Hindawi Publishing Corporation
Scientific Programming
Volume 2015, Article ID 371620, 6 pages
http://dx.doi.org/10.1155/2015/371620

Research Article

Network Latency in Systems Event Monitoring for Multiple Locations

Lukasz Kufel

Institute of Computing Science, Poznan University of Technology, 60-965 Poznan, Poland

Correspondence should be addressed to Lukasz Kufel; lukasz.kufel@hotmail.com

Received 9 June 2015; Accepted 30 August 2015

Academic Editor: Wan Fokkink

Network latency is one of the key parameters to consider when designing and implementing remote monitoring for security and system events. This paper describes how network latency may impact monitoring over wide area networks, especially when the monitoring system is hundreds or thousands of miles away from the monitored servers. Furthermore, an idea of local distributor is proposed to reduce the time of events data collection from multiple geographic locations.

1. Introduction

One of the key challenges for today's IT operations departments is to keep all crucial applications up and running on a 24 × 7 basis and, in case of any failure, to identify an incident's root cause. To fulfill these requirements, a monitoring platform needs to be designed and implemented. Moreover, efficient deployment of that platform can significantly reduce the costs of IT in an organization, due to the minimization of systems' downtime, and maintain a company's reputation as a result of the earlier detection of security threats [1, 2].

As discussed in our previous paper, there are three monitoring approaches that may be considered: agent-based, agentless, and proposed order-based monitoring (OBM) [3]. The agent-based approach offers in-depth monitoring but requires additional software for each monitored system. The agentless approach is a lightweight solution that uses a system's built-in monitoring protocols and technologies, such as Simple Network Management Protocol (SNMP) and Windows Management Instrumentation (WMI), but it only offers a general overview of the monitored environment. Novel OBM approach proposed by us integrates the advantages of both agent-based and agentless implementations and focuses on customization and setting up monitoring based on user demand. The key in the OBM approach is to first identify the valuable metrics of systems and applications that define their availability status and, secondly, to develop a mechanism to gather those metrics.

Our previous study was limited to one location [3]. In this paper, it is extended towards multiple geographic locations, and we demonstrate how WAN network latency may impact the overall monitoring process. In our previous work network latency was omitted due to the fact that systems events were collected using LAN infrastructure. Finally, we introduce the idea of local "Distributor" to minimize the impact of network latency when a monitoring system is deployed hundred miles away from the systems being monitored. The terms "network latency" and "network delay" are considered to be interchangeable; we use "network latency" in this discussion.

This paper is organized as follows. In Section 2 we described motivations for this study. Related work is given in Section 3. Section 4 presents definition of network latency and describes network parameters that need to be considered when designing events collection. The experimental results of events collection from multiple geographic locations will be obtained in Section 5. Proposal of Distributor concept is discussed in Section 6. Finally, summary and conclusions are in Section 7.

2. Motivation

Fast-growing international companies, as well as established organizations with offices around the world, have recently focused on providing unified IT services for their businesses. The standardization of IT services is dictated by economic factors, such as the consolidation of servers and their

maintenance costs, as well as by meeting the expectations of users, such as their desire for highly reliable and always available business applications. On the other hand, not all IT services can be moved to one central location due to the amount of transferred data over the network, network latency, data law restrictions, or legacy architecture of the applications. Examples of local services include print servers, file servers, domain controllers, domain name systems, HR systems, and applications specific to the local office.

From an organization's strategic perspective, IT departments can still provide guidance for locally running IT services, including how those services can be monitored and how operation departments are notified when support is required. The challenge we found in this area is that the current literature lacks design guidance for monitoring solutions for servers located in multiple geographic locations, sometimes thousands of miles away from the monitoring system.

3. Related Work

Many researchers have studied systems monitoring, but their scope was different than that of this paper. Scientists analyzed monitoring solutions regarding areas of network management [4], cloud monitoring [5–8], monitoring of grid systems [9, 10], and distributed systems located in one data center [11–15]. Many of these papers represent monitoring operations in logical layers, such as collection, representation, report, analysis, and presentation, as mentioned in the work of Lee et al. This paper focuses on the collection layer and discusses the impact of network latency while collecting raw measurements from systems located in multiple geographic locations.

4. Events Collection

The collection of events data from multiple geographic locations represents a nontrivial problem [11]. In this section we present definition of network latency and focus on key network parameters that should be considered when implementing monitoring system in multiple locations.

Apart from the additional costs an organization needs to spend on international or transcontinental links, there is a tangible network latency that can significantly impact the transmission of packets [16, 17]. According to Svoboda and his colleagues, network latency "*is a metric comprised of the sum of all small delay contributions along the data path between the two interfaces defined (…). Many parameters and variables influence the delay, especially if the measurement interfaces are separated by many hops*" [18]. In addition to latency, there are also other network parameters that should be considered when designing remote events collection. Table 1 presents a comparison of sample network parameters in local and remote events collection used by the monitoring system.

Network type represents how servers are connected to the monitoring system. This determines further parameters, such as *ping response/latency* and *bandwidth*, because devices are connected over a relatively short distance (LAN) or long links (WAN). Connections to systems over wide area networks require the use of Virtual Private Networks (VPNs) over public internet. When a highly reliable connection to remote offices is required, organizations may additionally use the Multiprotocol Label Switching (MPLS) mechanism, which is protocol-independent and provides a more efficient means for transmitting data.

Table 1: Network parameters in local and remote events collection.

Parameter	Local collection	Remote collection
Network type	LAN	WAN (VPN, MPLS)
Ping response/latency	Less than 2 ms	More than 2 ms
Bandwidth	100 Mbps–10 Gbps	1 Mbps–1 Gbps
Interruptions	Negligible	Multiple, external
Costs and maintenance	Internal	Third party

In the LAN type of network, *Interruptions* are negligible because the entire infrastructure is owned and maintained by the company. In this type of network, failures are usually related to hardware and configuration settings. In the WAN type of network, *Interruptions* occur much more frequently because, apart from hardware and configuration failures, the length of the network cable, multiple tenants, and exterior factors (such as ground excavation and submarine maintenances) also need to be taken into account.

Cost and maintenance is another parameter that may impact the decision regarding the geographic location of the monitoring system when considering remote events collection. WAN circuits are usually provided by third party companies known as Global Network Service Providers [19], and the WAN bandwidth is strictly correlated to service charges.

5. Experiment

To measure how network latency may impact remote events data collection, the following experiment was conducted:

(i) A sample location with servers to be remotely monitored was chosen in the USA (the state of Arizona).

(ii) Two additional locations, where the monitoring system could be installed, were identified, one in Europe (Ireland) and the other in the USA (the state of Washington).

The second goal of the experiment was to understand how far from the remotely monitored servers the monitoring system could be located to collect events at least every 5 minutes, as recommended by the US National Institute of Standard and Technology [20].

A set of 100 servers was selected from a local office datacenter, named M, that was located in Arizona in the USA. The Microsoft Windows Server operating system was used on all the servers. To simplify the test, each server had a shared folder with three files of various sizes: 1 kB, 4 kB, and 8 kB. Each file contained sample security and system events in a key equals value format (e.g., server = server_name). Access to those files was secured by a share folder and file system access control list (ACL). The second and third offices were chosen as places for running the monitoring system. The second office's local datacenter, named A, was located in Washington

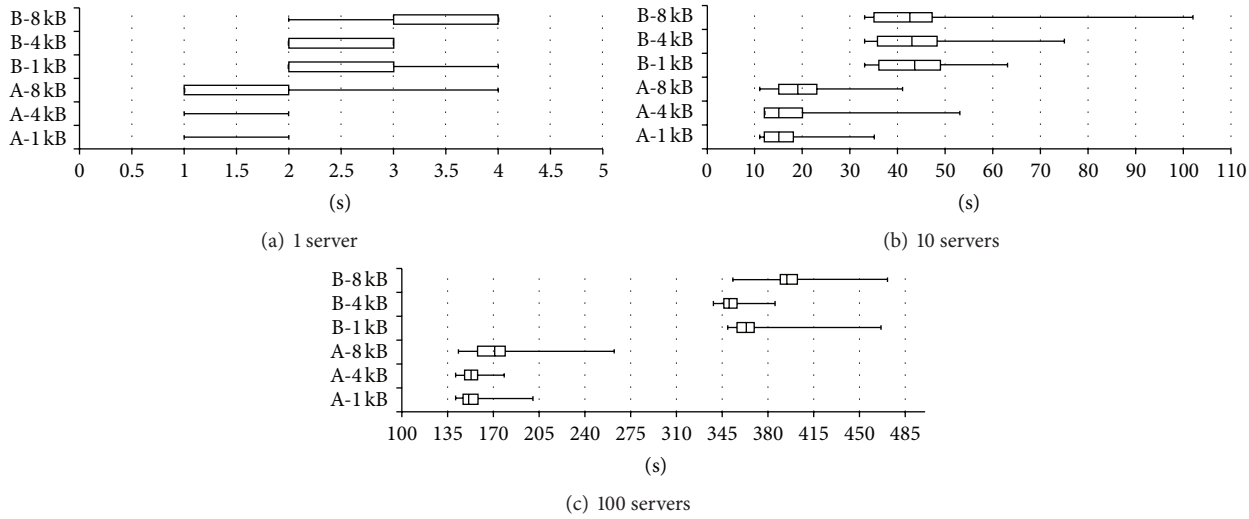

Figure 1: Results of events collection experiment. (a) The average time in seconds for events data collection from sets on one server. A-1 kB represents the time for the monitoring system server in office A (USA) for the 1 kB events file. B-8 kB represents the time for the server located in office B (Europe) for the 8 kB events file. (b, c) The average collection time in seconds from sets on 10 and 100 servers, respectively.

in the USA, and the third, named B, in Ireland, approximately 1800 km and 8000 km away from office M, respectively.

The first goal of the experiment was to measure the round trip response time from offices A and B to office M by running the *ping* command. The results of those tests are presented in Table 2.

The results from the ping command revealed that office B, theoretically located 4.4 times further away than office A, has a network latency that is approximately 3.7 times lower. Because the exact measurement of the distance that the packets traversed over the computer network is complex to calculate, we decided to evaluate the network latency by transferring the files of various sizes from office M to the monitoring system servers located at A and B. This was the second goal of the experiment.

The sequential process of copying sample event files was initiated simultaneously on the monitoring system servers in the local datacenters of both offices A and B. The copying process was executed by using the *xcopy* command to download the shared files over TCP from the monitored servers. This part of the experiment was divided into three stages. Stage one was prepared to measure the collection duration for one server, stage two for 10 servers, and stage three for 100 servers. Each stage contained three sets of sample files of different sizes. To run the experiment simultaneously for two days from both locations, each set was executed 192 times from the monitoring system server in office A (USA) and 96 times from the server in office B (Europe). During the experimental period, the WAN traffic optimization feature was disabled on the network devices.

The results (Figure 1) show that events data collection is dependent on the distance between the monitoring system and servers being monitored. They also revealed that an events data file of size 4 kB needed the same amount of time to complete the sequential collection as a 1 kB file. An 8 kB file required only approximately 10% more time to complete the copy process than 1 kB and 4 kB files. Comparison of the distance to the collection time of sample files (see Table 3) demonstrates that for the location that is more than four times further away the sample events data collection took two times longer.

Table 2: Distance and ping response times between offices M, A, and B.

	From location A (USA)	From location B (Europe)
Distance to office M, USA*	1,815 km (1,128 miles)	8,036 km (4,993 miles)
Number of hops to office M	8	8
Average ping response time	38.81 ms (63,386 samples)	144.31 ms (31,966 samples)

*Calculated using Google Distance Measurement Tool (https://maps.google.com/).

Additionally, the experiment's results prove that remote monitoring of 100 servers located in the USA from the local datacenter in the European office will take more than five minutes, regardless of the file size. This is due to the network latency and the number of connections that the monitoring system needs to establish. To accomplish a five-minute polling interval in the scenario where the monitoring system is in the European office, the number of monitored servers in the USA office is recommended to be approximately 50. Moreover, that frequency may not be achievable for every monitoring poll because the size of events data files may vary, as well as network performance, due to multiple external dependencies, such as time of day or day of the week.

To add to the experimental data, the average latency figures from one of the Global Network Service Providers,

Table 3: Average events data collection time for various file sets copied from 100 servers.

Location	1 kB collection	4 kB collection	8 kB collection	Distance	Ping
From A (USA) to M (USA)	155 sec (192 samples)	154 sec (192 samples)	170 sec (192 samples)	1,815 km (1,128 miles)	38.81 ms (63,386 samples)
From B (Europe) to M (USA)	366 sec (96 samples)	351 sec (96 samples)	396 sec (96 samples)	8,036 km (4,993 miles)	144.31 ms (31,966 samples)
Ratio B/A	2.36x	2.28x	2.33x	4.43x	3.72x

Table 4: Monthly average round trip times in milliseconds [ms] for various worldwide locations.

Link description	Feb. 15	Jan. 15	Dec. 14	Nov. 14	Oct. 14	Sep. 14
Within Europe hubs	11.65	11.71	11.78	11.74	11.85	11.68
Within North America hubs	35.79	35.77	36.45	35.83	35.39	35.74
London to New York	71.85	72.34	72.54	72.49	75.02	72.85
Singapore to Tokyo	83.42	77.95	77.33	76.51	76.47	79.42
Within Asia Pacific hubs	94.43	102.92	114.28	96.07	95.54	97.38
Trans-Pacific	109.82	109.82	109.80	109.80	109.76	109.70
India to UK	119.91	119.93	136.30	155.93	142.15	124.43
Within Latin America hubs	137.30	137.47	136.88	137.23	142.07	144.06
Australia to US	154.62	154.58	154.58	154.58	154.80	154.33
Singapore to US	182.58	178.74	182.88	180.74	181.86	175.58
North America to India	252.01	253.39	266.73	293.39	285.62	253.94
Australia to UK	296.36	335.35	315.85	290.23	286.13	305.87
New Zealand to UK	315.62	343.42	328.56	296.48	296.76	310.57

that is, Verizon, were attached. The results in Table 4 refer to monthly average round trip response times gathered by using the ping command via the Internet Control Message Protocol (ICMP). The data were collected from designated routers at key network hubs in worldwide locations in 5-minute intervals [21].

The details from Table 4 should be taken into account when designing global monitoring solutions and deciding where the monitoring system needs to be deployed. The other factor to consider is monitoring software license limitations and the availability of hardware where the monitoring system would be installed. Because network latency in remote monitoring cannot be avoided or reduced, one recommendation is to minimize the number of remote connections the monitoring system needs to establish with all monitored servers. In the basic approach, the number of connections from the monitoring system is based on the number of servers being monitored; in some cases, that number might be even higher due to multiple metrics gathered from the same server. For example, one request will gather CPU utilization, another will collect system logs from the last ten minutes, and a third may verify a list of running processes. Minimizing the number of connections would reduce the amount of SYN, SYN-ACK, and ACK messages between computers, DNS, and authorization requests; thus, the overall event collection duration will be shorter. To verify this theory, we propose the concept of a Distributor in the events collection process.

6. Distributor Concept in Events Data Collection

During the experiment, we realized that the security as well as system events data collection time is more dependent on the number of remote network connections than on the size of the files. Establishing a network connection with each individual server takes some time because of the participating components, such as the server's name resolution in DNS, authentication and authorization process to the remote resource, local- and wide-area network performance, and internet architecture when servers are geographically distributed. In this experiment, the process of establishing a network connection was performed by monitoring system located at different geographic sites.

To examine this problem, we reviewed recent articles that describe monitoring implementations for multiple geographic locations [22, 23]. Based on our research and experience, we propose the solution of a Distributor in events data collection (see Figure 2).

The Distributor is an aggregation point that minimizes the number of remote connections to monitored servers. It connects with each monitored server using only the local-area network rather than the wide-area network. Once the data are collected, only one data transfer connection to the monitoring system is required over the wide-area network. Because one data transfer connection introduces

Table 5: Number of remote connections from monitoring system to monitored servers. Data collection occurs every 5 minutes.

Scenario	10 servers	100 servers	200 servers
1 day of monitoring	2,880	28,800	57,600
1 day of monitoring with Distributor	288 (1x Distributor)	288 (1x Distributor)	576 (2x Distributor)
Reduction of connections (1 day)	90.00%	99.00%	99.00%
30 days of monitoring	86,400	864,000	1,728,000
30 days of monitoring with Distributor	8,640 (1x Distributor)	8,640 (1x Distributor)	17,280 (2x Distributor)
Reduction of connections (30 days)	90.00%	99.00%	99.00%

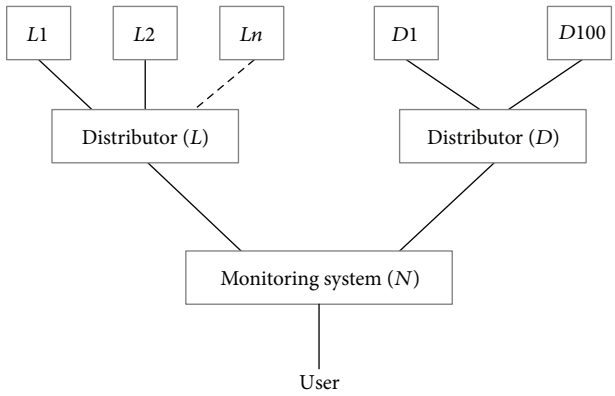

Figure 2: An example of a Distributor in events data collection. *L*1, *L*2, and *Ln* are servers placed in location *L*, while servers *D*1, *D*100 are in location *D*. Distributors are placed in local, respective locations. The monitoring system is in location *N*.

a single point of failure, an alternative link should be defined between the Distributor and monitoring system. In case any connection link is unresponsive, the Distributor is equipped with a mechanism that tries to resend the data when the network becomes available. The Distributor also introduces some delays because the data needs to first be downloaded to its local server and then transferred to the monitoring system. However, the benefit of establishing only one connection is tangible in scenarios where more than ten servers are being monitored remotely.

Additionally, in the Distributor concept, the data collected from all monitored servers is compressed before being sent to the monitoring system. Hence, there is an extra step; that is, the monitoring system must uncompress the received files. This overhead should negligibly impact the resources' utilization of the Distributor and the monitoring system but notably reduce the duration needed to transmit the data over the wide-area network because only one file is sent.

Further reductions can be achieved when a set of *thresholds* are established [4]. These thresholds will classify whether the collected event should be dealt with locally or sent immediately to the monitoring system or be transmitted in the next cumulative batch, which, for example, occurs once per hour or once per day. The threshold and event urgency may depend on an application's business criticality.

Sedlar and his colleagues discussed a similar approach to the Distributor idea in [23]. Events data collection in their work was implemented through 6500 commercial nodes providing IPTV in Slovenia. Their broadband network gateway (BNG) ran as a Distributor and gathered events from 100 local devices. All the systems were located in one country and were using infrastructure provided by one company, Telekom Slovenije. In our paper [3] sequential events data collection process from 130 nodes took on average over a three-month period 79 seconds. All nodes were located in the same location as the monitoring system and were using local gigabit network. In that case, the estimated duration for 100 servers is 61 seconds; thus, a size of 100 nodes per Distributor appears to be an appropriate choice in the proposed idea.

The Distributor concept extends the design of current monitoring approaches from being based in one location or in one country to be available in multiple locations, countries, and continents. It also gives a mechanism to resend the data when network link becomes unresponsive, for example, due to the fact that transcontinental links would always have dependencies on various network providers. Implementation of Distributor idea for 10 and 100 servers reduces the number of connections by 90% and 99%, respectively (see Table 5 for details).

7. Summary

Events data collection from multiple locations is aligned to organization growth and expansion to new markets. These changes encourage IT departments to review their support operation procedures and introduce new requirements. One of the aspects to consider is the expansion of existing monitoring solutions.

In this paper, we discussed the impact of network latency on security and system events data collection when the monitoring process is in operation for multiple geographic locations. The results gathered in the experiment show that simple tools such as *ping* and *xcopy* can be used to measure wide area network capacity and performance. The same tools can also help in deciding where monitoring system should be deployed when designing remote servers monitoring. The experiment also highlights that network latency has a tangible impact on establishing multiple remote connections and may significantly slow down the process of system failure notification.

The idea of an aggregation point, described in this work as the Distributor, should be considered when the current monitoring solution needs to monitor servers and applications located in remote offices. The aggregation would minimize the impact of network latency in the entire monitoring process because a smaller number of remote connections will be established and only one compressed file will be transferred at each polling time.

Conflict of Interests

The author declares that there is no conflict of interests regarding the publication of this paper.

References

[1] S. Stolfo, S. M. Bellovin, and D. Evans, "Measuring security," *IEEE Security & Privacy*, vol. 9, no. 3, pp. 60–65, 2011.

[2] H. Saiedian and G. Wishnie, "A complex event routing infrastructure for distributed systems," *Journal of Parallel and Distributed Computing*, vol. 72, no. 3, pp. 450–461, 2012.

[3] L. Kufel, "Security event monitoring in a distributed systems environment," *IEEE Security and Privacy*, vol. 11, no. 1, pp. 36–43, 2013.

[4] S. Lee, K. Levanti, and H. S. Kim, "Network monitoring: present and future," *Computer Networks*, vol. 65, pp. 94–98, 2014.

[5] G. Aceto, A. Botta, W. de Donato, and A. Pescapè, "Cloud monitoring: a survey," *Computer Networks*, vol. 57, no. 9, pp. 2093–2115, 2013.

[6] J. Montes, A. Sánchez, B. Memishi, M. S. Pérez, and G. Antoniu, "GMonE: a complete approach to cloud monitoring," *Future Generation Computer Systems*, vol. 29, no. 8, pp. 2026–2040, 2013.

[7] G. Katsaros, G. Kousiouris, S. V. Gogouvitis, D. Kyriazis, A. Menychtas, and T. Varvarigou, "A Self-adaptive hierarchical monitoring mechanism for Clouds," *Journal of Systems and Software*, vol. 85, no. 5, pp. 1029–1041, 2012.

[8] K. Fatema, V. C. Emeakaroha, P. D. Healy, J. P. Morrison, and T. Lynn, "A survey of cloud monitoring tools: taxonomy, capabilities and objectives," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2918–2933, 2014.

[9] S. Zanikolas and R. Sakellariou, "A taxonomy of grid monitoring systems," *Future Generation Computer Systems*, vol. 21, no. 1, pp. 163–188, 2005.

[10] S. Andreozzi, N. de Bortoli, S. Fantinel et al., "GridICE: a monitoring service for Grid systems," *Future Generation Computer Systems*, vol. 21, no. 4, pp. 559–571, 2005.

[11] M. L. Massie, B. N. Chun, and D. E. Culler, "The ganglia distributed monitoring system: design, implementation, and experience," *Parallel Computing*, vol. 30, no. 7, pp. 817–840, 2004.

[12] Nagios, http://www.nagios.org/.

[13] Zabbix, The Enterprise-class Monitoring Solution for Everyone, http://www.zabbix.com/.

[14] Ganglia, A scalable distributed monitoring system for high-performance computing systems, http://ganglia.sourceforge.net/.

[15] Hyperic, System Monitoring Software, http://hyperic-hq.sourceforge.net/.

[16] Submarine Cable Map, http://www.submarinecablemap.com/.

[17] L. De Vito, S. Rapuano, and L. Tomaciello, "One-way delay measurement: state of the art," *IEEE Transactions on Instrumentation and Measurement*, vol. 57, no. 12, pp. 2742–2750, 2008.

[18] P. Svoboda, M. Laner, J. Fabini, M. Rupp, and F. Ricciato, "Packet delay measurements in reactive IP networks," *IEEE Instrumentation & Measurement Magazine*, vol. 15, no. 6, pp. 36–44, 2012.

[19] Magic Quadrant for Global Network Service Providers, http://www.gartner.com/technology/reprints.do?id=1-1T96X35&ct=140417.

[20] K. Kent and M. Souppaya, *Guide to Computer Security Log Management*, US National Institute of Standards and Technology, 2006, http://csrc.nist.gov/publications/nistpubs/800-92/SP800-92.pdf.

[21] Verizon Network Latency, http://www.verizonenterprise.com/about/network/latency/.

[22] E. Casalicchio, M. Caselli, and A. Coletta, "Measuring the global domain name system," *IEEE Network*, vol. 27, no. 1, pp. 25–31, 2013.

[23] U. Sedlar, M. Volk, J. Sterle, A. Kos, and R. Sernec, "Contextualized monitoring and root cause discovery in IPTV systems using data visualization," *IEEE Network*, vol. 26, no. 6, pp. 40–46, 2012.