EmpMonitor On AWS

Deployment Guide

EmpMonitor
Keep your workforce united

Table of Contents

EmpMonitor on AWS Overview

All the EmpMonitor solutions are ready to deploy on Amazon's distributed, highly scalable, and reliable Amazon Web Services platform (AWS) as a Private Cloud option. This deployment guide will help you discover what you can expect from your EmpMonitor on AWS deployment and provide you with installation prerequisites, step by step instructions, technical, and support information.

Benefits of Deploying EmpMonitor on AWS

If you prefer on-premise deployments but unwilling to incur the cost and hassle of an in-house infrastructure or you want to move to the Cloud but have been concerned about compliance, security, or have made the transition and already host on AWS, then EmpMonitor on AWS is the right choice for you. With our AWS Private Cloud hosting option, enjoy the benefits of both worlds: cost and scalability comparable to a Cloud deployment, and security and flexibility similar to an On-Premise solution. Here are some infrastructure benefits you can expect if you choose to deploy EmpMonitor on AWS:

Flexible & Competitive Pricing
You need to pay only for the AWS resources you consume (ie. CPU, storage, memory). EmpMonitor's recommended configuration for a standard deployment (m4.xlarge instance, up to 100 users).

Optimized Environment
Choose the region, OS, storage, database, etc. For example, you can use Amazon S3 storage for EmpMonitor's session recordings and Amazon RDS databases for session logs that are highly optimized for performance on AWS.

Easy Deployment
Create an EmpMonitor server with a single click by launching the EmpMonitor AMI (Amazon Machine Image) from the AWS Marketplace.

Reliability
Support for High Availability (HA), redundancy with multi-geo replications, on-demand backup, and disaster recovery.

Scaling
Vertical and horizontal scaling with optional auto-scaling that adjusts capacity based on demand.

Central Management Console
Configure and manage all your deployments from one central location.

Security and Compliance

Firewall, encryption at rest, SSL encryption, VLAN, SSH tunnels, 2FA, IP whitelisting, and encrypted disks allowing for easy regulatory compliance for HIPAA, GDPR, PCI DSS and more make it ideal for EmpMonitor customers in government, healthcare, finance, and other regulated industries.

Prerequisites

To get started, you will need:
- An AWS account
- Your EmpMonitor license key, available from EmpMonitor Self-hosted portal at:
- An SSH client like Putty if you are using Windows

Step by Step Instructions

1. Creating an AWS Instance and Accessing it with SSH

Step 1-1

Visit: https://www.teramind.co/deployment/aws and click on the **Check out EmpMonitor on Amazon Marketplace** button, which will redirect you to EmpMonitor's deployment page on AWS.

Once you are there, click on the **Continue to Subscribe** button at the top right corner of the screen.

Step 1-2

Once you reach 'Subscribe to this software' screen, click on the **Accept Terms** button.

Step 1-3

Click on the **Continue to Configuration** button.

Step 1-4

In the 'Configure this software' screen, you can select the Fulfilment Option, Software Version, and the Region. Make sure to choose a region that is closest to you for better performance.

Click the **Continue to Launch** button when done.

Step 1-5

In the 'Launch this software' page, tap on the Choose Action pull-down menu and select Launch through EC2.

Click the **Launch** button.

Step 1-6

On the 'Choose an Instance Type' screen, select an instance type from the list of available instances. EmpMonitor recommends at least an m4.xlarge for a typical deployment of up to 300 users.

Once done, click on the **Review and Launch** button, which you'll find at the bottom of the page.

Step 1-7

The 'Review Instance Launch' screen shows all the details associated with the instance you have selected. You can edit any configuration if needed.

Click **Launch** to accept the configurations and launch the instance.

Step 1-8

Once the instance is ready, a window will pop-up where you can choose an existing key pair or create a new one.

From the first pull-down menu, select Create a new key pair and give it a name. Tap on the **Download Key Pair** button and save the file in a secure place. After saving, click **Launch** Instances.

The instance will be initialized in a few moments.

Step 1-9

In the Launch Status window, click the **View Instances** button at the bottom to see your newly created instance(s).

Step 1-10

Now, we need to assign an elastic IP to the instance. Otherwise, each time you restart the instance, it's going to assign a new IP address causing the EmpMonitor Agent to lose connection.

To assign an elastic IP, click on the **Elastic IPs** (under NETWORK & SECURITY) from the left sidebar. Then tap on the **Allocate new address** button at the right side of the screen.

Step 1-11

We will use the default settings on the Allocate new address screen. Just click on the **Allocate** button.

Step 1-12

After allocating the IP, click on the **IP address** to select it.

Step 1-13

Click the **Actions** menu and select the **Associate address** from the pop-up menu.

Step 1-14

Tap on the **Instance** box and select the EmpMonitor EC2 instance you created earlier and then click on the **Associate** button.

Step 1-15

Go back to the **Instance** menu. When the Instance State is running, it means it's ready for accepting connections. Copy Public DNS (IPv4) address for the instance you created and copy or write it down.

Step 1-16

Launch an SSH session. If you are using Windows, you can use a tool like Putty or a similar utility for the SSH. But, make sure you have administrative access.

Type: ssh -i "path/to/your/keypair.pem file" admin@

Paste/type the Public DNS you copied in the previous step just after the @ symbol in the command prompt and press Enter. The server will be ready in a few minutes.

Step 1-17

Once the server is ready, you will see the following message, 'Welcome to EmpMonitor!.' Below the message, you will be presented options to assign roles for the nodes/hosts, beginning with the Master role. However, before you can assign a role, you will need to finish the rest of the steps.

Keep this command window open as you will use it to finish the setup later.

2. Database Considerations

Before proceeding any further with the setup process, you need to determine if you would like to use external databases or storage with the instance. External databases or storage aren't mandatory. While, you may want to use Amazon RDS as the database, and Amazon S3 for screenshots. RDS and S3 will improve the scalability of your platform. And they are also recommended for deployments over 100 concurrent users.

In this guide you will learn how to use an external database and storage as an example using an AWS RDS database and the AWS S3 storage service. First, it will show you how to create an RDS database. If you already know how to create an RDS database, you can skip the next section.

3. Creating the RDS Database

Step 3-1

From your EC2 Management Console, click on the Services from the top left hand side. A few options will show up.

Type RDS in the search box.

Select the **RDS** under the Database category.

Step 3-2

Select **Databases** from the left-hand sidebar.

Then, tap on the **Create database** button on the top-right corner of the screen.

Step 3-3

On the 'Select engine' screen, click on PostgreSQL. Then tap on - Next button to go to the next screen.

Step 3-4

On the 'Choose use case' screen, select either **Production** or **Dev/Test** depending on your own choice. In this case, we will use the Dev/Test option.

Then click on Next to go to the next screen.

Step 3-5

On the 'Specify DB details' screen, you will notice that there's a **Free Tier** option. You can enable this if you are running test databases and don't want to pay for its usage.

Scroll down and you will find the 'Settings' section.

Type a name in the **DB instance identifier** field. Assign a **Master user name** and a **password**.

When done, tap on the Next button.

Step 3-6

On the 'Configure advanced settings' screen you have options to configure Network & Security, Log exports, encryption, backup etc. You can adjust them as per your requirement.

In this guide, we will leave them to their default settings.

Tap on the **Create database** button.

Step 3-7

Once your DB instance is created you will see a **View DB instance** details button. Click on it.

Step 3-8

On the database 'Summary' screen, you will find the details like the Engine, DB instance class, CPU Utilization, DB Connections, Free Storage Space, etc.

When the **DB instance status** changes to 'Available', which means the database will be ready for use.

4. StorageConsiderations

EmpMonitor doesn't need external storage. However, EmpMonitor can still utilize a file storage system for screenshots, captured attachments, documents, and metadata such as user images. AWS S3 is an efficient, secure, scalable, and redundant solution to store objects like these. The S3 storage will help you to improve the scalability of your platform and is recommended for deployments of over 100 concurrent monitored users.

5. Creating an IAM Role for the S3 Storage

For using the AWS S3 storage, you will need an Identity and Access Management or IAM profile. IAM enables you to manage access to your AWS services and resources securely. Using IAM, you can create and manage AWS users and groups and control user permissions. This guide will show you how to create an IAM profile. If you already know how to create one, you can skip this section.

Step 5-1

Go to the **Instances** screen from your EC2 Management Console. Make sure to select the instance you want to use.

Tap on the Actions button on top of the instance names.

From the menu, select **Instance Settings>Attach/Replace IAM Role**.

Step 5-2

On the 'Attach/Replace IAM Role' screen, tap on the **Create new IAM role** link.

Step 5-3

On the 'Roles' screen you will find an explanation of the IAM roles and links to some resources and documentations.

Tap on the **Create role** button to continue.

Step 5-4

Select the default AWS service as the type for the trusted entity.

Click the **Next: Permissions** button.

Step 5-5

From the list of permission policies, select the one named **'AmazonS3FullAccess'** (you can quickly locate it using the Search field).

Tap on the **Next: Tags** button to continue.

Step 5-6

AlthoughTags are optional and we will not use any tags in this guide. You can still add some tags if you want.

Click on the Next: Review button to go to the next screen.

Step 5-7
On the 'Review' screen, give a role name and tap on the **Create role** button.

6. Setting Up the Elastic Load Balance (optional)

Optionally, you can also add an Elastic Load Balancer or ELB in AWS. EmpMonitor supports the AWS Classic ELB so that all HTTP, HTTPS, and TCP traffic can be used for the same Virtual Machine (VM).

One of the benefits of using ELB is that you can manage your domain certificates from the AWS portal directly. At the moment it's only supported on single-node deployments. However, you can use a single master plus additional elastic or mining nodes with ELB.

For the setup of ELB, do the following:

Step 6-1
Click on the **Load Balancer** link found under the LOAD BALANCING section on the left sidebar.

Tap on the **Create Load Balancer** button on top.

Step 6-2
On the Select load balancer type screen, Click on the **Create** button under the Classic Load Balancer. It will take you to a step-by-step wizard for the setup of the load balancer.

Step 6-3
On the Step 1: Define Load Balancer screen, give a name for your load balancer. For example, **'testloadbalancer'**.

Use the **Add** button to add 3 protocol rows and populate them as follows:

LB Protocol	LB Port	Instance Protocol	Instance Port
HTTP	80	HTTP	80
TCP	443	TCP	443
TCP	10000	TCP	10000

Note: to use domain certificates, replace the TCP protocols for port 443 with HTTPS.

Step 6-4
Tap on the **Next: Assign Security Groups** button at the bottom-right corner of the screen to go to the next step.

Step 6-5
On the Assign Security Groups screen, select the Virtual Private Cloud (VPC) security group for your instance, for example: **TestGroup**.

If you need help with creating a security group, please consult Amazon's AWS documentation.

Tap on the **Next: Configure Security Settings** button at the bottom-right corner of the screen to go to the next step.

Step 6-6

On the Configure Security Settings screen, you might see a **warning** that your load balancer isn't using a secure listener. You can ignore the message as EmpMonitor will handle it.

Click the **Next: Configure Health Check** button at the bottom-right corner of the screen to go to the next step.

Step 6-7

On the Configure Health Check screen, change the Ping Port to **443**.
And set the Ping Path to root, i.e. just put a / (slash).

Click on the **Next: Add EC2 Instance** button at the bottom-right corner of the screen to go to the next step.

Step 6-8

On the Add EC2 Instances screen, select the instance(s) you need from the list of running instances.

Tap on the **Next: Add Tags** button at the bottom-right corner of the screen to go to the next step.

Step 6-9

Tags are key-value pairs used for metadata you can assign to an AWS resource. Tags make it easier to manage, search for, and filter resources. For this demo, we will not use any tags.

Click the **Next: Review and Create** button at the bottom-right corner of the screen to go to the next and last step of the wizard.

Step 6-10

On the Review screen, check to make sure all the settings are correct. Once satisfied, you can tap on the Create button.

Step 6-11

Once the load balancer is created, you will see a message saying, 'Successfully created load balancer'. Click the **Close** button on the right side of the screen to go back to the **Load Balancers** screen.

Step 6-12

On the Load Balancers screen, there you will find the status of your newly created load balancer. Note that, it might take a few minutes for your load balancer to become active.

7. Finishing the Deployment

Now that we have created the RDS database, the IAM profile for the S3 storage and optionally an ELB, we are ready to continue with the rest of the deployment process.

Step 7-1

Go to the Instances screen from your EC2 Management Console. Make sure to select the instance you want to use.
Click on the Actions button on top of the instance names.

From the menu, select Instance Settings>Attach/Replace IAM Role.

Step 7-2

On the 'Attach/Replace IAM Role' screen, click on the **IAM** role drop-down menu and select the IAM you created in the previous steps.

Tap on Apply to assign the IAM to the selected instance.

Step 7-3

Go back to the command prompt window you used in 'Step 1-12'. Give the master role a name.

On the 'use external database?' prompt enter Yes. Supply the RDS database IP / hostname, user name, and password you created before. Press enter when done.

EmpMonitor will initialize the database, tables, and required indexes automatically.

Step 7-4

After EmpMonitor has finished the database initializations, it will ask, 'Use S3 for storing data?'. Type Yes.

Type Yes to the 'Create new buckets?' prompt.

You will be creating three buckets: one for screenshots, one for user content i.e. attachments or docs, and the third one for application objects. Give unique names for each bucket.

Please see the S3 Bucket Naming Requirements below for information on rules related to the naming of the buckets.

Finally, you will be presented with a **https link**. Click this link to continue

S3 Bucket Naming Requirements

The S3 buckets must have a name that conforms with the naming requirements for non-US Standard regions. Amazon S3 defines a bucket name as a series of one or more labels, separated by periods, that must adhere to the following rules:

- The bucket name can be between 3 and 63 characters long, and can contain only lower-case characters, numbers, periods and dashes.
- Each label in the bucket name must start with a lowercase letter or number.
- The bucket name cannot contain underscores, end with a dash, have consecutive periods or use dashes adjacent to periods.
- The bucket name cannot be formatted as an IP address i.e. 198.51.100.24.

Step 7-5

When you click on the EmpMonitor server link from the previous step, a browser window will open. You may be displayed a warning message. This is because you are using HTTPS connection without SSL certificate. Most browsers will allow you to continue with an override action.

In case you are using Google Chrome, click on the ADVANCED link on the page and select the Proceed to option.

Step 7-6

You will get the "Welcome To EmpMonitor" screen. There you need to select the language and timezone of your country.

Then click on the CONTINUE button.

Step 7-7

Next 'Let's secure your EmpMonitor installation' screen will open, you need to create an Admin account with an email and a password.

Step 7-8

In a new browser, open (EmpMonitor portal link). There you can login with your registered admin email and password.

In the license tab, you will find a list of licenses. Select the **Key** link under the **Action** column of the list. A pop-up will be shown, which displays your license key.

Copy and save that license key in your records. Or you can also write it down to remember when needed.

Step 7-9

Return to the EmpMonitor dashboard page and enter your license key and then tap on the continue button.

Step 7-10

Now, you are ready to install the EmpMonitor agent and you can start monitoring the targeted computer user(s).

You can also do this later when you have time.

For the installation of EmpMonitor agent here are the steps:

You can also check our user guide or watch video to learn the detailed process of Downloading and Installing the EmpMonitor Hidden Agent

8. Setting up the OCR (optional)

If you want to use the OCR feature, you will need to set up two nodes along with a master node.

Step 8-1

Create the master VM as usual.

Then, create two additional VMs. The first node will store the OCR videos. In EmpMonitor, we will refer to it as teracv. The second node is for the OCR database. In EmpMonitor, we will refer to it as elastic.

Make sure all the nodes are in the same Security Group.

Step 8-2

From your AWS portal's Instances page, copy the Public DNS (IPv4) address for the teracv instance.

Step 8-3

Launch a SSH session. If you are on Windows, you can use a tool like Putty or a similar utility for the SSH. Make sure you have administrative access.

Type: ssh -i "path/to/your/keypair.pem" admin@

and paste the Public DNS (IPv4) you just copied after the @ symbol in the command prompt.

If prompted to confirm connection, enter yes.

Step 8-4

On the AWS portal's Instances page, select the master instance.

At the bottom of the screen, you will notice a Description panel is shown.

Copy or write down the Private IPs address from the Description panel.

Step 8-5

Go back to the command prompt window.

When asked, enter teracv at the Role (Master) prompt.

Paste the Private IP address you copied earlier at the Enter master address prompt.

It will take a few minutes for EmpMonitor to set up the teracv node.

Step 8-6

From your AWS portal's Instances page, copy the Public DNS (IPv4) address for the elastic instance

Step 8-7

Launch an SSH session. If you are on Windows, you can use a tool like Putty or something similar. Make sure you have administrative access.

Type: ssh -i "path/to/your/keypair.pem" admin@

and paste the Public DNS (IPv4) you just copied after the @ symbol in the command prompt.

If prompted to confirm connection, enter yes.

Step 8-8

On the AWS portal's Instances page, select the **master** instance.

At the bottom of the screen, you will see a Description panel.

Copy or write down the Private IPs address from the Description panel.

Step 8-9

Go back to the command prompt window.

When asked, enter elastic at the Role (Master) prompt.

Paste the Private IP address you copied earlier at the Enter master address prompt.

It will take a few minutes for EmpMonitor to set up the elastic node.

Step 8-10

We will now need to approve the two nodes on the EmpMonitor Dashboard.

Login to your EmpMonitor Dashboard on the master server.

Tap the **Cog Wheel** icon near the top-right corner of the dashboard.

Select **Settings**.

Step 8-11

In the Settings screen, click the **Server management** tab.

In the Nodes section, you will see the Screen mining node and the Screen mining database node.

Approve both the nodes by clicking on the **APPROVE** buttons.

You are now all set up for the OCR.

Architecture

Prerequisites

To get started, you will need:
- An AWS account
- Your EmpMonitor license key, available from EmpMonitor Self-hosted portal at:
- An SSH client like Putty if you are using Windows

Step by Step Instructions

1. Creating an AWS Instance and Accessing it with SSH

Step 1-1

Visit: https://www.teramind.co/deployment/aws and click on the **Check out EmpMonitor on Amazon Marketplace** button, which will redirect you to EmpMonitor's deployment page on AWS.

Once you are there, click on the **Continue to Subscribe** button at the top right corner of the screen.

Step 1-2

Once you reach 'Subscribe to this software' screen, click on the **Accept Terms** button.

Step 1-3

Click on the **Continue to Configuration** button.

Step 1-4

In the 'Configure this software' screen, you can select the Fulfilment Option, Software Version, and the Region. Make sure to choose a region that is closest to you for better performance.

Click the **Continue to Launch** button when done.

Step 1-5

In the 'Launch this software' page, tap on the Choose Action pull-down menu and select Launch through EC2.

Click the **Launch** button.

Step 1-6

On the 'Choose an Instance Type' screen, select an instance type from the list of available instances. EmpMonitor recommends at least an m4.xlarge for a typical deployment of up to 300 users.

Once done, click on the **Review and Launch** button, which you'll find at the bottom of the page.