our proposed architecture, these ToR switches are excellent candidates to become edge gateways due to their proximity to the servers that are being monitored.

4 Impact on Network Load

To understand the significance of the additional load that is associated with real time monitoring of a data center, a number of steps have to be taken. First, the number of servers per rack and the number of racks per data center have to be identified. Next, the data types that can be collected from a server have to be investigated, as well as their data size. And finally, the load on the network that is generated by real time monitoring has to be calculated.

The number of servers that can be placed inside a rack is not only limited by the size of the servers, but also by the data center's cooling capacity and power limitations. A standard full height rack offers space for up to 40 servers, leaving 2U for other equipment. In practice this number is between 25 to 35 servers per rack. Using blade servers, the density of a rack can be increased much further. A typical high performance 10U blade server enclosure contains 16 servers. This increases the density to 64 servers per rack. There are also 3U blade servers enclosures for low performance blade servers that house 20 blade servers. This results in a maximum density of 260 servers per rack. In all cases at least 2U are left for the ToR switch and a Keyboard Video Mouse switch.

The largest data center in the world is China's Range International Information Group data center, covering over 500 000 m². More commonly, data centers are between 10 000 and 20 000 m² in size. For example, Google's Dallas data center is 18 000 m² and contains 9090 server racks [3]. Applying the previously determined server density numbers, it can be extrapolated that a data center containing 9090 server racks can house anywhere between 318 000 and 2 363 400 servers. A report from Gartner estimates that Google had around 2.5 million servers in July 2016, spread across 13 data centers, which equates to around 192 000 servers per data center [13].

There are two categories of sensors required to monitor a data center: hardware sensors and virtual sensors. The hardware sensors are used to monitor the temperature and humidity, as well as power consumption. These measurements can be made on a global level for the whole data center, as well as on individual server level. Virtual sensors are software-based sensors, they can be agents interacting with the operating system to gather information about the CPU, memory, networking interfaces, storage devices, and more. There are software agents available that can collect and publish this type of data, popular solutions include: Telegraf, StatsD, collectd, Zabbix, Prometheus, and Nagios.

In this work, Telegraf is used to represent the virtual sensors, because of its popularity and its ability to integrate with a multitude of platforms. Telegraf is a plugin-based software solution for collecting and transmitting a wide variety of data. It consists of four plugin types: input plugins, processor plugins, aggregator plugins, and output plugins. Input plugins collect data from the system, processor plugins transform the data, aggregator plugins aggregate the data, and output plugins transmit the data to other systems. Only the input plugins that collect generic system information are included in our experiments.

To determine the bandwidth required to monitor the generic metrics measured by Telegraf, experiments are performed using a real server. The server

is a Dell PowerEdge R7425 with dual AMD EPYC 7551 32-core processors, 512 GB of RAM, and six 960 GB Intel S4510 SSD's. The operating system is Proxmox, a Debian-based virtualization environment. Telegraf is installed on the operating system and configured to collect the selected metrics. MQTT, a lightweight publish-subscribe network protocol, is configured as the output plugin. An MQTT broker is deployed on a second host. Wireshark, a network packet analyser, is also installed on this second host in order to monitor the network usage. The traces produced by Wireshark are analysed to calculate the required bandwidth for real time monitoring of a data center. An overview of the setup is shown in Fig. 2.




Fig. 2. Setup to analyse the bandwidth usage when performing real time monitoring.

To determine the load on the infrastructure, network packets were collected for a duration of 600 s. During this period, 185 400 messages were sent to the MQTT broker. In total, 55.3 megabytes of data were transmitted, an average of 92.2 kB/s. While seemingly insignificant for one server, however when we extrapolate this and use Google's Dallas data center and a rack density of 25 servers per rack as an example, the total bandwidth would equal $25 \text{ servers per rack} \times 9090 \text{ racks} \times 92.2 \text{ kB/s} = 167.62 \text{ Gbit/s}$. In practice this number is conservative, as the servers per rack density is ever increasing, and data centers are becoming ever larger.

5 Proposed Edge-Based Architecture

One method to reduce the overall load on a data center's network is bringing the computations closer to the source of the data. This reduces the amount of hops required for the data to reach their destination, and in turn limits the load to the access layer instead of overloading the aggregation and core layers. The architecture we propose is shown in Fig. 3. As each rack as a ToR switch, the goal is to leverage the computational power of the switch to turn it into an edge gateway. Every edge gateway is responsible for processing and analysing the data of their rack only. Therefore, the edge gateway would only have to handle the network traffic of a limited amount of servers. The network load for the gateway ranges between 18 Mbit/s and 47 Mbit/s, for 25 servers and 64 servers per rack respectively. At these loads the impact on the switch itself is minimal.




Fig. 3. Proposed edge-based architecture using Top of the Rack switches.

Because edge gateways are close to the source of the data, the network latency is also greatly reduced. This is crucial for real-time monitoring, as the data center operator should be informed as soon as possible about critical events. The edge gateway can also be used to automatically interact with the servers. For example, when a server is overheating, the gateway could inform the server to reduce the load, or even lower the frequency at which the CPU cores are operating. This allows the edge gateways to act as autonomous agents. The proposed architecture also improves the scalability of the data center. As the data center grows and more racks are placed and filled with servers, the impact that monitoring these new servers has will be minimized as the majority of the data remains at the ToR switch. It is also possible for multiple racks to be clustered together, such that the edge gateways of these racks communicate with each other in a peer-to-peer fashion. Another benefit of this approach concerns the privacy. In case a rack is dedicated to processing sensitive data, the edge gateway will ensure that monitoring data collected from these sensitive servers does not leave the rack. Or, when the data does have to be transmitted outside the rack, it is anonymised and privacy sensitive data is removed before it is sent across the network.

Using edge computing instead of traditional cloud computing to perform real-time monitoring in data centers has a number of benefits. From reducing the network load, to increasing the responsiveness, enabling autonomous control, as well as improved scalability and privacy. These advantages come at the cost of increased deployment complexity, and more complex ToR switches.

6 Conclusion

Real-time monitoring of a data center comes at a cost: the increase in network traffic is significant enough to influence the performance of a data center. We estimated the additional load that is placed on a data center's network, and have

shown that the additional load is significant. To counteract this problem, we proposed an architecture based on edge computing that enables real-time monitoring while reducing the required bandwidth, leveraging the network infrastructure of the data center by relying on ToR switches. In our future work, we aim to implement the proposed architecture and perform a quantitative evaluation of the performance of the architecture, compared to monitoring based on a traditional cloud computing approach.

References

1. Al-Fares, M., Loukissas, A., Vahdat, A.: A scalable, commodity data center network architecture. *SIGCOMM Comput. Commun. Rev.* **38**(4), 63–74 (2008)
2. Barroso, L.A., Clidaras, J., Hölzle, U.: The datacenter as a computer: an introduction to the design of warehouse-scale machines (2013)
3. Chen, T., Gao, X., Chen, G.: The features, hardware, and architectures of data center networks: a survey. *J. Parallel Distrib. Comput.* **96**, 45–74 (2016)
4. Farrington, N., Andreyev, A.: Facebook’s data center network architecture, pp. 49–50 (2013)
5. Hassan, N., Gillani, S., Ahmed, E., Yaqoob, I., Imran, M.: The role of edge computing in Internet of Things. *IEEE Commun. Mag.* **56**(11), 110–115 (2018)
6. Huang, Y., Lu, Y., Wang, F., Fan, X., Liu, J., Leung, V.C.M.: An edge computing framework for real-time monitoring in smart grid. In: 2018 IEEE International Conference on Industrial Internet (ICII), pp. 99–108, October 2018
7. Levy, M., Hallstrom, J.O.: A new approach to data center infrastructure monitoring and management (DCIMM). In: 2017 IEEE 7th Annual Computing and Communication Workshop and Conference (CCWC), pp. 1–6, January 2017
8. Medina-Santiago, A., et al.: Adaptive model IoT for monitoring in data centers. *IEEE Access* **8**, 5622–5634 (2020)
9. Mills, M.P.: The cloud begins with coal-an overview of the electricity used by the global digital ecosystem. Technical report, Digital Power Group (2013)
10. Ngoko, Y., Cérin, C.: Reducing the number of comatose servers: automatic tuning as an opportunistic cloud-service. In: 2017 IEEE International Conference on Services Computing (SCC), pp. 487–490, June 2017
11. Setz, B., Rao, G.S.V., Lazovik, A., Aiello, M.: A data-driven approach to monitoring colocation data centers. In: 2019 IEEE International Conference on Big Data Intelligence and Computing (2019, to appear)
12. Vemula, D., Setz, B., Rao, G.S.V., Gangadharan, G.R., Aiello, M.: Metrics for sustainable data centers. *IEEE Trans. Sustain. Comput.* **PP**(99), 1 (2017)
13. Zaman, S.K., Khan, A.U.R., Shuja, J., Maqsood, T., Mustafa, S., Rehman, F.: A systems overview of commercial data centers: initial energy and cost analysis. *Int. J. Inf. Technol. Web Eng.* **14**, 42–65 (2019)
14. Zeck, A., Bouroudjian, J.: Real-world experience with a multicloud exchange. *IEEE Cloud Comput.* **4**(4), 6–11 (2017)