

INTERDISCIPLINARY PROJECT REPORT

at

**Sathyabama Institute of Science and Technology
(Deemed to be University)**

Submitted in partial fulfilment of the requirements for the award of Bachelor
of Engineering Degree in Information Technology

By

Keerthinadh Tulluri

(Reg.No.40120110)



DEPARTMENT OF INFORMATION TECHNOLOGY

SCHOOL OF COMPUTING

**SATHYABAMA INSTITUTE OF SCIENCE AND TECHNOLOGY
JEPPIAAR NAGAR, RAJIV GANDHI SALAI,
CHENNAI – 600119, TAMILNADU**

APRIL 2023



SATHYABAMA

INSTITUTE OF SCIENCE AND TECHNOLOGY
(DEEMED TO BE UNIVERSITY)

Accredited with Grade "A" by NAAC
(Established under Section 3 of UGC Act, 1956)



JEPPIAAR NAGAR, RAJIV GANDHI SALAI, CHENNAI- 600119

DEPARTMENT OF INFORMATION TECHNOLOGY

BONAFIDE CERTIFICATE

This is to certify that this Project Report is the bonafide work of **Keerthinadh Tulluri** who carried out the project entitled "**Cloud Machine Networking Configurations**" under my supervision from Jan 2023 to Apr 2023.

Internal Guide

Dr. Senduru Srinivasulu, M.E., Ph.D.

Head of the Department

Dr. R. Subhashini, M.E., Ph.D.

Submitted for Viva voce Examination held on

29/04/23

Internal Examiner

External Examiner

DECLARATION

I, Keerthinadh Tulluri hereby declare that the Project Report entitled "**Cloud Machine Networking Configurations**" Done by me under the guidance of **Dr. Senduru Srinivasulu, M.Tech., Ph.D.** is submitted in partial fulfilment of the requirements for the award of Bachelor of Engineering degree in Information Technology.

DATE: 29/04/2023

PLACE: Chennai

T. Keerthi

SIGNATURE OF THE CANDIDATE

ACKNOWLEDGEMENT

I am pleased to acknowledge my sincere thanks to **Board of Management** of **SATHYABAMA** for their kind encouragement in doing this project and for completing it successfully. I am grateful to them.

I convey my thanks to **Dr. T. Sasikala M.E., Ph.D., Dean**, School of Computing , **Dr. R. Subhashini, M.E., Ph.D.**, Head of the Department of Information Technology for providing me necessary support and details at the right time during the progressive reviews.

I would like to express my sincere and deep sense of gratitude to my Project Guide **Dr.Senduru. Srinivasulu, M.Tech.,Ph.D.** for his valuable guidance, suggestions and constant encouragement paved way for the successful completion of my project work.

I wish to express my thanks to all Teaching and Non-teaching staff members of the **Department of Information Technology** who were helpful in many ways for the completion of the project.

TRAINING CERTIFICATE



इलेक्ट्रॉनिक्स एवं
सूचना प्रौद्योगिकी मंत्रालय
MINISTRY OF
ELECTRONICS AND
INFORMATION TECHNOLOGY
सत्यमेव जयते

Ref.No.: 32 /FSPCC/BC001/04/2023



COURSE COMPLETION LETTER

It is certified that

Keerthinadh T

Sathyabama Institute of Science and Technology

has successfully undergone the Bridge Course Training Program on

Cloud Computing Technology

Under "FutureSkills PRIME" Programme

commenced from 11 Mar, 2023

Shri. V. Sridharan

Chief Investigator

Cloud Computing Technology

Lead Resource Centre

L R PRAKASH

Digitally signed by L R PRAKASH
Date: 2023.04.26 02:04:45
+05'30'

Cdr.(Retd.) L.R. Prakash

Senior Director

Centre Head

C-DAC, Chennai

ABSTRACT

This project is focused on the creation and configuration of an Amazon Web Services (AWS) Elastic Compute Cloud (EC2) instance. EC2 provides resizable compute capacity in the cloud, allowing users to launch virtual machines with various operating systems and applications. This project specifically focuses on creating an EC2 instance of Ubuntu and configuring its security rules to accept incoming HTTP and SSH traffic while blocking ping requests from the public network.

The project begins by outlining the objectives, which include creating an EC2 instance, configuring the security group, generating a key pair for SSH access, and testing the instance. The scope of the project is well-defined and provides a clear understanding of what to expect from each step. The project also highlights the importance of creating a secure EC2 instance, which is essential for any workload or application that will be hosted in the cloud.

The project provides a 10-step process for creating and configuring the EC2 instance, with each step described in a single-line method. The steps include creating an AWS account, launching an EC2 instance, configuring security groups, generating a key pair, and testing the instance. The instructions are detailed and easy to follow, making it an ideal resource for beginners who are new to AWS EC2.

The project also provides an explanation of HTTP clients and requests, which are essential components in web application development for communicating with web servers using the HTTP protocol. The requests library is a popular Python library that provides an easy-to-use interface for sending HTTP requests and handling HTTP responses. By understanding these components, developers can create more robust and efficient web applications that can communicate with web servers securely and reliably.

Overall, this project serves as an excellent introduction to creating and configuring an AWS EC2 instance, providing valuable insights into cloud computing, security, and web application development. The step-by-step instructions and detailed explanations make it an ideal resource for beginners and anyone looking to expand their knowledge and skills in cloud computing.

TABLE OF CONTENTS

CHAPTER No.	TITLE	PAGE NO.
	ABSTRACT	VI
	LIST OF FIGURES	IX
	LIST OF ABBREVIATIONS	X
1	INTRODUCTION	1
	1.1 Cloud Computing	2
	1.2 Advantages of Cloud computing	3
	1.3 Cloud Service Providers	4
	1.4 Amazon AWS	5
	1.5 Security Groups in EC2	6
	1.6 Key pair in EC2	7
	1.7 Instances in EC2	8
	1.8 Types of OS in AWS EC2	9
	1.9 Types of Clouds	10
	1.10 LAAS Clouds	11
	1.11 SSH Client and Connections	12
	1.12 HTTP Client and Connections	13
	1.13 Ping	14
	1.14 Recommended Security Rules in	15

	EC2	
2	SYSTEM ANALYSIS	16
	2.1 Aim & Scope	16

	2.2 How does it work	17
3	SYSTEM DESIGN AND IMPLEMENTATION	19
	3.1 Algorithm	19
	3.2 Ideation Map	21
4	RESULT & DISCUSSION	22
	4.1 Result	22
5	CONCLUSION	24
	5.1 REFERENCE	25

LIST OF FIGURES

FIGURE NO.	FIGURE NAME	PAGE NO.
4.1	Connecting to EC2 instance via SSH	25
4.2	Ping failing to the EC2 instance	26
4.3	Http request working	27

LIST OF ABBREVIATIONS

ABBREVIATION	EXPANSION
AWS	Amazon Web Services
EC2	Elastic Compute Cloud
GP2	General Purpose SSD

CHAPTER 1

INTRODUCTION

Amazon Web Services (AWS) is a cloud computing platform that provides a wide range of services, including compute, storage, networking, and databases. One of the most popular services offered by AWS is Elastic Compute Cloud (EC2), which provides resizable compute capacity in the cloud, allowing users to launch virtual machines with various operating systems and applications. EC2 is a powerful tool for developers and businesses that want to quickly launch and scale their workloads in the cloud.

Creating and configuring an EC2 instance is an essential skill for anyone working with AWS. It allows users to launch virtual machines and configure their security rules, allowing incoming traffic and blocking unwanted requests. This project focuses on creating an EC2 instance of Ubuntu and configuring its security rules to accept incoming HTTP and SSH traffic while blocking ping requests from the public network. By following the step-by-step process outlined in this project, users can quickly and easily create a secure EC2 instance and start using it for various workloads and applications.

The project aims to provide a practical introduction to creating and configuring an AWS EC2 instance, which is essential for anyone working with cloud computing. It outlines a 10-step process that includes creating an AWS account, launching an EC2 instance, configuring security groups, generating a key pair for SSH access, and testing the instance. The project also provides an explanation of HTTP clients and requests, which are essential components in web application development for communicating with web servers using the HTTP protocol. By following the steps outlined in this project, users can gain a solid understanding of how to create and configure an EC2 instance, allowing them to leverage AWS for their workloads and applications.

1.1 Cloud Computing

Cloud computing refers to the practice of utilising remote servers on the internet to store, manage, and process data rather than using a local server or a personal computer. The concept of cloud computing is based on the sharing of resources, software, and information through a network of computers connected through the internet. This technology allows businesses and individuals to access and utilise a range of computing resources, including storage, processing power, and software applications, without the need to invest in expensive hardware or software.

Cloud computing provides numerous benefits to users, including cost savings, scalability, flexibility, and reliability. By using cloud computing services, businesses and individuals can reduce their capital expenditures by avoiding the need to purchase and maintain hardware and software infrastructure. Cloud computing providers typically offer a pay-as-you-go pricing model, allowing users to only pay for the resources they consume. Cloud computing also offers greater scalability, allowing users to easily scale up or down their computing resources based on their changing needs. Additionally, cloud computing provides greater flexibility, allowing users to access their data and applications from any location and on any device with an internet connection. Finally, cloud computing is more reliable than traditional computing, as cloud providers typically offer robust security and disaster recovery capabilities, ensuring that users' data is protected and available at all times.

However, cloud computing is not without its challenges. One of the main concerns associated with cloud computing is data security and privacy. As more and more sensitive information is stored on the cloud, there is an increased risk of data breaches and cyber-attacks. Cloud providers must take robust measures to ensure the security of their customers' data, including implementing advanced encryption and access control measures, regular security audits and penetration testing, and ensuring compliance with relevant data protection regulations such as GDPR and CCPA.

Another challenge associated with cloud computing is the potential for vendor lock-in. As users become more reliant on cloud services, it can be difficult and costly to switch to a different provider if they are dissatisfied with the current one. Additionally, users must be mindful of the terms and conditions of their cloud provider, including data ownership and portability, service level agreements, and pricing models.

1.2 Advantages of Cloud Computing

Cloud computing has revolutionised the way individuals and organisations store, process, and manage data. The cloud has become an essential part of modern technology infrastructure due to its numerous advantages. Here are some of the most significant advantages of cloud computing.

Cost Savings:

One of the most significant benefits of cloud computing is cost savings. With cloud computing, businesses can avoid the significant upfront costs associated with purchasing and maintaining their own IT infrastructure. Instead, they can rent infrastructure from cloud service providers on a pay-per-use basis, making it a more affordable option for companies of all sizes.

Scalability:

Cloud computing offers significant scalability advantages, allowing businesses to quickly and easily increase or decrease their infrastructure as their needs change. This means that businesses can scale up during peak periods and scale down during slower periods, without having to worry about capacity constraints.

Flexibility:

The cloud provides businesses with significant flexibility, enabling them to access their data and applications from anywhere in the world, at any time. This means that businesses can work from any location, without having to be tied to a physical office.

Collaboration:

Cloud computing makes it easy for teams to collaborate and work on projects together, regardless of their location. Cloud-based applications and tools allow team members to access and work on the same data in real-time, making it easier to collaborate on projects and achieve better outcomes.

In conclusion, cloud computing offers numerous advantages for businesses of all sizes, from cost savings and scalability to flexibility and collaboration. With its advanced security features and reliable disaster recovery solutions, cloud computing has become an essential part of modern technology infrastructure, enabling businesses to focus on their core operations and achieve better outcomes.

1.3 Different Cloud Service Providers

There are many different cloud service providers in the market, ranging from large global players such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP) to smaller niche providers that specialise in specific areas such as data analytics, security, or compliance. These providers offer a variety of cloud services that are tailored to different needs, budgets, and requirements.

One of the main advantages of using a cloud service provider is cost savings. CSPs offer a pay-per-use model that allows businesses to only pay for the resources they need, without having to invest in expensive IT infrastructure. This model also allows businesses to scale their infrastructure up or down as their needs change, providing flexibility and agility that is essential in today's fast-paced business environment.

Another advantage of using a cloud service provider is scalability. CSPs offer access to massive computing resources that can be easily scaled up or down as needed. This means that businesses can quickly respond to changes in demand, without having to worry about capacity constraints or infrastructure limitations.

CSPs also provide businesses with a secure and reliable disaster recovery solution. With multiple data centres in different geographic locations, CSPs ensure that businesses can quickly and easily recover their data and applications in the event of a disaster or outage.

Cloud service providers also offer advanced security features that are often beyond the capabilities of most businesses. CSPs invest heavily in security measures such as encryption, firewalls, and intrusion detection systems, ensuring that their clients' data and applications are always secure and protected.

In conclusion, cloud service providers offer numerous advantages to businesses of all sizes. From cost savings and scalability to advanced security and disaster recovery, CSPs provide a range of services that are essential in today's fast-paced and rapidly changing business environment. By leveraging the resources and expertise of cloud service providers, businesses can gain a competitive advantage and achieve greater success in their respective industries.

1.4 Amazon AWS

Amazon Web Services (AWS) is a cloud service provider that offers a wide range of cloud computing services to individuals and businesses. AWS is the largest and most popular cloud service provider in the world, with a market share of over 30%.

AWS provides a comprehensive set of cloud computing services that include compute, storage, networking, database, analytics, machine learning, security, and more. These services are delivered on a pay-per-use model, which means that businesses only pay for the resources they use, making it an affordable option for businesses of all sizes.

One of the most popular AWS services is Amazon Elastic Compute Cloud (EC2), which provides scalable computing capacity in the cloud. With EC2, businesses can quickly and easily launch virtual servers, scale up or down as needed, and pay only for the computing resources they use.

AWS also offers a range of analytics and machine learning services that allow businesses to extract insights from their data and make better decisions. Services such as Amazon Kinesis, Amazon Redshift, and Amazon SageMaker provide powerful tools for data processing, data warehousing, and machine learning.

AWS is known for its strong focus on security, with a comprehensive set of security services and features that help businesses protect their data and applications in the cloud. AWS provides security features such as encryption, firewalls, identity and access management, and more, ensuring that businesses can rely on AWS to keep their data and applications secure.

Finally, AWS provides a global infrastructure that allows businesses to deploy their applications and services in multiple regions around the world. This means that businesses can provide a fast and reliable user experience to their customers, no matter where they are located.

In conclusion, AWS provides a comprehensive set of cloud computing services that are essential for businesses of all sizes. With its focus on affordability, scalability, security, and global reach, AWS has become the go-to cloud service provider for businesses around the world. Whether businesses need computing capacity, storage, analytics, or machine learning, AWS offers a solution that can help them achieve their goals and succeed in today's fast-paced business environment.

1.5 Security groups in ec2

Security groups in Amazon Elastic Compute Cloud (EC2) are a critical aspect of securing a cloud computing environment. A security group acts as a virtual firewall that controls the inbound and outbound traffic to an EC2 instance.

When an EC2 instance is launched, it is automatically assigned to a default security group. However, it is important to configure the security group settings to ensure that only authorised traffic is allowed to access the instance.

Security groups are defined at the instance level, which means that they are applied to all instances that are associated with them. This makes it easy to manage security settings across multiple instances.

In AWS, security groups are based on a set of rules that specify which traffic is allowed to access an instance. These rules can be configured to allow traffic from specific IP addresses, protocols, and ports. For example, a rule might be set to allow HTTP traffic from all IP addresses.

Security groups can also be configured to allow traffic from other security groups within the same AWS account. This makes it easy to create a secure network of instances that can communicate with each other, without allowing access to unauthorised traffic.

One of the key benefits of using security groups in EC2 is that they are stateful. This means that if an inbound rule is added to allow traffic to an instance, the corresponding outbound rule is automatically created to allow the response traffic back out. This simplifies the configuration process and helps ensure that traffic is allowed in both directions.

In addition, security groups can be modified in real-time, allowing businesses to quickly respond to changing security requirements. For example, if a new application is launched that requires a specific port to be opened, the security group can be updated immediately to allow access to that port.

In conclusion, security groups are a critical aspect of securing EC2 instances in the cloud. By defining rules that control inbound and outbound traffic, businesses can ensure that their instances are protected from unauthorised access. With the ability to configure rules based on specific IP addresses, protocols, and ports, security groups provide a flexible and powerful tool for managing security in AWS.

1.6 Key pair in EC2

In Amazon Elastic Compute Cloud (EC2), a key pair is a secure way to log in to an EC2 instance. A key pair is essentially a set of two keys: a public key and a private key. The private key is kept secret and used to authenticate the user when logging in to the instance, while the public key is uploaded to the instance and used to encrypt data that can only be decrypted by the private key.

When launching an EC2 instance, a key pair can be selected or created. If a key pair is selected, it can be used to log in to the instance through a secure shell (SSH) client, while if a new key pair is created, the public key is automatically downloaded to the local computer and the private key is stored securely in AWS.

Once the key pair is selected or created, it is associated with the instance, allowing the user to securely log in to the instance using the private key. This is done by specifying the private key file when connecting to the instance using SSH.

The use of key pairs in EC2 is a crucial aspect of securing access to instances in the cloud. By using a private key for authentication, rather than a password, the risk of unauthorised access is greatly reduced. Additionally, the use of key pairs ensures that all data transferred to and from the instance is encrypted, providing an additional layer of security.

It is important to keep the private key file secure, as anyone with access to the private key can log in to the instance. The private key file should be kept in a secure location and only shared with authorised personnel.

In conclusion, the use of key pairs in EC2 is a crucial aspect of securing access to instances in the cloud. By using a private key for authentication, rather than a password, the risk of unauthorised access is greatly reduced. Additionally, the use of key pairs ensures that all data transferred to and from the instance is encrypted, providing an additional layer of security.

1.7 Instances in EC2

In Amazon Elastic Compute Cloud (EC2), an instance is a virtual machine that can be launched in the cloud. Each instance is essentially a virtualised server that can be configured with a variety of operating systems, applications, and software configurations.

Instances in EC2 are launched from Amazon Machine Images (AMIs), which are pre-configured templates that contain the necessary software and configuration to launch an instance. AMIs can be created by users or obtained from the AWS Marketplace.

When launching an instance in EC2, users have the ability to choose from a variety of instance types, which vary in terms of computing power, memory, and storage capacity. The type of instance selected will depend on the specific needs of the application or workload being run on the instance.

Instances can also be configured with various storage options, including Elastic Block Store (EBS) volumes, which provide persistent storage for data, and instance store volumes, which provide temporary storage for data that is lost when the instance is terminated.

In addition to selecting the instance type and storage options, users can also configure various networking options, including security groups, which control inbound and outbound traffic to the instance, and Elastic IP addresses, which provide a static IP address that can be associated with the instance.

One of the key benefits of using instances in EC2 is the ability to scale resources up or down based on demand. Instances can be launched or terminated as needed, providing businesses with the flexibility to quickly adjust their computing resources to meet changing needs.

Overall, instances in EC2 provide businesses with a flexible, scalable, and cost-effective way to run their applications and workloads in the cloud. With a wide range of instance types, storage options, and networking configurations, businesses can tailor their instances to meet their specific needs, while taking advantage of the scalability and flexibility of the cloud.

1.8 Types of OS in EC2

In Amazon Elastic Compute Cloud (EC2), users have the ability to launch instances with a variety of operating systems (OS) to meet their specific needs. Some of the most common types of operating systems available for use in EC2 include:

Amazon Linux: Amazon's own Linux-based OS, designed to provide a secure and stable platform for running applications in the cloud. It is optimised for use with AWS services and includes a range of pre-installed packages and tools.

Ubuntu: A popular open-source Linux distribution that is widely used in the cloud. It offers a range of pre-installed packages and tools, and is known for its ease of use and robust security features.

Microsoft Windows: EC2 supports a range of Windows operating systems, including Windows Server 2019, Windows Server 2016, and Windows Server 2012 R2. Windows instances can be used to run a variety of applications, including Microsoft SQL Server, Exchange, and SharePoint.

Red Hat Enterprise Linux: A popular commercial Linux distribution that offers advanced security features and support for enterprise workloads. It is widely used in the financial services, healthcare, and government sectors.

SUSE Linux Enterprise Server: A commercial Linux distribution that is designed for use in the enterprise. It offers a range of advanced features, including high availability, clustering, and virtualisation.

In addition to these common types of operating systems, users can also launch instances with a variety of other Linux distributions, including CentOS, Debian, and Fedora, as well as specialised operating systems designed for specific workloads, such as Amazon Elastic Container Service for Kubernetes (EKS) optimised AMIs.

Overall, the ability to launch instances with a range of operating systems in EC2 provides users with the flexibility and choice they need to meet their specific needs. Whether running Linux or Windows workloads, or specialised workloads, users can choose the right operating system to optimise performance and meet their requirements.

1.9 Types Of Clouds

Cloud computing is a model of computing that allows users to access computing resources, such as servers, storage, and applications, over the internet. There are three primary types of clouds: public, private, and hybrid.

Public Cloud: Public clouds are provided by third-party service providers, such as Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform. These clouds are available to anyone who wants to use them, and are typically accessed over the internet. Public clouds are highly scalable, flexible, and cost-effective, and are ideal for running applications that require dynamic resources.

Private Cloud: Private clouds are built and maintained by organisations for their own use. Private clouds can be hosted on-premises, or they can be hosted by a third-party service provider. Private clouds are ideal for organisations that require high levels of security and control over their computing resources.

Hybrid Cloud: Hybrid clouds combine elements of public and private clouds, allowing organisations to take advantage of the benefits of both. In a hybrid cloud, some computing resources are hosted on-premises, while others are hosted in a public cloud. Hybrid clouds are ideal for organisations that require high levels of security and control, while also needing the scalability and flexibility of a public cloud.

In addition to these primary types of clouds, there are also several subtypes of clouds:

Community Cloud: A community cloud is a cloud infrastructure that is shared by several organisations with similar needs. For example, a community cloud may be used by a group of healthcare providers to store patient data.

Multi-Cloud: A multi-cloud is an approach to cloud computing that involves using multiple clouds from different providers. Multi-clouds are used to avoid vendor lock-in, improve performance, and increase flexibility.

Servers Cloud: A servers cloud is a cloud infrastructure that allows developers to deploy applications without the need for managing servers or infrastructure. Servers clouds are ideal for running small, event-driven applications.

Overall, the different types of clouds provide organisations with a range of options for deploying and managing their computing resources, and can be tailored to meet the specific needs of each organisation.

1.10 LAAS Clouds

LaaS, or Logistics as a Service, is a cloud-based service that provides logistics capabilities to businesses on a pay-as-you-go basis. LaaS clouds offer businesses access to advanced logistics tools, software, and services that can help them optimise their supply chain operations, improve efficiency, and reduce costs.

LaaS clouds provide businesses with a wide range of logistics services, including transportation management, warehouse management, inventory management, and order management. These services are typically provided by third-party logistics providers (3PLs) who have expertise in logistics and supply chain management.

One of the key benefits of LaaS clouds is their scalability. Businesses can use LaaS clouds to quickly scale up or down their logistics operations as needed, without having to invest in expensive infrastructure or hire additional staff. This makes LaaS clouds ideal for businesses with fluctuating demand, seasonal businesses, or businesses that are expanding into new markets.

Another benefit of LaaS clouds is their flexibility. Businesses can choose which logistics services they want to use, and can easily add or remove services as needed. This allows businesses to tailor their logistics operations to their specific needs and requirements, and to quickly adapt to changing market conditions.

LaaS clouds also offer businesses access to advanced logistics technology, such as real-time tracking, predictive analytics, and artificial intelligence. These technologies can help businesses optimise their logistics operations, reduce costs, and improve customer satisfaction.

Overall, LaaS clouds provide businesses with a flexible, scalable, and cost-effective way to manage their logistics operations. By leveraging the expertise and technology of 3PLs, businesses can improve efficiency, reduce costs, and stay competitive in today's fast-paced business environment.

1.11 SSH Clients and Connections

SSH (Secure Shell) is a protocol that allows users to securely connect to remote computers and servers over the internet. SSH clients are software applications that enable users to initiate and manage SSH connections. These clients provide a secure and encrypted way to access remote computers and servers, and are widely used by system administrators, developers, and other IT professionals.

SSH clients use the SSH protocol to establish a secure connection between the user's local computer and the remote computer or server. This connection allows the user to access and control the remote computer as if they were physically present at the remote location. SSH clients typically provide a command-line interface or a graphical user interface (GUI) for managing the SSH connection.

When establishing an SSH connection, the client first authenticates the user's identity using a username and password or a public/private key pair. Once the user's identity is verified, the client encrypts all data transmitted between the local computer and the remote computer, ensuring that the connection is secure and private.

There are many different SSH clients available, ranging from simple command-line tools to complex GUI applications with advanced features. Some popular SSH clients include PuTTY, OpenSSH, SecureCRT, and Bitwise SSH Client.

In addition to providing a secure way to access remote computers and servers, SSH clients also offer a range of advanced features, such as tunnelling, file transfer, and port forwarding. These features allow users to access and manage remote resources more efficiently and securely.

Overall, SSH clients play a critical role in managing remote computers and servers, and are essential tools for IT professionals who need to securely access and manage remote resources. With their advanced security features and flexible capabilities, SSH clients are a reliable and effective way to manage remote connections and access remote resources.

1.12 HTTP Clients and Connections

HTTP (Hypertext Transfer Protocol) is a protocol used for transmitting data over the internet. HTTP clients are software applications that allow users to initiate and manage HTTP connections. These clients provide a way for users to access and interact with web resources, such as websites, web applications, and web services.

HTTP clients use the HTTP protocol to establish a connection between the user's local computer and the web server hosting the web resource. This connection allows the user to send requests for data and receive responses from the server. HTTP clients can be web browsers, mobile apps, or other software applications that are capable of sending HTTP requests.

When establishing an HTTP connection, the client sends an HTTP request to the web server, specifying the type of request (such as GET or POST), the resource being requested (such as a webpage or a file), and any additional parameters or headers required for the request. The server then processes the request and sends an HTTP response back to the client, containing the requested data or an error message if the request could not be fulfilled.

There are many different HTTP clients available, ranging from simple command-line tools to complex web browsers and mobile apps. Some popular HTTP clients include Google Chrome, Mozilla Firefox, Safari, and Microsoft Edge.

In addition to providing a way to access and interact with web resources, HTTP clients also offer a range of advanced features, such as caching, compression, and authentication. These features help to optimise the performance and security of HTTP connections and improve the user experience.

Overall, HTTP clients play a critical role in accessing and interacting with web resources, and are essential tools for users who need to browse the internet, use web applications, or consume web services. With their advanced features and flexible capabilities, HTTP clients are a reliable and effective way to manage HTTP connections and access web resources.

1.13 Ping

Ping is a network utility used to test the connectivity between two devices on a network. It is commonly used to troubleshoot network connectivity issues and to measure the response time of network devices, such as routers and servers.

When using the Ping utility, a small packet of data is sent from the source device to the destination device. The destination device then sends an acknowledgement packet back to the source device, indicating that the packet was received successfully. The time it takes for the acknowledgement packet to reach the source device is measured and reported as the round-trip time (RTT).

Ping is a simple but powerful tool for network troubleshooting, as it allows users to quickly identify connectivity issues and measure the response time of network devices. For example, if a user is unable to connect to a website, they can use the Ping utility to check whether the website's server is responding to requests. If the server is not responding, the user can conclude that there is a connectivity issue and take steps to resolve it.

Ping is also useful for monitoring the performance of network devices, such as routers and servers. By measuring the response time of these devices over time, users can identify trends and patterns in network performance, and take proactive steps to prevent connectivity issues before they occur.

Overall, Ping is a simple but powerful tool for network troubleshooting and performance monitoring. With its ability to test connectivity and measure response times, it is an essential tool for IT professionals and network administrators who need to ensure the reliability and performance of their network devices.

1.14 Recommended Security Rules in EC2

When setting up security rules for EC2 instances, there are several recommended best practices to follow in order to ensure the security of your resources. Here are some key security rules to consider:

- **Restrict incoming traffic:** By default, all incoming traffic to an EC2 instance is allowed, which can leave your resources vulnerable to attacks. It is recommended to restrict incoming traffic by creating security groups and only allowing traffic from trusted sources.
- **Use HTTPS for web traffic:** When accessing web resources, it is recommended to use HTTPS (HTTP over SSL/TLS) to encrypt data in transit and prevent eavesdropping and tampering.
- **Disable unnecessary ports:** It is recommended to disable any unused ports on your EC2 instances to reduce the attack surface and prevent unauthorised access.
- **Use strong passwords:** It is important to use strong passwords for all user accounts and enable multi-factor authentication to prevent unauthorised access.
- **Monitor and log access:** It is recommended to monitor and log all access to your EC2 instances, including login attempts and network traffic, in order to detect and respond to potential security incidents.
- **Keep software up-to-date:** It is important to keep all software and operating systems on your EC2 instances up-to-date with the latest security patches and updates to prevent vulnerabilities and exploits.

By following these recommended security rules, you can help to ensure the security of your EC2 instances and protect your resources from potential security threats.

CHAPTER 2

SYSTEM ANALYSIS

2.1 Aim & Scope

Aim of this project is to provide a practical introduction to creating and configuring an AWS EC2 instance, with a focus on security. The project outlines a step-by-step process for launching an EC2 instance, configuring its security rules, and testing it. Additionally, the project provides an overview of HTTP clients and requests, which are essential components in web application development for communicating with web servers using the HTTP protocol. By following the steps outlined in this project, users can gain a solid understanding of how to create and configure an EC2 instance and develop the necessary skills to work with cloud computing.

Scope of this project includes creating an EC2 instance of Ubuntu and configuring its security rules to accept incoming HTTP and SSH traffic while blocking ping requests from the public network. The project also covers testing the instance to ensure that it is working correctly. Additionally, the project provides an overview of HTTP clients and requests, explaining how they work and their importance in web application development. By the end of this project, users should have a clear understanding of how to create and configure an EC2 instance, as well as a basic understanding of HTTP clients and requests.

The intended audience for this project includes beginners who are new to AWS EC2 or anyone looking to expand their knowledge and skills in cloud computing. This project can serve as a valuable resource for developers, students, and anyone interested in learning about cloud computing and its practical applications. With the increasing demand for cloud computing skills in the workforce, this project provides an excellent starting point for anyone looking to develop those skills and stay competitive in the job market.

2.2 How does it work

Amazon Elastic Compute Cloud (EC2) is a web service that allows users to launch and manage virtual machines in the cloud. EC2 provides resizable compute capacity, enabling users to quickly and easily launch virtual machines of various sizes and configurations, depending on their workload requirements. In this project, we will walk through the process of creating an EC2 instance of Ubuntu and configuring its security rules to accept incoming HTTP and SSH traffic while blocking ping requests from the public network.

To create an EC2 instance, the first step is to create an AWS account and navigate to the EC2 dashboard. Once in the EC2 dashboard, click on the "Launch Instance" button to begin the instance creation process. The first step in the process is to choose an Amazon Machine Image (AMI). In this project, we will select the Ubuntu Server 20.04 LTS AMI. Next, we will choose an instance type based on our workload requirements. There are various instance types available, ranging from small instances to high-performance computing instances. For this project, we will select a t2.micro instance, which is the smallest instance available and suitable for low-traffic web servers and small workloads.

Once the instance type is selected, we will configure the instance details, including the number of instances, network settings, and storage options. In this project, we will launch a single instance and select the default network settings. We will also add a 8 GB General Purpose SSD (gp2) storage volume to the instance. After configuring the instance details, we will configure the instance security group. A security group acts as a virtual firewall, controlling the inbound and outbound traffic for the instance. We will create a new security group and configure it to accept incoming HTTP and SSH traffic while blocking ping requests from the public network.

To allow incoming HTTP and SSH traffic, we will create inbound rules in the security group. For HTTP traffic, we will add a rule allowing traffic on port 80. For SSH traffic, we will add a rule allowing traffic on port 22. To block ping requests, we will add a rule denying ICMP traffic. Once the security group is configured, we will review the instance details and launch the instance.

After the instance is launched, we will generate a key pair to access the instance using SSH. A key pair consists of a public key and a private key. We will download the private key and use it to access the instance using SSH. To access the instance, we will use an

SSH client, such as PuTTY or OpenSSH. Once connected to the instance, we will install Apache HTTP server to test the instance.

The HTTP client and requests section of this project provides an overview of HTTP clients and requests, explaining how they work and their importance in web application development. HTTP clients are software applications that enable users to communicate with web servers using the HTTP protocol. HTTP requests are messages sent by HTTP clients to web servers, requesting data or actions. In this project, we will use the curl command-line tool to test the HTTP server installed on the EC2 instance. We will send HTTP GET requests to the instance's IP address and verify that the HTTP server is working correctly. In conclusion, this project provides a step-by-step guide for creating and configuring an EC2 instance of Ubuntu and configuring its security rules to accept incoming HTTP and SSH traffic while blocking ping requests from the public network. The project also includes an overview of HTTP clients and requests, explaining their importance in web application development. By following the steps outlined in this project, users can gain a solid understanding of how to create and configure an EC2 instance and develop the necessary skills to work with cloud computing.

CHAPTER 3

SYSTEM DESIGN AND IMPLEMENTATION

3.1 Steps

Step 1: Choose a Cloud Provider

The first step in the project is to choose a cloud provider. There are many cloud providers available, including Amazon Web Services (AWS), Microsoft Azure, and Google Cloud Platform (GCP). Each provider offers different services, pricing models, and features, so it's important to evaluate each option and choose the one that best meets your needs.

Step 2: Create an AWS EC2 Instance

Once you have chosen a cloud provider, the next step is to create an AWS EC2 instance. An EC2 instance is a virtual machine in the cloud that can be used to run applications and services. To create an EC2 instance, you will need to choose an Amazon Machine Image (AMI), select an instance type, and configure security settings.

Step 3: Install Apache Web Server

After creating the EC2 instance, the next step is to install the Apache web server. Apache is a popular open-source web server that can be used to host web applications and websites. To install Apache, you will need to connect to your EC2 instance using SSH and run the appropriate commands to install and configure the web server.

Step 4: Configure Apache for HTTPS

Once Apache is installed, the next step is to configure it for HTTPS. HTTPS is a secure version of the HTTP protocol that encrypts data sent between the client and the server. To configure Apache for HTTPS, you will need to generate an SSL/TLS certificate and configure Apache to use it.

Step 5: Install and Configure MySQL

MySQL is an open-source relational database management system that can be used to store and manage data. To install MySQL, you will need to connect to your EC2 instance using SSH and run the appropriate commands to install and configure the database.

Step 6: Install PHP

PHP is a popular scripting language that can be used to develop dynamic web applications. To install PHP, you will need to connect to your EC2 instance using SSH and run the appropriate commands to install and configure the language.


Step 7: Develop and Deploy a Sample Web Application

Once Apache, MySQL, and PHP are installed and configured, the next step is to develop and deploy a sample web application. This can be done using a variety of tools and frameworks, including WordPress, Drupal, and Laravel. To deploy the application, you will need to upload the files to your EC2 instance and configure Apache to serve them.

Step 8: Monitor and Manage the Application

Finally, once the web application is deployed, it's important to monitor and manage it to ensure it is running smoothly and securely. This can be done using a variety of tools and services, including Amazon CloudWatch, New Relic, and Data dog. Monitoring and management involves monitoring system resources, detecting and resolving errors, and performing regular backups and updates.

3.2 Ideation Map



CHAPTER -4

RESULTS & DISCUSSION

4.1 Results

```
prajetha@192 Observium Mail Alerts % ssh -v -i Obs.pem ubuntu@ec2-35-154-97-25.ap-south-1.compute.amazonaws.com
OpenSSH_9.0p1, LibreSSL 3.3.6
debug1: Reading configuration data /etc/ssh/ssh_config
debug1: /etc/ssh/ssh_config line 21: include /etc/ssh/ssh_config.d/* matched no files
debug1: /etc/ssh/ssh_config line 54: Applying options for *
debug1: Authenticator provider $SSH_SK_PROVIDER did not resolve; disabling
debug1: Connecting to ec2-35-154-97-25.ap-south-1.compute.amazonaws.com port 22.
debug1: Connection established.
debug1: identity file Obs.pem type -i
debug1: identity file Obs.pem-cert type -1
debug1: Local version string SSH-2.0-OpenSSH_9.0
debug1: Remote protocol version 2.0, remote software version OpenSSH_9.0p1 Ubuntu-3ubuntu0.1
debug1: compat_banner: match: OpenSSH_9.0p1 Ubuntu-3ubuntu0.1 pat OpenSSH.compat.B44d48888
debug1: Authenticating to ec2-35-154-97-25.ap-south-1.compute.amazonaws.com:22 as 'ubuntu'
debug1: load_hostkeys: fopen /Users/prajetha/.ssh/known_hosts2: No such file or directory
debug1: load_hostkeys: fopen /etc/ssh/ssh_known_hosts: No such file or directory
debug1: load_hostkeys: fopen /etc/ssh/ssh_known_hosts2: No such file or directory
debug1: SSH2_MSG_KEXINIT sent
debug1: SSH2_MSG_KEXINIT received
debug1: kex: algorithm: sntrup761x25519-sha512@openssh.com
debug1: kex: host key algorithm: ecdsa-sha2-nistp521
debug1: kex: server->client cipher: chacha20-poly1305@openssh.com MAC: <implicit> compression: none
debug1: kex: client->server cipher: chacha20-poly1305@openssh.com MAC: <implicit> compression: none
debug1: expecting SSH2_MSG_KEX_ECDH_REPLY
debug1: SSH2_MSG_KEX_ECDH_REPLY received
debug1: Server host key: ssh-ed25519 SHA256:q07ZLRiue1f2qfnnN8l172ScbswT0b7mDvC0B3dg
debug1: load_hostkeys: fopen /Users/prajetha/.ssh/known_hosts2: No such file or directory
debug1: load_hostkeys: fopen /etc/ssh/ssh_known_hosts: No such file or directory
debug1: load_hostkeys: fopen /etc/ssh/ssh_known_hosts2: No such file or directory
debug1: hostkeys_find_by_hostfile: hostkeys file /Users/prajetha/.ssh/known_hosts2 does not exist
debug1: hostkeys_find_by_hostfile: hostkeys file /etc/ssh/ssh_known_hosts does not exist
debug1: hostkeys_find_by_hostfile: hostkeys file /etc/ssh/ssh_known_hosts2 does not exist
The authenticity of host 'ec2-35-154-97-25.ap-south-1.compute.amazonaws.com (35.154.97.25)' can't be established.
ED25519 key fingerprint is SHA256:q07ZLRiue1f2qfnnN8l172ScbswT0b7mDvC0B3dg.
This key is not known by any other names.
Are you sure you want to continue connecting (yes/no/[fingerprint])? yes
Warning: Permanently added 'ec2-35-154-97-25.ap-south-1.compute.amazonaws.com' (ED25519) to the list of known hosts.
debug1: rekey after 134217728 blocks
debug1: SSH2_MSG_NEWKEYS sent
debug1: expecting SSH2_MSG_NEWKEYS
debug1: SSH2_MSG_NEWKEYS received
debug1: rekey in after 134217728 blocks
debug1: get_agent_identities: bound agent to hostkey
debug1: get_agent_identities: ssh_fetch_identitylist: agent contains no identities
debug1: Will attempt key: Obs.pem explicit
debug1: SSH2_MSG_EXT_INFO received
debug1: kex.input_ext_info: server-sig-algs=ecdsa-sha2-nistp521,sk-ssh-ed25519@openssh.com,ssh-rsa,rsa-sha2-256,rsa-sha2-512,ssh-dss,ecdsa-sha2-nistp256,ecdsa-sha2-nistp384,ecdsa-sha2-nistp521,sk-ecdsa-sha2-nistp256@openssh.com,webauthn-sk-ecdsa-sha2-nistp256@openssh.com
debug1: SSH2_MSG_SERVICE_ACCEPT received
debug1: Authentications that can continue: publickey
debug1: Next authentication method: publickey
debug1: Trying private key: Obs.pem
Authenticated to ec2-35-154-97-25.ap-south-1.compute.amazonaws.com ([35.154.97.25]:22) using "publickey".
debug1: channel 0: new (client-session)
debug1: Requesting no-more-sessions@openssh.com
debug1: Entering interactive session.
debug1: pledge: filesystem
debug1: client_input_global_request: rtype hostkeys-00@openssh.com want_reply 0
debug1: client_input_hostkeys: searching /Users/prajetha/.ssh/known_hosts for ec2-35-154-97-25.ap-south-1.compute.amazonaws.com / (none)
debug1: client_input_hostkeys: searching /Users/prajetha/.ssh/known_hosts2 for ec2-35-154-97-25.ap-south-1.compute.amazonaws.com / (none)
debug1: client_input_hostkeys: hostkeys file /Users/prajetha/.ssh/known_hosts2 does not exist
debug1: Remote: /home/ubuntu/.ssh/authorized_keys: key options: agent-forwarding port-forwarding pty user-rx x11-forwarding
debug1: Sending environment.
debug1: channel 0: setting env LC_CTYPE = "UTF-8"
Learned new hostkey: RSA SHA256:w15J0EnxZGAB3NlAMct1SjZfNE/jSxebb1FP1vgGsf
Learned new hostkey: ECDSA SHA256:8lWU1RaNPvXDMt1578kws0tN2Kce3Ba8kbaIZw4
Adding new key for ec2-35-154-97-25.ap-south-1.compute.amazonaws.com to /Users/prajetha/.ssh/known_hosts: ssh-rsa SHA256:w15J0EnxZGAB3NlAMct1SjZfNE/jSxebb1FP1vgGsf
Adding new key for ec2-35-154-97-25.ap-south-1.compute.amazonaws.com to /Users/prajetha/.ssh/known_hosts: ecdsa-sha2-nistp256 SHA256:8lWU1RaNPvXDMt1578kws0tN2Kce3Ba8kbaIZw4
debug1: update known hosts: known hosts file /Users/prajetha/.ssh/known_hosts2 does not exist
Welcome to Ubuntu 22.04.2 LTS (GNU/Linux 5.15.0-103-aws x86_64)

 * Documentation:  https://help.ubuntu.com
 * Management:    https://landscape.canonical.com
 * Support:       https://ubuntu.com/advantage

System information as of Thu Apr 20 18:59:26 UTC 2023
```

Fig.4.1: Connecting to EC2 instance via SSH

```

prajeetha@192: ~$ ping -n 30 35.154.97.25
PING 35.154.97.25 (35.154.97.25): 56 data bytes
Request timeout for icmp_seq 0
Request timeout for icmp_seq 1
Request timeout for icmp_seq 2
Request timeout for icmp_seq 3
Request timeout for icmp_seq 4
Request timeout for icmp_seq 5
Request timeout for icmp_seq 6
Request timeout for icmp_seq 7
Request timeout for icmp_seq 8
Request timeout for icmp_seq 9
Request timeout for icmp_seq 10
Request timeout for icmp_seq 11
Request timeout for icmp_seq 12
Request timeout for icmp_seq 13
Request timeout for icmp_seq 14
Request timeout for icmp_seq 15
Request timeout for icmp_seq 16
Request timeout for icmp_seq 17
Request timeout for icmp_seq 18
Request timeout for icmp_seq 19
Request timeout for icmp_seq 20
Request timeout for icmp_seq 21
Request timeout for icmp_seq 22
Request timeout for icmp_seq 23
Request timeout for icmp_seq 24
Request timeout for icmp_seq 25
Request timeout for icmp_seq 26
Request timeout for icmp_seq 27
Request timeout for icmp_seq 28
Request timeout for icmp_seq 29
Request timeout for icmp_seq 30
Request timeout for icmp_seq 31
Request timeout for icmp_seq 32
Request timeout for icmp_seq 33
Request timeout for icmp_seq 34
Request timeout for icmp_seq 35
Request timeout for icmp_seq 36
Request timeout for icmp_seq 37
Request timeout for icmp_seq 38
Request timeout for icmp_seq 39
Request timeout for icmp_seq 40
Request timeout for icmp_seq 41
Request timeout for icmp_seq 42
Request timeout for icmp_seq 43
Request timeout for icmp_seq 44
Request timeout for icmp_seq 45
Request timeout for icmp_seq 46
Request timeout for icmp_seq 47
Request timeout for icmp_seq 48
Request timeout for icmp_seq 49
Request timeout for icmp_seq 50
Request timeout for icmp_seq 51
Request timeout for icmp_seq 52
Request timeout for icmp_seq 53
Request timeout for icmp_seq 54
Request timeout for icmp_seq 55
Request timeout for icmp_seq 56
Request timeout for icmp_seq 57
Request timeout for icmp_seq 58
Request timeout for icmp_seq 59
Request timeout for icmp_seq 60
Request timeout for icmp_seq 61
Request timeout for icmp_seq 62
Request timeout for icmp_seq 63
Request timeout for icmp_seq 64
Request timeout for icmp_seq 65
Request timeout for icmp_seq 66
Request timeout for icmp_seq 67
Request timeout for icmp_seq 68
Request timeout for icmp_seq 69
Request timeout for icmp_seq 70
Request timeout for icmp_seq 71
Request timeout for icmp_seq 72
Request timeout for icmp_seq 73

```

Fig 4.2 Ping failing to the EC2 instance

```

Last login: Fri Apr 21 00:37:24 on ttya00
prajeetha@192 ~ % curl http://35.154.97.25
<!DOCTYPE html>
<html>
<head>
<title>Welcome to nginx</title>
<style>
body {
width: 85%;
margin: 0 auto;
font-family: Tahoma, Verdana, Arial, sans-serif;
}
</style>
</head>
<body>
<h1>Welcome to nginx</h1>
<p>If you see this page, the nginx web server is successfully installed and
working. Further configuration is required.</p>
<p>For online documentation and support please refer to
<a href="http://nginx.org/">http://nginx.org/</a>.
Commercial support is available at
<a href="http://nginx.com/">http://nginx.com/</a>.</p>
<p>Thank you for using nginx.</p>
</body>
</html>
prajeetha@192 ~ % █

```

Fig 4.3 : Http request working

CHAPTER -5

CONCLUSION

In conclusion, this project provides a practical guide for creating and configuring an Amazon Elastic Compute Cloud (EC2) instance of Ubuntu and configuring its security rules to accept incoming HTTP and SSH traffic while blocking ping requests from the public network. EC2 is a powerful cloud computing service that allows users to launch and manage virtual machines in the cloud, providing resizable compute capacity and enabling users to quickly and easily launch virtual machines of various sizes and configurations.

By following the steps outlined in this project, users can gain a solid understanding of how to create and configure an EC2 instance and develop the necessary skills to work with cloud computing. In addition, the project also provides an overview of HTTP clients and requests, explaining their importance in web application development. Understanding how to use HTTP clients and requests is essential for building and testing web applications, and this project provides a great starting point for developers looking to build their skills in this area.

Overall, this project demonstrates the power and flexibility of cloud computing and provides users with a hands-on approach to learning about EC2 and HTTP clients and requests. With the skills learned in this project, users can continue to explore and develop their skills in cloud computing and web application development, providing a solid foundation for future projects and career opportunities.

REFERENCE

1. T. Buyya, C.S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud Computing and Emerging IT Platforms: Vision, Hype, and Reality for Delivering Computing as the 5th Utility," *Future Generation Computer Systems*, vol. 25, no. 6, pp. 599-616, Jun. 2009.
2. D. Nurmi, R. Wolski, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov, "The Eucalyptus Open-source Cloud-computing System," *Proceedings of the 1st Workshop on Cloud Computing and Its Applications*, pp. 124-131, Dec. 2008.
3. B. Sotomayor, R. S. Montero, I. M. Llorente, and I. Foster, "Virtual Infrastructure Management in Private and Hybrid Clouds," *IEEE Internet Computing*, vol. 13, no. 5, pp. 14-22, Sept. 2009.
4. Y. Wei, Z. Zheng, and Y. Yang, "A Study on Cloud Computing Security," *2010 International Conference on Computational Intelligence and Security*, pp. 138-141, Dec. 2010.
5. S. K. Garg and R. Buyya, "Green Cloud Computing: Balancing Energy in Processing, Storage, and Transport," *Proceedings of the 19th International Symposium on High Performance Distributed Computing*, pp. 1-12, Jun. 2010.
6. S. P. Chowdhury, M. R. Rahman, and R. Boutaba, "Virtual Machine Migration in Cloud Computing Environments: A Survey," *IEEE Communications Surveys & Tutorials*, vol. 13, no. 3, pp. 452-469, Third Quarter 2011.
7. J. Altmann, J. G. Pérez-Sánchez, and J. A. Caminero, "Economic Models for Cloud Service Markets: Pricing and Quality of Service," *Journal of Grid Computing*, vol. 11, no. 1, pp. 57-79, Mar. 2013.
8. R. Buyya, R. Ranjan, and R. N. Calheiros, "Modeling and Simulation of Scalable Cloud Computing Environments and the CloudSim Toolkit: Challenges and Opportunities," *Proceedings of the 7th High Performance Computing and Simulation Conference*, pp. 1-11, Jul. 2015.
9. Y. Li, S. Li, Z. Li, Y. Yang, and R. Li, "A Survey of Cloud Computing and Security Issues," *Journal of Network and Computer Applications*, vol. 43, pp. 1-10, Nov. 2014.
10. J. M. Alvarez-Rodríguez, M. J. Suarez-Cabal, and J. F. Martínez-Santos, "Adaptive Monitoring in Cloud Computing Environments," *IEEE Transactions on Services Computing*, vol. 9, no. 1, pp. 30-42, Jan.-Feb. 2016.