Harri Sillanpää

NETWORK PERFORMANCE MONITORING IN HYBRID INFRASTRUCTURE

Bachelor's thesis

South-Eastern Finland
University of Applied Sciences

South-Eastern Finland University of Applied Sciences

Degree title	Bachelor of Engineering
Author	Harri Sillanpää
Thesis title	Network performance monitoring in hybrid infrastructure
Commissioned by	Telia Cygate Oy
Year	2024
Pages	29 pages
Supervisor	Heikki Jantunen

ABSTRACT

Monitoring network performance in a modern hybrid environment that utilizes cloud computing resources together with on-premises infrastructure is crucial to provide quality of service for users and maintain continuous network operations for business-critical services. Monitoring performance in these environments possesses challenges due to the complexity of networks and with the large number of technologies involved.

The objective of this thesis was to study how to effectively monitor network performance in a hybrid infrastructure, what kind of monitoring solutions existed in the market, and to compare their characteristics, as well as to provide aspects of what organizations need to consider when deploying the solutions.

Qualitative methods were used to select the monitoring solutions for this study by interviewing networking and monitoring specialists from Telia Cygate Oy and researching the current market selection. Information about the monitoring technologies was gathered from the product documentation provided by the manufacturers.

The study showed that monitoring performance in a hybrid infrastructure was possible with the presented monitoring solutions. However, each of them is best suited for a specific environment and the decision of which solution might be suitable for an organization should be based on the current and future technologies in use within the network. Combining cloud providers native monitoring tools with an external monitoring solution might provide the widest overall picture of network performance in a hybrid environment.

Keywords: hybrid, performance, network, monitoring, cloud

CONTENTS

1 INTRODUCTION

Modern networking environments utilize cloud computing resources together with on-premises infrastructure, creating a hybrid networking environment for organizations. Monitoring network performance in these environments is crucial to provide quality of service for users and maintain continuous network operations for business-critical services. Hybrid infrastructure, as a term is not defined, and it is used vaguely in different sources to describe several hybrid environments. In this thesis the term describes an IT networking environment that combines on-premises networking infrastructure with one or more cloud providers' resources.

Performance monitoring in a hybrid infrastructure possesses challenges due to the complexity of the network combining several technologies, of which many have native monitoring solutions without the capability of monitoring connections across the technologies. This leads to the main questions that this thesis aims to answer, how do we effectively monitor network performance in a hybrid infrastructure, what kind of monitoring solutions are available in the market, and how well do they operate, and what aspects the organizations need to consider when deploying the monitoring solutions.

The topic for this thesis was presented by Heikki Jantunen who is working at Telia Cygate Oy as a senior cloud network architect, and the challenges of hybrid network monitoring are acknowledged within the company. My own interest in this topic comes from working at Telia Cygate Oy's IT Operation Center and troubleshooting connectivity issues related to a wide range of network technologies.

The scope of this study focuses on the performance monitoring of connections between networking devices within the hybrid infrastructure, consisting of on-premises data center and Azure public cloud environment. The theory sections give background information on network monitoring, public cloud operations, and data center network environments as well as introduce data collection protocols and metrics used in performance monitoring. Monitoring techniques and

challenges in a hybrid infrastructure are discussed and in the last section, the characteristics, and capabilities from different monitoring solutions available in the market are reviewed and compared. Product documentation and interviews with monitoring specialists from Telia Cygate is used for the source of information about the monitoring solutions.

2 NETWORK MONITORING

Network monitoring is the practice of consistently overseeing all the network equipment and connections for any failures to ensure continued network performance. A variety of data is collected from the devices and filtered for analysis to identify any issues in the networks. Network monitoring can be categorized into availability and performance monitoring while Application performance monitoring (APM) focuses more on application performance and user experience. (Gillis & Slattery n.d.)

2.1 Data collection

Simple network management protocol (SNMP) is the most widely used monitoring protocol for collecting information from devices and it has been in active use since 1988 with widespread compatibility. SNMP software runs on the hardware or service being monitored and collects data to a text based MIB database. A centralized SNMP manager software queries the agents MIB with Object Identifier (OID) requests to gather specific information. There are currently three versions of the SNMP protocol and in most cases it functions in a synchronous model, where communication is initiated by the SNMP manager and the agent responding. (Scarpati n.d.)

A more modern approach is streaming telemetry, which operates as a push-based model. Streaming telemetry supported network devices constantly send system metrics to the management system, whereas SNMP polls the devices at certain time intervals. A push-based model is generally more efficient and scalable in more extensive networks without affecting the devices' performance, and it allows a near real-time analysis of performance. A drawback to this

technology is that it is not widely supported yet, and some vendors use proprietary telemetry, requiring vendor's monitoring systems for analysis. (Gervasi 2024.)

Network flow monitoring provides more insight into traffic traversing through devices and networks. There are different standards and formats such as NetFlow, sFlow, and Internet Protocol Flow Information Export (IPFIX). Each of these work in a slightly different way, but all are distinct from port mirroring and packet capture which capture the contents of every packet that is passing. NetFlow is the original solution, developed by Cisco in the late 1990s, and it operates by capturing a set of packets from the flow of traffic that share a common set of characteristics such as source and destination address, port, and protocol type. The flow records are then exported to a flow collector and analyzed with a flow analyzer, these are often operating as a single entity and combined into a larger monitoring solution. IPFIX is very similar to NetFlow with few additional fields added and version nine of NetFlow was used as the basis for its development. It is an open standard and supported by many networking vendors apart from Cisco. sFlow, on other hand, captures deeper levels of information by randomly sampling the full packet headers and partial packet payloads. This can reduce CPU and bandwidth utilization on the collecting devices but may affect the accuracy of collected information. (Grimmick 2021.)

Another means of collecting data from the network devices is to utilize the system logging protocol (syslog), defined by Request for Comments (RFC) 3164 document. It allows devices to transfer event messages to a logging server, with defined numerical facility and severity values about the event. Facility value corresponds to the system of origin whereas severity values, ranging from zero to seven, describe the level of severity of an event, zero meaning an emergency where system is unusable (Lonvick 2001).

Monitoring software systems can utilize Application Programming Interface (API) protocols in a callback method, in which web-based messages are sent to other

systems after an event occurs. These are called webhooks and they use HTTP POST messages to trigger actions on the receiving end. (Heusser 2023.)

2.2 Availability monitoring

Availability monitoring is the basic layer of network monitoring, and it refers to the continuous monitoring of device hardware components and interfaces. The simplest monitoring method is a ping tool that sends Internet Control Message Protocol (ICMP) echo request messages to a target IP address and gets an echo reply, signaling a successful connection. SNMP protocol offers much more information from the target devices, and it is used to track vital system metrics such as CPU and memory, status of the power units and fans, and device uptime information. It also provides interface metrics which are vital for detecting any link failures in a redundant networking environment.

2.3 Performance monitoring

Network performance monitoring is the process of monitoring and measuring the quality of service of a network. Gathered metrics can aid in troubleshooting connectivity issues and help to optimize the network performance. Monitoring tools combine various types of network data such as metrics from infrastructure devices, flow data, packet data, and results from synthetic tests to analyze the performance over time.

Common performance metrics include bandwidth and utilization, which represents the maximum capacity of a network connection and the percentage of its usage at a given time, whereas throughput measures the actual amount of data that is or has been transferred. Packet loss, latency, and jitter metrics reveal insights on the performance of connections between a source and a destination. Latency measures delays, and it is often measured as a round trip time of packet travelling from source to destination and back. Packet loss refers to the number of packets that fail to reach the destination and jitter measures the inconsistency of arriving packet data or the variation in latency over time. (Kentik n.d.e.)

One of the most used tools for monitoring network paths is the traceroute. It was created in the 1980s and it uses a series of ICMP packets that are sent from source to destination IP address. Packets travel from router to router along the path with increasing Time to Live values (TTL), and each router reduces the value by one and replies to the source with an ICMP message TTL Exceeded when the value reaches zero. The first packet has a TTL value of one, so the first router replies, then a second packet is sent with TTL value of two and it is forwarded to the second router, process is repeated until the maximum number of hops is reached or the destination router replies. This operation reveals the routers along the path with round-trip times of each packet. However, latency values might not be accurate using ICMP since routers are prioritizing traffic and ICMP packets might be dropped. Firewalls may also drop ICMP packets along the path, but this can be avoided using TCP as the communication protocol. (Harris 2022.)

3 HYBRID INFRASTRUCTURE

Information technology (IT) requires infrastructure to operate, and IT infrastructure is simply the sum of all the necessary components, such as personal computers, servers, and networks that form the basis to run a business's IT services (Wright 2020). There are two primary types of IT infrastructures, traditional and cloud. In a traditional type, the company's infrastructure is located on their premises, and it is a private environment accessible only from the company network. (IBM n.d.a.)

A cloud computing IT infrastructure can be further divided into private, public, hybrid, and multi-cloud solutions. A private cloud is a single tenant environment in which all the hardware and software resources are dedicated exclusively to, and accessed only, by a single customer. The environment is typically hosted from an on-premises data center but can also be hosted on an independent cloud provider's infrastructure or built on rented infrastructure hosted from offsite data center by a service provider. A public cloud on the other hand is a multi-tenant environment where the computing resources are shared among multiple customers and the independent cloud service provider owns and maintains the

infrastructure. Access to those resources is available over the internet and they are provided on a subscription basis with pay-as-you-use pricing. A hybrid cloud combines private and public cloud environments whereas multi-cloud solution combines multiple independent cloud providers public cloud services. (IBM n.d.b.)

A hybrid infrastructure is a hybrid IT environment which works at the device, system and application levels in different locations and platforms. Part of the company's systems and applications can be in public cloud and some either in their own or in different vendors on-premises data centers. The hybrid environment consists of computation infrastructure that is located both on-premises and public cloud and it includes the overall architecture that governs the traditional IT systems and platforms. (Teräväinen 2022.)

3.1 Importance and benefits

Hybrid environment responds to the growing needs of organizations to make more use of data, enhance the flow of information, and enables agile development and scalability. (Teräväinen 2022.)

Benefits of opting for a hybrid infrastructure include scalability and security as well as ensuring business continuity and it can bring cost savings. By using public cloud services when demands spike, companies can seamlessly increase their operational capacity while having business critical data and operations running in their private cloud or an on-premises data center. Upgrading operational hardware and software is much faster in a cloud environment compared to traditional network infrastructure. (vmWare, n.d.)

Leveraging public cloud storage options can bring operational cost savings by having non-mission critical operations running in a public cloud instead of a private cloud or in on-premises data centers. Most of the public cloud storage services are charged only when they are used. In Hybrid infrastructure, public clouds can also be used to absorb workload surges to protect private servers

from overloading as well as duplicating data in case of disaster or computing failure ensuring business continuity. (vmWare, n.d.)

However, not all data can be stored in the cloud as some organizations may have strict policies about their data storage and access control, and some might be obligated by law to store their data within country limits at a certain location with secure access policies. The demand for enterprise and colocation owned data centers is increasing, based on Uptime Institute's Global Data Center Survey Results from 2023 (Pärssinen 2024).

3.2 Data centers and public cloud computing

"There's no cloud without a data center" (Sudipto 2023). A data center is a facility that consists of computing, storage, and networking infrastructure and a business relies heavily on the applications, services, and data located in the data center, making it a critical asset for all operations. There are many types of data centers and service models enterprises can choose from. Data centers can be smaller, completely owned and managed by an enterprise, or they can be colocation data centers by a third-party service provider, or they can be huge hyperscalers owned and managed by public cloud providers such as Amazon Web Services, Microsoft Azure, or Google Cloud. (Yasar 2022.)

According to Yasar (2023), public cloud providers deliver service models that can be categorized into three main categories: Infrastructure-as-a-service (IaaS), Platform-as-a-service (PaaS), and Software-as-a-service (SaaS). These models differ in the number of services and underlying infrastructure the cloud provider and the customer govern. This is called a shared responsibility model (Figure 1).

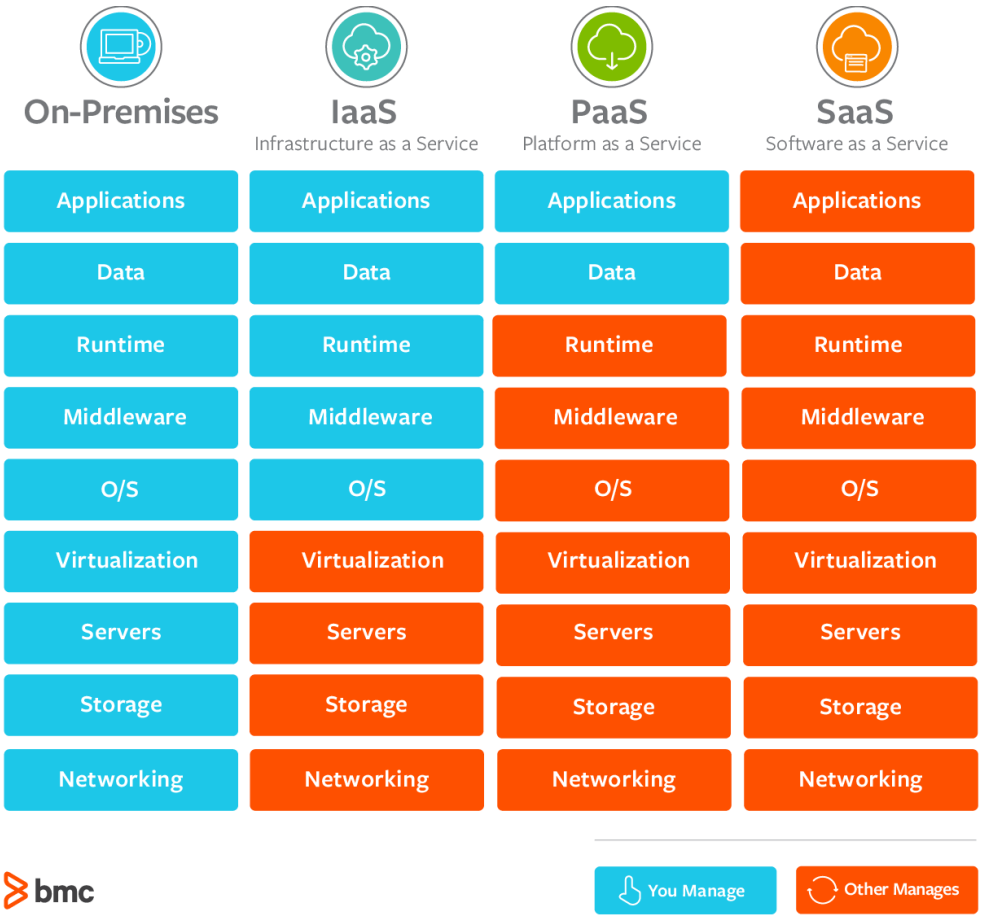

Figure 1. Shared responsibility model (Watts & Raza, 2019)

IaaS service model delivers compute, storage, and network resources on a pay-as-you-go basis. The customer is responsible for deploying, maintaining, and supporting the applications while IaaS provider is responsible for maintaining the physical infrastructure. (Aws, n.d.a.)

Public cloud providers' data centers are located around the world and clustered together in geographical locations called regions. Each region contains one or more availability zones (AZ), and each contains one or more data centers. AZs are close enough to have very low latency between them, but far enough to reduce the likelihood of local outages or weather affecting one or more at the same time. (Azure n.d.)

Resources in the cloud are deployed under private networks that are separate from those of other customers. Microsoft calls this Virtual Network (VNet) and Amazon Virtual Private Cloud (VPC), and both segregate the networks with

subnets. VPC subnets are mapped to an AZ and a subnet can belong to only one AZ and there can be public or private subnets. A subnet is public if it has an internet gateway attached to it allowing access to the internet from the deployed resources. VNet does not separate subnets to private and public and resources connected to it have access to the internet by default. (Ekezue 2017.)

3.3 Data center networking

At the core of a data center there are the servers, often called hosts and each consisting of one or more processors, memory, network interface and local high-speed disk or flash for storage. These resources are then packaged into racks and allocated as clusters that can consist of thousands of hosts all connected with a high-bandwidth network. (Abts & Felderman 2012.)

The data center network (DCN) interconnects all the resources, and it is like what the central nervous system is to a human (Abst & Felderman 2012). DCN consists of physical network devices such as switches, routers, load balancers, firewalls, and applications delivery controllers and it provides secure connections to the data center via internet or a network and supports virtualization (Sudipto 2023).

There are many architectural designs of data center networks and the traditional is a three-tier design model. It consists of three network layers: core, aggregate, and access. The core layer at the top connects the data center to the internet and routes the traffic. The aggregation layer below provides uplinks from the access layer switches and includes load balancers and firewalls. At the bottom layer, the access layer switches connect the servers to the network. The three-tier model is an ideal solution for north-to-south traffic but lacks fault tolerance, scalability and packets experience latency passing through multiple hops especially in server-to-server communication. (Sudipto 2023.)

A modern data center architectural design addresses scalability and redundancy limitations and increases network performance. In a Spine-leaf architecture there are two network layers: spine and leaf. The spine layer devices provide routing

and act as a core for the network while the leaf layer switches connect the servers and provide access to the network. In this design model, each leaf switch is connected to each spine switch, enabling fast server-to-server communication by reducing the number of hops between any server to two. (Margaret 2022.)

An illustration of the two architectures can be seen in Figure 2.

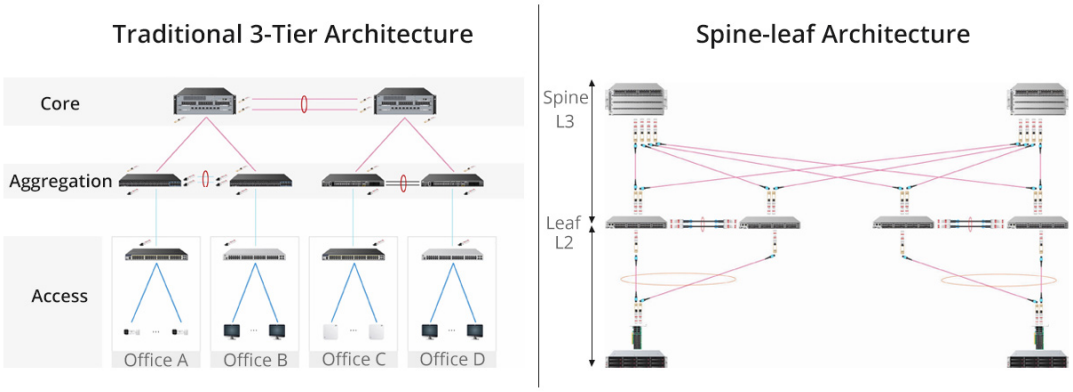

Figure 2. Three-tier and spine-leaf architectures (Margaret 2022)

In a hybrid infrastructure connecting to the data centers can be established by using a virtual private network (VPN), direct connection, or by software-defined wide area network (SD-WAN). A VPN is a popular option, and it enables organizations to connect securely using encrypted tunnels. A site-to-site VPN connects branch offices to companies' main network as well as to the cloud. Remote access VPN enables remote users to connect to the company network and access services and resources. A limitation of using VPN connections is that the traffic goes through the internet affecting the speed of connections and can cause congestion. (Vargas & Bielicki 2021.)

Users can connect to the applications residing in the data centers either by using mobile or fixed networks. Traffic goes through packet switched network and then through the internet core and enters the data centers. (Pärssinen 2024.)

Connecting on-premises or colocation data centers with cloud data centers can be established using direct physical connections on a private circuit, or virtual for example WAN cloud exchange (WAN-CX). Direct connections provide fast and

reliable communication between data centers avoiding the internet in between. Direct physical connections require that both the enterprise and cloud service provider (CSP) have their networking infrastructure in the same facility because network cable needs to be connected from enterprise router to CSP controlled router. These connections can be made in meet-me rooms provided by colocation data centers or by leasing ports from WAN provider. CSP's have different names for direct connection services e.g. Amazon Direct Connect and Microsoft ExpressRoute. (Burke 2018.)

VPCs can also be directly connected by using VPC peering. According to AWS, VPC peering connection is a connection between two VPCs, and traffic is routed between them using private IPv4 addresses and IPv6 addresses as if the instances were within the same network. Peering connections can be established between own account VPCs, or with a VPC in another account and they can be in different Regions. (Aws n.d.b.)

3.4 Performance monitoring techniques

Monitoring network performance in a hybrid infrastructure requires monitoring of several key components due to the scattered and complex nature of the networks. Connections between the core routers and switches near the servers in on-premises data center must be monitored as well as connections from the core routers to remote office networks and to cloud resources. An organization's network architectures can vary between the number of services hosted in the cloud or on-premises environments and whether users connect through office networks or use VPN connections from remote locations. These variances need to be considered when planning an effective performance monitoring strategy and technologies to be used. Monitoring can be focused on a user traffic or connectivity between network devices or a combination of both.

Performance monitoring techniques can be categorized into two main groups, active and passive. An active, also called synthetic monitoring, does not use real traffic as the source of data, instead it uses a series of network tests generated by monitoring probes or agents towards other probes or services. The tests can

simulate various types of network traffic and provide real-time insight into network performance. Active monitoring can be used for continuous testing and benchmarking network performance as well as pinpointing bottlenecks and troubleshooting connectivity issues. (Charest 2023.)

Passive monitoring, on the other hand, captures real network traffic that passes through the network devices and is stored for analysis without interfering with network performance. Capturing traffic can be done using hardware-based network taps that copy all the passing traffic or enabling port mirroring on the monitored devices or by using flow monitoring software. Passive monitoring allows administrators to examine the flow of traffic and enables deep insight into network utilization, performance, and security. It is used for identifying long-term trends and historical traffic analysis, as well as revealing top talkers in the networks and monitoring application usage. (Charest 2023.)

Path monitoring is also a form of performance monitoring, and it ensures that traffic is routed using optimal routes and can reveal possible hotspots and vulnerable bridge connections in the networks. Border gateway protocol (BGP) is used for routing traffic on the internet and it connects Autonomous system (AS) networks together on a global scale. Monitoring BGP routes provides valuable information on latencies between hops, detects changes in routing and ensures redundancy.

3.5 Performance monitoring challenges

One of biggest challenges when monitoring performance in a hybrid environment is the lack of visibility across technologies. Organizations may have a mix of monitoring technologies, one for on-premises and one or more for cloud environments, plus vendor specific management systems e.g., Aruba Central, Cisco Viptela and DNA Center which provide certain performance metrics.

When a cloud environment is added to an organization’s existing on-premises infrastructure or vice versa, it often requires an additional monitoring technology to get a complete picture of network performance in the hybrid environment.

Traditional monitoring technologies are often designed for on-premises environments and are not capable of monitoring cloud native services and resources. Cloud resources are highly dynamic and traditional monitoring tools may not adapt well to the on-demand scaling. (Lamberti 2023.)

A baseline of network performance must be established to set the correct thresholds for network tests because without it is difficult to troubleshoot connectivity issues and measure performance. Any added monitoring system must support integration with an organization's IT Service management (ITSM) system to enable ticketing and responses to alerts.

4 EXAMPLES OF MONITORING TECHNOLOGIES

Some of the popular network performance monitoring tools were chosen to research what is possible to monitor with these tools, how and where they can be deployed, and what kind of monitoring techniques they use. Microsoft Azure monitoring section showcases Azure public cloud monitoring tools, techniques, and possibilities in a hybrid environment. Possible disadvantages and advantages as well as additional thoughts of each monitoring technology are discussed in the comparison and thoughts section.

4.1 ScienceLogic

ScienceLogic SL1 platform offers a scalable and secure monitoring system that can be deployed on-premises, in customer managed clouds, or as ScienceLogic-managed SaaS. SL1 supports multiple separate client organizations in a single platform and data is fully partitioned to enable integrity. SL1 uses both agent and agentless data collection and serves availability, performance, events and configurations via SNMP, API, SSH, and Syslog to a unified operational data lake. Over 400 supported SL1 Powerpack monitoring solutions enable monitoring from e.g., major public cloud providers hybrid and multi-cloud environments, virtualization solutions (e.g., VMware), software defined networks, servers, and storage. SL1 supports integrations with several ITSM tools, providing automation to ticketing and troubleshooting as well as enabling bi-directional synchronization

of data with companies Configuration management database (CMDB). (ScienceLogic n.d.a.)

A distributed system of the SL1 platform consists of a user interface, a database server, and one or more message and data collectors. These operations can be run from an all-in-one appliance or from dedicated nodes each performing one function. The distributed system can be upgraded to an extended system, which brings four new appliances: a compute cluster, a load balancer, a storage cluster, and a management node. (ScienceLogic n.d.b.)

Data collectors are responsible for discovering new devices in the SL1 system and collecting agentless monitoring data from the devices. Message collectors collect SNMP, Trap messages, syslog messages, and agent-based data from installed managements agents. A database server is at the heart of the operation, and it pushes data to and from the collectors as well as processes and normalizes the data. It allocates tasks to the other nodes in the SL1 system, executes automation actions in response to events, and stores all configuration and policy data. The user interface can be accessed through a web browser and administrators as well as users can view collected data, reports, and events, define policies, and it provides access to the ScienceLogic API. (ScienceLogic n.d.b.)

From a standpoint of network monitoring, SL1 supports availability and latency monitoring, visibility to IP networks and system vital metrics such as CPU and physical memory. At the discovery phase, all network interfaces are detected from the devices and monitoring protocols are decided. Protocols used for availability monitoring of hardware-based devices are ICMP, SNMP, TCP, and UDP and for component devices, meaning an entity that runs under the control of a management device, Dynamic Applications are used from the data collectors to poll the devices for availability at a defined frequency. Latency monitoring refers to the amount of time it takes for an end device to respond and allow communication initiated by the SL1 and it is measured in milliseconds. SL1 can trigger events based on the monitoring data it receives by using set threshold

values. The values, for example for availability monitoring with ICMP, can be set to a percentage value from 0 to 100, 100 percentage meaning every ICMP packet must go through without triggering an event. (ScienceLogic n.d.c.)

SL1 extends its monitoring capabilities to a wide range of technologies with predefined monitoring packets, called PowerPacks. These include Dynamic Applications for data collection, event policies, device templates, custom reports, and dashboard views for system status monitoring, among others. All major public cloud providers are supported, and the PowerPack for Microsoft Azure enables monitoring for example to services such as ExpressRoute, Load balancer, Virtual network, and Domain name system (DNS). It is possible to install a collector into Azure Vnet and gather monitoring metrics with webhooks. (ScienceLogic n.d.d.).

A graphical representation of relations between the discovered devices in SL1 is presented in different topology Maps. There are Layer-2 maps, discovered with either Cisco discovery protocol (CDP) or with Link layer discovery protocol (LLDP), and Layer-3 maps. These are discovered by running a traceroute to the end devices from the collectors, and all devices which have layer-3 collection enabled, are visible in the map. (ScienceLogic n.d.b.)

4.2 Kentik

Kentik is a cloud-based platform for collecting, analyzing, and visualizing health and performance data of organization’s networks. Data can be collected from both on-premises infrastructure and cloud resources, utilizing actual traffic with flow metrics and synthetic tests. The Kentik portal is a Web-based user interface for management, traffic metrics analyzing, and viewing logical and network maps. (Kentik n.d.a.)

The main data source for Kentik metrics is the traffic flow that passes through the devices, such as routers, switches, and endpoints. The primary protocols in use are sFlow, IPFIX, and NetFlow versions nine and five. Flow metrics can be enriched with additional data with e.g., SNMP, BGP, and GeoIP to provide

interface names and descriptions, source, and destination details, as well as used protocols. Depending on the device, data can be collected or pushed to Kentik cloud storage from flow enabled devices, installed host agents or from locally hosted proxy agents. (Kentik n.d.a.)

Supported physical devices include Cisco's ASA, SD-WAN vEdge and Meraki, as well as Juniper PFE switches and Palo Alto firewalls, while host agent software runs only on Linux hosts (Kentik n.d.b.).

Kentik enables public cloud monitoring by exporting flow logs about resources in the cloud from the largest providers AWS, Azure, IBM, and Google Cloud Platform. The resources can be a VPC, a subnet, a virtual machine (VM), or an interface. In Azure, flow logs are generated by network security groups (NSGs) and consist of a set of records about the flow of traffic that originated from or destined to a resource. For AWS, Kentik offers a more detailed network monitoring called Cloud Performance Monitor with a user interface and it enables Direct connect and Site-to-Site VPN service monitoring as well as guidance on where to place synthetic agents for network testing. (Kentik n.d.c.)

Various synthetic network tests can be run with Kentik's software agents that are either private agents, deployed on-premises or on cloud infrastructure, or global agents maintained by Kentik and deployed around the world in key Internet hub locations. Network tests can be agent-to-agent, agent-to-server, or autonomous tests based on gathered flow logs. Ping and traceroute (TCP or UDP) tests provide latency, jitter, and loss metrics while application and routing performance tests focus on HTTP, DNS, and BGP. (Kentik n.d.d.)

4.3 Obkio

Obkio is a cloud-based SaaS application for synthetic monitoring, and it utilizes monitoring agents for either APM or NPM. There are both hardware and software monitoring agents available which gather latency, jitter, and packet loss metrics as well as measure Voice over IP (VoIP) quality and perform on-demand or scheduled network tests. Agents can be installed to end-point devices, servers,

or close to the perimeter of the network to monitor Wide area network (WAN), or they can be preconfigured and installed public monitoring agents deployed in public cloud providers locations. Monitoring agents can monitor other network devices with SNMP to provide metrics such as CPU usage, interface bandwidth, error, and availability metrics. Traceroute and speed tests can be run between agents and Obkio provides various visualization graphs of these results and from gathered device metrics. (Obkio n.d.a.)

There are currently eight public monitoring agents in Microsoft Azure located in four countries, hosted, and maintained by Obkio (Obkio n.d.b.). These public monitoring agents can monitor cloud deployed applications, services, and provide network metrics such as response time, throughput, latency, bandwidth utilization, and packet loss. It is possible to set threshold values and get alert notifications from triggered events through SMS, email, or by integration with an incident management system. (Obkio n.d.c.)

4.4 ThousandEyes

ThousandEyes is part of Cisco networking company, and it offers an agent based NPM and APM software delivered as a SaaS with synthetic monitoring strategy. It is possible to monitor BGP routing, hybrid, and multi-cloud environments of major cloud providers as well as on-premises key networking devices and end points. Internet service providers (ISP) broadband routers can also be monitored with ISP monitoring tool which is currently available in the USA with selected providers (ThousandEyes n.d.a.). Routing and network traffic path visualization can be observed with hop-by-hop graphical maps of both public internet and enterprises WAN.

Cloud environment monitoring is enabled by leveraging the global network of pre-deployed Cloud agents maintained by ThousandEyes. Microsoft Azure being the most covered with agents installed in 36 regions vantage points to provide loss, latency, and jitter metrics as well as service delivery paths. ThousandEyes is compliant with OpenTelemetry, Sharelinks and Terraform to assist with APM and triggering alerts from tests. (ThousandEyes n.d.b.)

End point and network device monitoring is enabled by installing Endpoint and Enterprise monitoring agents to end-user laptops and key network devices in the network. Enterprise agents can also be installed into virtual cloud environment VPCs and Vnets and supported deployment options are presented in Figure 3.

Supported Hardware Devices and Models		
Cisco 4000 Series Integrated Service Routers (ISRs)*	Cisco Catalyst 8300 and 8200 Series Edge Platforms*	Cisco Catalyst 9300 and 9400 Series Switches*†
Cisco Meraki MX Series (MX67/C/W, MX68/C/CW, MX75, MX85, MX95, MX105, MX250, MX450)‡	Cisco Nexus 9500 and 9300 Series Switches*	Raspberry Pi 4 Model B
Software and Virtual Environment Deployment Options		
Container for Docker	Installer for Intel® NUC	Linux Package (RHEL, CentOS, Ubuntu)
VM for ESXi, Oracle VirtualBox, Hyper-V		

Figure 3. Enterprise agent deployment options (ThousandEyes n.d.)

Enterprise agents auto-discover network infrastructure with LLDP and CDP protocols and poll the target devices with SNMP to gather specified telemetry data. It is also possible to run various network tests from any endpoint or enterprise agent to other agents as well as to service IPs and analyze the results. (ThousandEyes n.d.c.)

4.5 Microsoft Azure native monitoring tools

Public cloud providers underlying network infrastructure are completely owned and managed by the provider, therefore there is no visibility and access into the actual routers and switches. Network monitoring in public cloud environments is enabled through various services.

Azure Network Watcher provides tools to monitor, diagnose, and view metrics collected from Azure IaaS resources such as VMs, VNets, applications gateways, and load balancers. Diagnostic tools enable administrators to test next hop routes, verify IP flow, enable packet capture on VMs, and troubleshoot VPN connections. IP traffic to and from NSGs is monitored with flow logs which provide information on traffic levels, bandwidth, and identifies top talkers and undesired traffic in the networks. VNet flow logs tool is currently in a preview stage and will simplify log monitoring in the future by enabling log monitoring at

virtual networks. Traffic analytics tool is used to visualize network activity across the Azure subscriptions, help to optimize network deployments based on traffic flow patterns, as well as show open ports to the internet and pinpoint VMs that attempt to connect to rogue networks. (Kazwini 2023.)

Azure Monitor is responsible for collecting, analyzing, and responding to monitoring data from cloud and on-premises environments (Figure 4). It is possible to collect data from applications, virtual machines, guest operating systems, databases, containers, and network events in combination with Network Watcher. (Azure monitor overview 2024.)

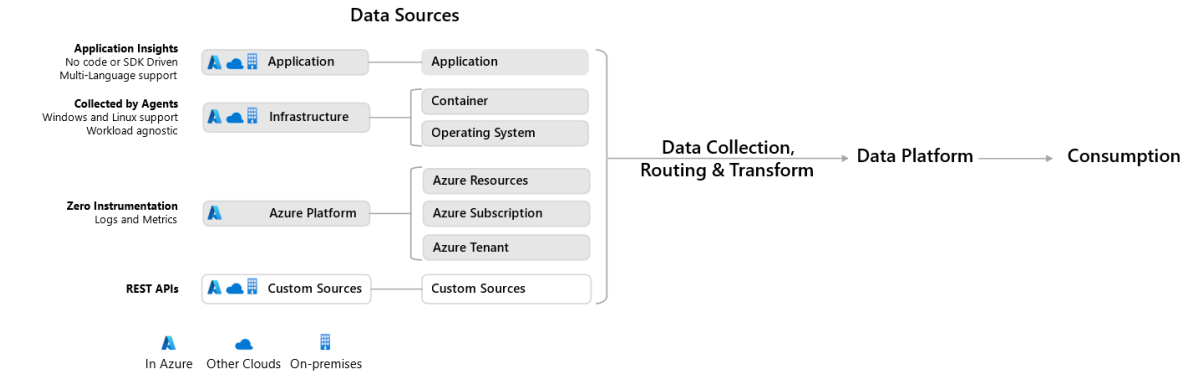

Figure 4. Azure Monitor data sources (Microsoft 2024)

Extending monitoring to on-premises and other clouds resources is made possible by enabling Microsoft Arc on target machines and installing Azure Monitor Agents to containers, virtual machines, or host operating systems. Supported systems include Windows and Linux virtual machines, workstations, and laptops. Agents collect syslog, performance metrics, and file-based logs from the host systems. (Azure Monitor Agent overview 2024.)

With Azure Monitor, it is possible to monitor ExpressRoute and Direct ExpressRoute, collect performance, circuit, and route metrics as well as ARP and BGP availability. Performance metrics include throughput, dropped bits and for the Direct connection, physical metrics such as Rx and Tx light levels, admin state, and line protocol status of ports are added. (ExpressRoute monitoring, metrics, and alerts 2023.)

Connection Monitor service enables various connectivity tests between monitoring agents towards endpoints by using HTTP, ICMP, and TCP protocols. Tests can verify connectivity and availability of hosts as well as continuously monitor, for example ExpressRoute connections. Connection Monitor can send triggered alerts to Azure Monitor, which then can automatically trigger scaling on resources or send alerts to integrated ticketing systems with webhooks. (Connection monitor overview 2024.)

The results of Connection Monitor tests, network topology, health, and collected metrics can be viewed with Azure Monitor Network Insight service. It provides a visual presentation of the above as well as an access to NSG flow logs, Traffic analytics and diagnostic tools. (Azure Network Insight overview 2023.)

4.6 Ixia Hawkeye

Hawkeye is a synthetic network and application monitoring software, and it is a part of Keysight's product line (Keysight n.d.). Hawkeye is a single tenant environment, and the platform consists of a main server with registration and license servers attached, and a web-based user interface. The main server can be deployed on-premises, and it collects the monitoring results and metrics. Network and application monitoring is enabled with Ixia's hardware probes and Hawkeye software agents distributed in the network or installed to the end devices. Network tests are configured in the main server and sent to the probes or agents which generate reports of the results. One to one test generates synthetic traffic from node to node, while continuous tests are real service tests where endpoints generate continuous traffic to server or network equipment and reports the response times. Hawkeye uses ICMP, TCP, UDP, and HTTP as well as real user traffic with video or voice type packets to produce performance metrics such as delay, loss, jitter, bandwidth, Mean Option Score (MOS), and throughput. Path discovery tests provide visibility into network topology from Hawkeye hardware end points to any remote location and they are run with TCP and UDP or ICMP packets. Tests are configurable to show alternative routes to

the destination, and results are presented as a graphical hop-by-hop map with added Autonomous System (AS) numbers of middle points. (Hawkeye n.d.)

Hawkeye endpoints can be virtual-, hardware-, or software-based and their monitoring capabilities vary depending on the model (Figure 5). Endpoints are easy to deploy and require only IP configurations and necessary firewall openings for the monitoring connections.

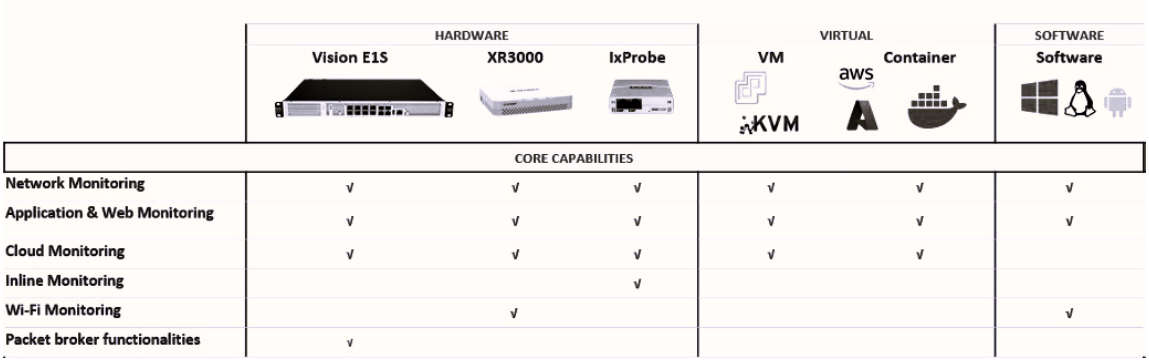

	HARDWARE			VIRTUAL		SOFTWARE
	Vision E1S	XR3000	IxProbe	VM	Container	Software
CORE CAPABILITIES						
Network Monitoring	√	√	√	√	√	√
Application & Web Monitoring	√	√	√	√	√	√
Cloud Monitoring	√	√	√	√	√	
Inline Monitoring			√			
Wi-Fi Monitoring		√				√
Packet broker functionalities	√					

Figure 5. Hawkeye endpoints (Keysight n.d.)

All endpoint types can perform basic network monitoring tests with TCP, UDP, ICMP, and web page downloads but hardware endpoints have the most tests available e.g., path discovery and additional real service tests. Hardware endpoints are designed to be inserted at key junctures in the network whereas virtual and software endpoints are for the servers and user workstations. Cloud monitoring in Figure 5 refers to a TCP ping test to any known AWS server with Path discovery. However, it is possible to do the basic network tests with a software or container endpoint located in AWS and Azure.

Keysight offers additional monitoring services including hardware and virtual network taps, packet brokers as well as Cloudlens service for packet capture in the cloud. These provide 100% visibility to link traffic and enable east-to-west traffic monitoring without impacting network performance. IxChariot is a lighter weight cloud-based monitoring solution from Keysight, delivered as a SaaS from AWS platform. All its endpoints are software-based, and they can be installed on any Windows or Linux PC/server, Android or IPhone as well as on VMs and

Internet of things (IoT) devices. IxChariot produced tests can be viewed from management console in the cloud or from Hawkeye main server.

4.7 Comparison of the solutions and additional thoughts

Sciencelogic SL1 monitoring platform is an excellent tool for monitoring availability, vital system metrics, and processes from network devices and servers. It is a highly scalable multitenant environment capable of monitoring thousands of devices with the ability to separate the devices by organizations and it can be installed on-premises. PowerPacks extend the monitoring capabilities to other environments, and it is possible to get limited number of metrics from public cloud providers. Performance monitoring is limited to latency metrics between a collector and an endpoint, and synthetic performance tests as well as flow monitoring is not supported by the SL1.

Kentik is a cloud-based solution, and the main data source is the flow metrics gathered from software agents and flow enabled network devices. Wide range of routers, switches, and firewalls are supported while host agents can only be installed to Linux hosts. Largest public cloud providers are well covered with public cloud agents and depending on the provider, different flow metrics can be gathered. AWS is the most covered with the additional monitoring of the DirectConnect and Site-to-Site VPNs. Kentik is a great network performance monitoring tool for monitoring connections between organizations key network devices and locations with flow metrics and synthetic tests. It is best suited for a hybrid infrastructure with AWS as the cloud provider and supported devices in use at the on-premises data center and remote offices.

Obkio and all the other cloud-based solutions are easy to set up and low in maintenance for administrators but require tight data retention policies to avoid extra costs for data storage. Obkio’s wide variety of monitoring agents can be installed to almost any host, being a network device or a server, and agents can monitor surrounding devices as well. Supported synthetic tests deliver basic performance metrics between the agents but Obkio does not support flow

metrics. Public monitoring agents are distributed among the largest providers; however, the number of agents is currently quite low.

ThousandEyes has a limited number of supported network devices for installing Enterprise monitoring agents, only Cisco routers and switches are supported. Endpoint agents can be installed to Windows and Mac operating systems while Linux systems are supported for server and docker agents. Public cloud monitoring agents are well distributed around the world and the most regions are covered in Azure. On top of the basic network performance metrics, it is possible to view graphical network path and topology maps from agent tests as well as monitor BGP routing. ThousandEyes is most suited for hybrid environment that combines on-premises Cisco's key network devices with one or more cloud providers.

Azure's native monitoring tools offer the most metrics and visibility to their public cloud environment and services. There are many tools for network monitoring, and it takes a skillful administrator to use those effectively and optimally to avoid unnecessary costs. On-premises and other cloud environments can be monitored with Microsoft Arc enabled on target servers and machines. Supported operating systems include Linux and Windows versions. Available network performance metrics are very limited from target servers, and on-premises key network devices cannot be monitored, except ExpressRoute with the Azure Monitor tool.

Hawkeye is a synthetic monitoring platform, and it can be installed to and managed from an on-premises environment. It offers a wide variety of network tests that can be run from node-to-node or from node-to-service with hardware and software-based monitoring probes that can be installed on most host systems. Synthetic tests are the most versatile with the added real user traffic simulations and MOS, on top of the network performance metrics. However, Hawkeye is a single tenant environment which means that the network tests and probes cannot be grouped by organizations. Hawkeye's Path discovery test includes additional routes as well as AS numbers of middle points by using TCP, UDP, and ICMP packets instead of just ICMP.

Public cloud monitoring is possible by installing Hawkeye software agents to VMs, containers, or workstations. However, AWS's cloud monitoring is supported with the additional capability to run network tests to any known AWS server from a probe without the need for installing an agent on the target server. Cloud environment and on-premises infrastructure monitoring can be further enhanced with the other monitoring products from Keysight. Hawkeye is well suited for monitoring network performance in on-premises data centers and connections to remote offices and cloud servers. High use of hardware probes on the other hand requires personnel for installation and maintenance.

5 CONCLUSION

In this thesis the main questions were how to effectively monitor network performance in a hybrid infrastructure, what kind of monitoring solution were available in the market, and how well do they operate, and what aspects the organizations need to consider when deploying the monitoring solutions.

Based on the research in this thesis, each monitoring solution is best suited for a certain type of environment and might offer a wider range of monitoring options for specific cloud providers. The decision of which solution might be suitable is related to the organization's current and future environment and which kind of devices and traffic needs to be monitored. It might be needed to combine Azure's native monitoring tools with a performance monitoring solution to effectively monitor performance in a hybrid infrastructure. Monitoring solutions store information about the networking devices and form topology maps of the environments, and this aspect must be considered when deciding a deployment option as on-premises based solutions offer more secure access and control over the data.

The selected monitoring products in this thesis represent a part of available solutions in the market and the information about each monitoring solution was gathered from the product documentations which might not include all the details about the software and how well it operates in a certain environment. Testing the

solutions in a real production environment would give more insight into the software and its features.

In the beginning of this work, I had a plan to implement and test the best suited solution in a testing environment. However, after the research it was hard to decide the most suitable one since the testing environment and cases were not specified clearly. For future testing and research, I would suggest testing the different solutions in a testing environment with a specific scenario. It could be troubleshooting a connectivity issue or measuring connection performance e.g., between on-premises data center and cloud provider. These tools could also be used for testing or measuring network performance after configuration or hardware changes.

REFERENCES

Abst, D., Felderman, B. 2012. A Guided Tour through Data-center Networking: A good user experience depends on predictable performance within the data-center network. *Association for Computing Machinery,* Queue 10 (5), 10—23. E-magazine. Available at: https://dl.acm.org/doi/abs/10.1145/2208917.2208919 [Accessed 21 January 2024].

Aws, n.d.a. What is IaaS (Infrastructure as a Service)? Web page. Available at: https://aws.amazon.com/what-is/iaas/ [Accessed 7 January 2024].

Aws, n.d.b. What is VPC peering? Web page. Available at: https://docs.aws.amazon.com/vpc/latest/peering/what-is-vpc-peering.html [Accessed 27 January 2024].

Azure Monitor Agent overview. 2024. Microsoft learn platform. Web page. Available at: https://learn.microsoft.com/en-us/azure/azure-monitor/agents/agents-overview [Accessed 10 March 2024].

Azure Monitor overview. 2024. Microsoft learn platform. Web page. Available at: https://learn.microsoft.com/en-us/azure/azure-monitor/overview [Accessed 10 March 2024].

Azure Network Insight overview. 2023. Microsoft learn platform. Web page. Available at: https://learn.microsoft.com/en-us/azure/network-watcher/network-insights-overview [Accessed 10March 2024].

Azure. n.d. Build solutions for high availability using availability zones. Web page. Available at: https://learn.microsoft.com/en-us/azure/architecture/high-availability/building-solutions-for-high-availability [Accessed 28 January 2024].

Burke, J. 2018. What cloud data center interconnect technologies can do. *Buyer's Handbook: Now data center interconnect lets traffic ride the cloud.* Web article. Available at: https://www.techtarget.com/searchnetworking/feature/What-cloud-data-center-interconnect-technologies-can-do [Accessed 27 January 2024].

Charest, F. 2023. Active vs. Passive Network Monitoring: Which Method is Right for You. Blog. 13 July 2023. Available at: https://obkio.com/blog/active-vs-passive-network-monitoring/ [Accessed 29 March 2024].

Connection monitor overview. 2024. Microsoft learn platform. Web page. Available at: https://learn.microsoft.com/en-us/azure/network-watcher/connection-monitor-overview [Accessed 10 March 2024].

Ekezue, F. 2017. Differentiating between Azure Virtual Network (VNet) and AWS Virtual Private Cloud (VPC). Blog. Available at: https://devblogs.microsoft.com/premier-developer/differentiating-between-azure-virtual-network-vnet-and-aws-virtual-private-cloud-vpc/ [Accessed 28 January 2024].

ExpressRoute monitoring, metrics, and alerts. 2024. Microsoft learn platform. Web page. Available at: https://learn.microsoft.com/en-us/azure/expressroute/expressroute-monitoring-metrics-alerts [Accessed 10 March 2024].

Gervasi, P. 2024. The Benefits and Drawbacks of SNMP and Streaming Telemetry. Blog. 10 January 2024. Available at: https://www.kentik.com/blog/the-benefits-and-drawbacks-of-snmp-and-streaming-telemetry/ [Accessed 30 March 2024].

Gillis, A. Slattery, T. n.d. Network monitoring. Web page. Available at: https://www.techtarget.com/searchnetworking/definition/network-monitoring [Accessed 30 March 2024].

Grimmick, R. 2021. Network Flow Monitoring Explained: NetFlow vs sFlow vs IPFIX. Blog. 17 June 2023. Available at: https://www.varonis.com/blog/flow-monitoring [Accessed 31 March 2024].

Harris, M. 2022. Traceroute Limitations Explained. Blog. 3 January 2022. Available at: https://www.netbraintech.com/blog/limitations-of-traceroute/ [Accessed 1 April 2024].

Hawkeye. n.d. User guide. Web page. Available only to product users. [Accessed 16 March 2024].

Heusser, M. 2023. Webhooks explained simply: What they do and how they work. Web page. Available at: https://www.techtarget.com/searchapparchitecture/tip/Webhooks-explained-simply-and-how-they-differ-from-an-API [Accessed 31 March 2024].

IBM. n.d.a. What is IT Infrastructure? Web page. Available at: https://www.ibm.com/topics/infrastructure [Accessed 3 February 2024].

IBM. n.d.b. What is private cloud? Web page. Available at: https://www.ibm.com/topics/private-cloud [Accessed 3 February 2024].

Keysight. n.d. Hawkeye Software Suite. Web page. Available at: https://www.keysight.com/us/en/products/network-visibility/hawkeye-software-suite.html [Accessed 16 March 2024].

Kazwini, H. 2023. What is Azure Network Watcher? Web page. Available at: https://learn.microsoft.com/en-us/azure/network-watcher/network-watcher-overview [Accessed 4 March 2024].

Kentik. n.d.a. About Kentik. Web page. Available at: https://kb.kentik.com/v0/Ab01.htm [Accessed 3 March 2024].

Kentik. n.d.b. Network Devices. Web page. Available at: https://kb.kentik.com/v4/Cb01.htm#Cb01-Supported_Device_Types [Accessed 3 March 2024].

Kentik. n.d.c. Cloud Overview. Web page. Available at: https://kb.kentik.com/v4/Na00.htm [Accessed 3 March 2024].

Kentik. n.d.d. Synthetics Overview. Web page. Available at: https://kb.kentik.com/v4/Ma00.htm [Accessed 3 March 2024].

Kentik. n.d.e. Network Performance Monitoring (NPM). Web page. Available at: https://www.kentik.com/kentipedia/network-performance-monitoring/ [Accessed 31 March 2024].

Lamberti, A. 2023. How to Monitor Hybrid Networks for End-to-End Visibility: Hybrid Network Monitoring. Blog. 5 October 2023. Available at: https://obkio.com/blog/hybrid-network-monitoring/ [Accessed 29 March 2024].

Lonvick, C. 2001. The BSD syslog Protocol. RFC3164 part 4.1. Available at: https://www.rfc-editor.org/rfc/rfc3164.html#section-4.1 [Accessed 31 March 2024].

Margaret. 2022. What Is Spine-leaf Architecture and How to Design It. Blog. Available at: https://community.fs.com/article/leaf-spine-with-fs-com-switches.html [Accessed 22 January 2024].

Obkio. n.d.a. Obkio Documentation Center. Web page. Available at: https://obkio.com/docs/ [Accessed 21 February 2024].

Obkio. n.d.b. Azure. Web page. Available at: https://obkio.com/public-monitoring-agents-directory/azure/ [Accessed 21 February 2024].

Obkio. n.d.c. A Deep Dive into Microsoft Cloud Monitoring for IT Pros. Web page. Available at: https://obkio.com/blog/microsoft-cloud-monitoring/ [Accessed 21 February 2024].

Pärssinen, M. 2024. Elements of sustainable ICT, Data Center. Lecture recording. Aalto University. 16 January 2024.

Scarpati, J. n.d. Simple Network Management Protocol (SNMP). Web page. Available at: https://www.techtarget.com/searchnetworking/definition/SNMP [Accessed 30 March 2024].

ScienceLogic. n.d.a. SL1 Platform Overview. Web page. Available at: https://sciencelogic.com/platform/overview [Accessed 12 February 2024].

ScienceLogic. n.d.b. Overview of SL1 Features. Product documentation. Available at: https://docs.sciencelogic.com/latest/Content/Web_General_Information/Overview_SL1/chapter_02_feature_overview.htm [Accessed 12 February 2024].

ScienceLogic. n.d.c. Monitoring Device Availability and Latency. Product documentation. Available at: https://docs.sciencelogic.com/latest/Content/Web_Monitoring_Tools/Infrastructure_Health/monitoring_availability_and_latency.htm [Accessed 13 February 2024].

ScienceLogic. n.d.d. Key Metrics Collected by the PowerPack. Product documentation. Available at: https://docs.sciencelogic.com/latest/Content/Web_Vendor_Specific_Monitoring/Microsoft_Azure/azure_key_metrics.htm#Azureh1_27 [Accessed 14 February 2024].

Sudipto, P. 2023. Data Center Networking: What It Is, Why It Matters, And Types. Web page. Available at: https://www.g2.com/articles/data-center-networking-dcn [Accessed 21 January 2023].

Teräväinen, T. 2022. Hybridi-IT puhuttaa – näistä syistä se voi olla organisaationne seuraava askel. Blog. 16 March 2022. Available at: https://www.cgi.com/fi/fi/blogi/pilvipalvelut-ja-hybridi-it/hybridi-it-puhuttaa [Accessed 30 December 2023].

ThousandEyes. n.d.a. Cloud Agents. Web page. Available at: https://www.thousandeyes.com/product/cloud-agents [Accessed 17 February 2024].

ThousandEyes. n.d.b. Microsoft Azure Monitoring. Web page. Available at: https://www.thousandeyes.com/solutions/azure-monitoring [Accessed 17 February 2024].

ThousandEyes. n.d.c. Visualize the Impact of Device Health on User Experience. Web page. Available at: https://www.thousandeyes.com/solutions/network-device-monitoring [Accessed 17 February 2024].

Vargas, D., Bielicki, B. 2021. A guide to networking for a hybrid infrastructure. *Next-Generation Infrastructure*. Web article. Available at: https://blog.shi.com/next-generation-infrastructure/a-guide-to-networking-for-a-hybrid-infrastructure/ [Accessed 27 January 2024].

vmWare, n.d. What is hybrid infrastructure?. Web page. Available at: https://www.vmware.com/nordics/topics/glossary/content/hybrid-infrastructure-service.html [Accessed 6 January 2024].

Wright, S. 2020. A Quick Guide to Enterprise IT Infrastructure. Web page. Available at: https://www.wrighttechnologies.com/guide-to-enterprise-it-infrastructure/ [Accessed 3 February 2024].

Yasar, K. 2022. Definition data center. Web page. Updated April 2022. Available at: https://www.techtarget.com/searchdatacenter/definition/data-center# [Accessed 7 January 2024].

Yasar, K. 2023. Cloud computing. Web page. Updated December 2023. Available at: https://www.techtarget.com/searchcloudcomputing/definition/cloud-computing [Accessed 7 January 2024].

Figure 1. Shared responsibility model. Watts, S., Raza, M. 2019. Available at: https://s7280.pcdn.co/wp-content/uploads/2017/09/saas-vs-paas-vs-iaas.png [Accessed 7 January 2024].

Figure 2. Three-tier and spine-leaf architectures. Margaret. 2022. Available at: https://community.fs.com/article/leaf-spine-with-fs-com-switches.html [Accessed 22 January 2024].

Figure 3. Enterprise agent deployment options. ThousandEyes. n.d. Available at: https://marketo-web.thousandeyes.com/rs/772-KGG-249/images/ThousandEyes-Product-Brief-Enterprise-Agent.pdf [Accessed 17 February 2024].

Figure 4. Azure Monitor data sources. Microsoft. 2024. Available at: https://learn.microsoft.com/en-us/azure/azure-monitor/overview [Accessed 10 March 2024].

Figure 5. Hawkeye endpoints. Keysight. n.d. Available at: https://www.keysight.com/us/en/assets/7019-0137/data-sheets/Hawkeye-Active-Network-Monitoring-Platform.pdf [Accessed 17 March 20204].