Server Fleet Management at Scale

AWS Implementation Guide

Rob Barnes

June 2018

Last updated: December 2019 (see revisions)

Contents

About This Guide

This implementation guide discusses architectural considerations and configuration steps for deploying the Server Fleet Management at Scale solution on the Amazon Web Services (AWS) Cloud. It includes links to an AWS CloudFormation template that launches, configures, and runs the AWS services required to deploy this solution using AWS best practices for security and availability.

The guide is intended for IT infrastructure architects, administrators, operations engineers, and DevOps professionals who have practical experience architecting on the AWS Cloud.

Overview

Amazon Web Services (AWS) customers who own a fleet of servers are sometimes unsure of how to best automate their fleet management for operational efficiency and maintenance. AWS Systems Manager provides a unified user interface so customers can view operational data from multiple AWS services and allows customers to automate operational tasks across their AWS resources. With Systems Manager, customers can maintain a consistent configuration of their Amazon Elastic Compute Cloud (Amazon EC2) or on-premises instances. They can also automate maintenance and deployment tasks, or automatically apply patches, updates, and configuration changes across any resource group.

To help customers more easily leverage the capabilities of Systems Manager, AWS offers the Server Fleet Management at Scale solution. This solution combines Systems Manager with Amazon Inspector, an automated security assessment service, to help simplify software inventory management, OS patch compliance, and security vulnerability assessments on managed instances. The solution is easy-to-deploy, and automatically provisions the services necessary to automate server fleet management.

Cost

You are responsible for the cost of the AWS services used while running this reference deployment. As of the date of publication, the cost for running this solution with default settings in the US East (N. Virginia) Region for 100 Amazon EC2 instances, and daily Amazon Inspector assessments is approximately **$562.50 per month**. This pricing does not include variable charges incurred from Amazon EC2 instances, Amazon Simple Storage Service (Amazon S3), AWS Lambda, or Amazon CloudWatch.

Prices are subject to change. For full details, see the pricing webpage for each AWS service you will be using in this solution.

Architecture Overview

Deploying this solution builds the following environment in the AWS Cloud:

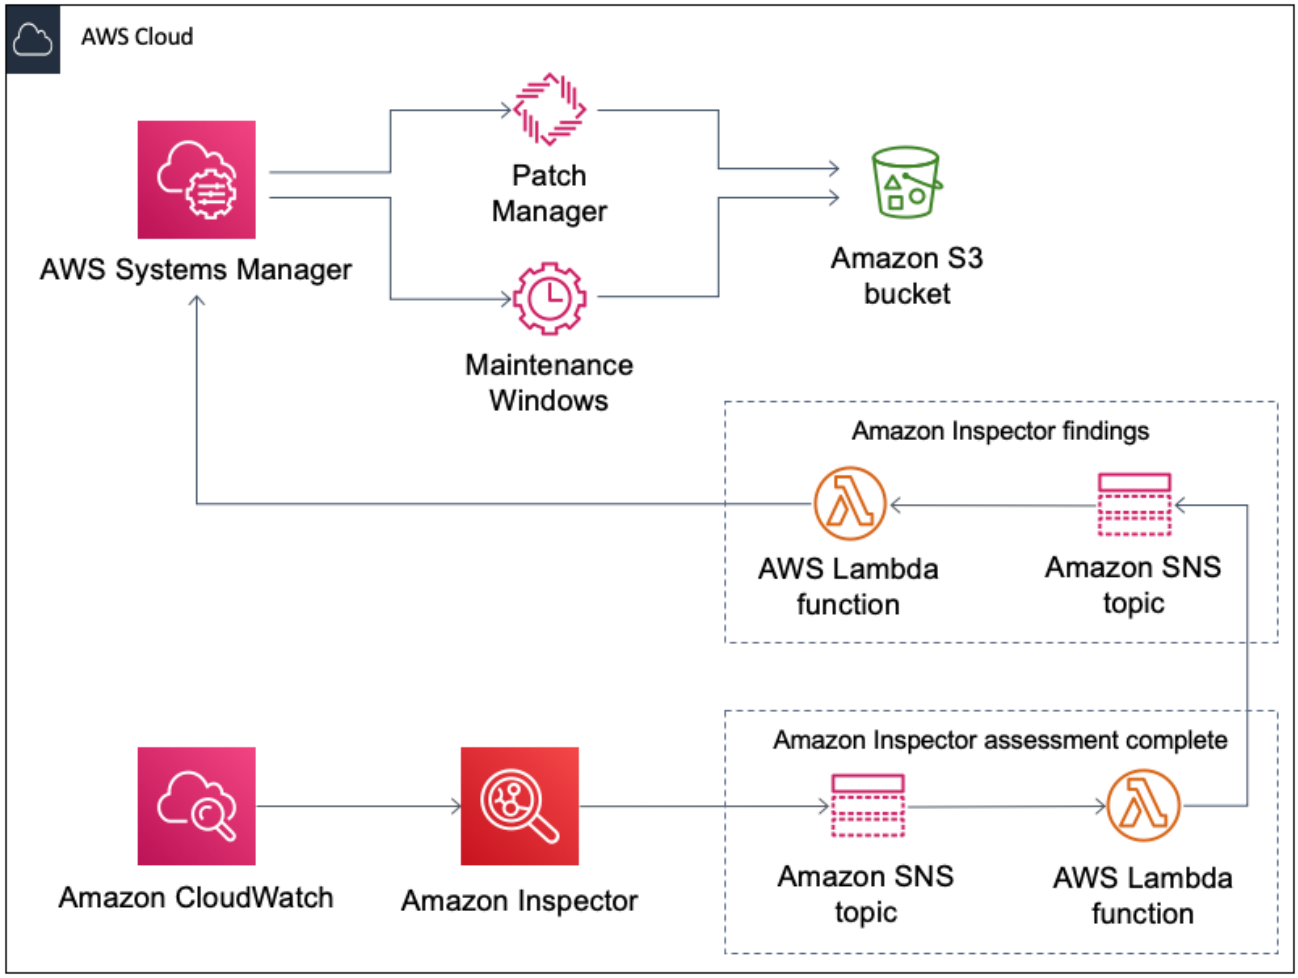

Figure 1: Server Fleet Management at Scale architecture

The AWS CloudFormation template deploys AWS Systems Manager, Amazon Inspector, an Amazon Simple Storage Service (Amazon S3) bucket, an AWS Key Management Service (KMS) key, an AWS Identity and Access Management (IAM) role, an Amazon CloudWatch event, an AWS Lambda function, and an Amazon Simple Notification Service (Amazon SNS) topic.

Systems Manager specifies patch compliance thresholds, defines the schedule for when patching tasks should be run, and defines the Systems Manager associations used to periodically ensure that servers remain in compliance with established configurations. Systems Manager artifacts, including patching and server execution histories and inventories, are stored in the Amazon S3 bucket and encrypted with an AWS KMS key.

A CloudWatch event triggers Amazon Inspector to run daily security assessments on your fleet of Amazon Elastic Compute Cloud (Amazon EC2) instances. Amazon Inspector defines the rules packages for assessments and identifies the target Amazon EC2 instances for assessment runs. When the assessment is complete, Amazon Inspector publishes a message to an Amazon SNS topic that has two subscribers; an AWS Lambda function, and the provided email address. The function then queries Amazon Inspector for the agent IDs of the agents within the assessment run, and sends a message for each agent ID to a second Amazon SNS topic. A second Lambda function receives a notification for each agent ID and queries Amazon Inspector for the findings for each agent, sorts them by vulnerabilities, and updates the Systems Manager Inventory data for the instance under management. Note that the maximum number of agents that can be included in the assessment target of an assessment run is 500.

This solution is designed to allow you to use your own server fleet, but it also includes a sample server fleet you can deploy for testing purposes. For more information, see Appendix A.

Deployment Considerations

Prerequisite: Server Fleet Configuration

To use your existing server fleet with this solution, you must complete the following prerequisite tasks:

- Create an instance profile role with the required AWS Systems Manager permissions
- Verify that your Amazon Elastic Compute Cloud (Amazon EC2) instances meet Systems Manager requirements
- Install AWS Systems Manager Agent (SSM Agent)
- Create `Environment` and `Patch Group` tags with the **Managed Instances Tag Value** that you specified during initial configuration. The `Environment` tag runs the Amazon Inspector assessments and runs Systems Manager inventory daily. The `Patch Group` tag by default is run weekly and directs Systems Manager OS patching for the instances.

Systems Manager uses the tag key to identify applicable Amazon EC2 instances. When an appropriately tagged instance is launched, Systems Manager performs the following tasks:

- Installs or updates the Amazon Inspector agent
- Creates the Systems Manager associations between the servers and Systems Manager documents to ensure the servers are continuously evaluated against defined baselines

- Adds the servers to a patch management regimen to ensure the servers are patched regularly
- Adds the servers to a group that Amazon Inspector will use to regularly run vulnerability assessments

Note that this solution is designed to work with the latest version of the SSM Agent. For more information on configuring and installing the SSM Agent on your instances, see Working with SSM Agent.

Regional Deployments

This solution uses Amazon Inspector which is available in specific AWS Regions only. If you deploy this solution in a region that does not support this service, the Amazon Inspector resources will not be deployed.

While AWS Systems Manger is available in the AWS GovCloud (US) Region, some solution features are not available in that region. We recommend that you deploy this solution in regions that support all solution features.

Solution Components

Systems Manager Associations

The `ManageInspectorAgent` association runs weekly to ensure that the Inspector agent is installed on the targeted managed instances.

The `GatherSoftwareInventory` association runs daily to gather the software inventory of the targeted managed instances. You can view a list of the managed instance's application in the Managed Instance Console **Inventory** tab.

To verify that your Amazon Elastic Compute Cloud (Amazon EC2) instances meet AWS Systems Manager requirements and currently supported Operating Systems, review the prerequisites listed here.

Maintenance Window

A maintenance window allows you to define tasks that will be run against a set of instances on a given schedule. This gives you flexibility and control for how you perform routine tasks. The solution's created maintenance window is scheduled to run weekly in a two-hour window, contains a Run Command task that uses the document `AWS-RunPatchBaseline` to perform patching, and updates the targets defined by the `Patch Group` tag key supplied in the **Managed Instances Tag Value** parameter.

To verify that your Amazon EC2 instances are supported by AWS Systems Manager Patch Manager, review the Operating Systems Supported by Patch Manager listed here.

Amazon Inspector Rules Packages

Amazon Inspector compares the behavior and the security configuration of the assessment targets to selected security *rules packages*. Currently, this solution uses the following rules packages:

- Common Vulnerabilities and Exposures
- Security Best Practices
- Network Reachability
- Center for Internet Security (CIS) Benchmarks

To verify that your Amazon EC2 instances meet Amazon Inspector requirements, review the Amazon Inspector Supported Operating Systems and Regions listed here.

AWS CloudFormation Template

This solution uses AWS CloudFormation to automate the deployment of the Server Fleet Management at Scale solution. It includes the following AWS CloudFormation template, which you can download before deployment:

View template

server-fleet-management-at-scale.template: Use this template to launch the solution and all associated components. The default configuration deploys AWS Systems Manager, Amazon Inspector, an Amazon Simple Storage Service (Amazon S3) bucket, an AWS Key Management Service (AWS KMS) key, an AWS Identity and Access Management (IAM) role, an Amazon CloudWatch event, an AWS Lambda function, and an Amazon Simple Notification Service (Amazon SNS) topic, but you can also customize the template based on your specific needs.

> **Important:** If you delete the solution stack, all of the resources created by the AWS CloudFormation template will be deleted, except the Resource Sync Amazon S3 bucket. You must manually delete the bucket.

Automated Deployment

Before you launch the automated deployment, please review the architecture, configuration, prerequisites, and other considerations discussed in this guide. Follow the step-by-step

instructions in this section to configure and deploy the Server Fleet Management at Scale solution into your account.

Time to deploy: Approximately four minutes

Prerequisites

To use your existing server fleet with this solution, you must complete the prerequisite tasks.

Launch the Stack

This automated AWS CloudFormation template deploys the Server Fleet Management at Scale solution.

> **Note:** You are responsible for the cost of the AWS services used while running this solution. See the Cost section for more details. For full details, see the pricing webpage for each AWS service you will be using in this solution.

1. Sign in to the AWS Management Console and click the button to the right to launch the `server-fleet-management-at-scale` AWS CloudFormation template.

   Launch Solution

   You can also download the template as a starting point for your own implementation.
2. The template is launched in the US East (N. Virginia) Region by default. To launch this solution in a different AWS region, use the region selector in the console navigation bar.

> **Note**: This solution uses the Amazon Inspector service, which is currently available in specific AWS Regions only.[1] If you launch this solution where the service is not available, the solution will deploy without the Amazon Inspector features.

3. On the **Specify Details** page, assign a name to your solution stack.
4. Under **Parameters**, review the parameters for the template and modify them as necessary. This solution uses the following default values.

Parameter	Default	Description
Managed Instances Tag Value	Sample Fleet	This tag value identifies applicable Amazon EC2 instances. When combined with the `Patch Group` tag name, AWS Systems Manager will patch the instance. When combined with the `Environment` tag name, Systems Manager will perform inventory scans. If Amazon Inspector is available in the deployed region, it will perform daily assessments.

[1] For the most current service availability by AWS Region, see https://aws.amazon.com/about-aws/global-infrastructure/regional-product-services/

Parameter	Default	Description
Create a Sample Fleet	Yes	Choose whether to launch the sample server fleet
Amazon SNS Email	*<Requires input>*	Email address that Amazon Inspector will use to send Amazon SNS notifications of assessment run completions **Note:** You must acknowledge the confirmation email from Amazon SNS to receive notifications.

5. Choose **Next.**
6. On the **Options** page , choose **Next**.
7. On the **Review** page, review and confirm the settings. Be sure to check the box acknowledging that the template will create AWS Identity and Access Management (IAM) resources.
8. Choose **Create** to deploy the stack.

   You can view the status of the stack in the AWS CloudFormation Console in the **Status** column. You should see a status of CREATE_COMPLETE in approximately four minutes.

Security

When you build systems on AWS infrastructure, security responsibilities are shared between you and AWS. This shared model can reduce your operational burden as AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. For more information about security on AWS, visit the AWS Security Center.

Amazon S3 Bucket Encryption

The Amazon Simple Storage Service (Amazon S3) bucket created by this solution requires objects to be encrypted prior to being stored in the bucket. The objects must be encrypted by the AWS Key Management Service (AWS KMS) encryption key this solution creates.

Amazon Inspector Findings

When Amazon Inspector runs assessments on your instances, it may produce **critical**, **high**, **important**, or **informational** findings. We recommend that you review these findings with your organization's security team and remediate the findings according to your organization's security policies.

IAM Service Roles

AWS Identity and Access Management (IAM) roles enable customers to assign granular access policies and permissions to services and users on the AWS Cloud. This solution creates optional IAM roles with least-privilege access that contain the permissions needed to perform the tasks required by their respective functions. We recommend that you review the role policies and further restrict them as needed once the deployment is up and running.

Security Group

If you choose to deploy the sample server fleet, this solution creates a security group that is designed to control and isolate network traffic for the sample instances. We recommend that you review the security group and further restrict access as needed once the deployment is up and running.

Additional Resources

AWS Services

- AWS Systems Manager
- Amazon Inspector
- AWS CloudFormation
- AWS Lambda
- Amazon Elastic Compute Cloud
- AWS Identity and Access Management
- Amazon Simple Storage Service
- Amazon Simple Notification Service

Appendix A: Sample Server Fleet

The Server Fleet Management at Scale solution includes a sample server fleet for testing purposes. Deploying this solution with the sample fleet builds the following environment on the AWS Cloud.

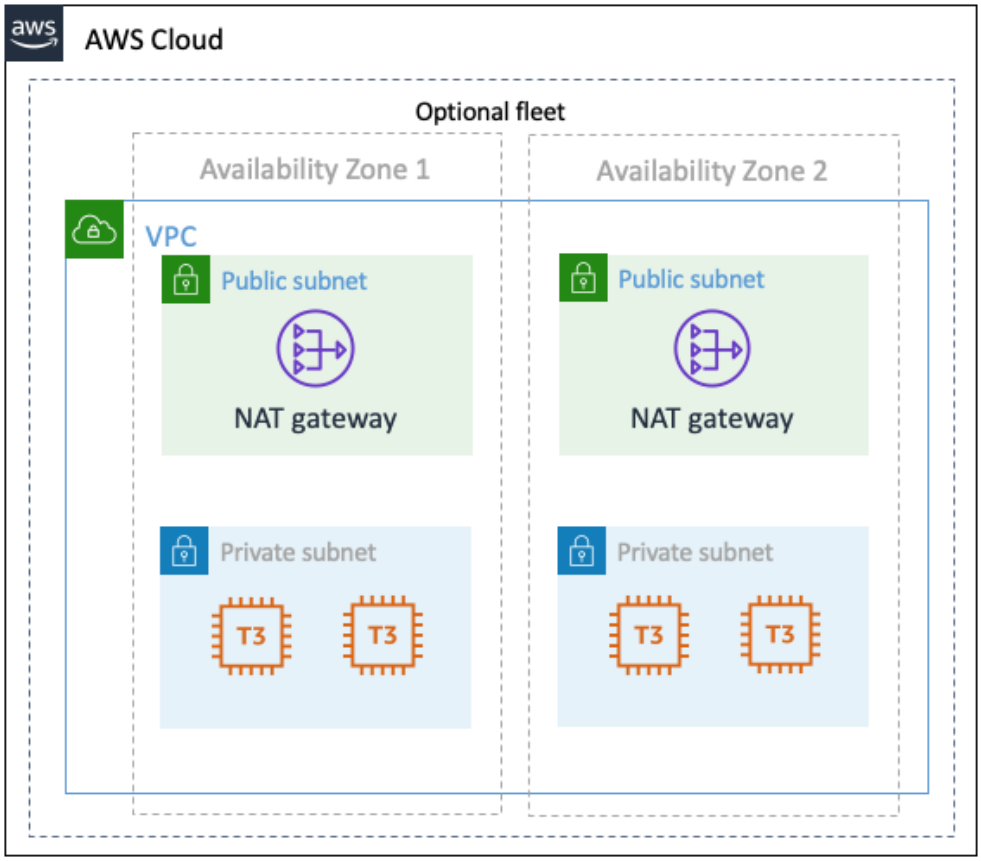

Figure 2. Server Fleet Management at Scale with sample fleet

If you choose to deploy the sample servers, the template launches an Amazon Virtual Private Cloud (Amazon VPC) network topology with two public and two private subnets. Four Amazon Elastic Compute Cloud (Amazon EC2) t3.large instances are deployed in the private subnet, and access the internet through a NAT gateway in the public subnet.

A security group restricts outbound network access to port 443 and port 80, and an AWS Identity and Access Management (IAM) role allows the instances to interact with AWS Systems Manager to receive commands.

Managing Your Fleet

Once the AWS CloudFormation template is deployed, you can test the solution with this sample scenario for managing your fleet of instances.

View Managed Instances

1. Sign in to the AWS Management Console, navigate to the **AWS Systems Manager** service, and select **Managed Instances.**
2. Select any **Instance ID**, and verify in the **Associations** tab that AWS Systems Manager documents `ManageInspectorAgent` and `GatherSoftwareInventory` have been applied to your instances.

Run Amazon Inspector

This solution creates a daily Amazon Inspector schedule that runs an assessment against a target of specially tagged instances using all available rules packages. Use this procedure to run the assessment manually.

1. In the console, navigate to the **Amazon Inspector Console**.
2. Select **Assessment templates**, and select the **Sample Fleet**.
3. Click **Run**.
4. Select the **Assessment runs** link on the left to view the progress of the assessment.

   The assessment is configured to run for 15 minutes, when it is completed an Amazon Simple Notification Service (Amazon SNS) notification will be sent to subscribed users with the Systems Manager inventory of findings for each affected instance.
5. After the assessment run completes, navigate to the **AWS Systems Manager** console, and select **Managed Instances**.
6. Select any **Instance ID**, and select the **Inventory** tab.
7. To view Amazon Inspector agent findings, in the **Inventory type** drop-down, select **Custom:InspectorFindings**.

   Note that you can extend the inventory of managed instances and create a custom inventory for the Amazon Inspector findings.

Remediate the Amazon Inspector Findings

If Amazon Inspector has any findings they will be remediated in the next maintenance window. However, you can manually update the maintenance window to remediate the findings before the next window.

1. In the AWS Systems Manager console, under **Actions**, select **Maintenance Windows**.

2. Select the **fleet-wide-weekly-patching** maintenance window.
3. Select **Edit**, navigate to the **Schedule** pane.
4. Select **Cron schedule builder**.
5. In the **Window starts**, select the **Every Day** radio button, and enter a **time** for the near future, and select **Save Changes**.

This updates the maintenance window at the time you specify. Note that the time should be represented in coordinated universal time (UTC) and in 24-hour format, so you'll need to adjust this time according to your time zone. For example, if you want to update the maintenance window at 12:30 p.m. ET, enter the time as 17:30.

6. Select the linked **Window ID**, and select the **History** tab to verify that the task was triggered. The task should show as **In Progress**.

The time it takes to complete the task will depend on the number of instances that are being patched and the size of the patches applied. For the sample fleet, the task can take up to 30 minutes to complete.

7. Navigate to the **Amazon Inspector** console to verify the applied patches were run on the targeted instances.
8. Select **Assessment templates**, and select the **Sample Fleet**.
9. Click **Run**.
10. Select the **Assessment runs** link on the left to view the progress of the assessment.
11. After the assessment run completes, navigate to the **AWS Systems Manager** console, and select **Managed Instances**.
12. Select any **Instance ID**, and select the **Inventory** tab.
13. In the **Inventory type** drop-down, select **Custom:InspectorFindings**.

The number of high severity vulnerabilities should be fewer than the initial Patch run.

Appendix B: Collection of Operational Metrics

This solution includes an option to send operational metrics to AWS. We use this data to better understand how customers use this solution and related services and products. When enabled and Amazon Inspector is deployed, the following information is collected and sent to AWS:

- **Solution ID:** The AWS solution identifier
- **Unique ID (UUID):** Randomly generated, unique identifier for each solution deployment
- **Timestamp:** Data-collection timestamp
- **Managed Instance Count:** The number of Amazon Inspector Agents within the Assessment Run

Note that AWS will own the data gathered via this survey. Data collection will be subject to the AWS Privacy Policy. To opt out of this feature, complete one of the following tasks:

a) Modify the AWS CloudFormation template mapping section as follows:

```
Send:
     AnonymousUsage:
        Data:”Yes”
```

to

```
Send:
     AnonymousUsage:
        Data: “No”
```

OR

b) After the solution has been launched, find the `SFM-RespondToInspectorCompleteFunction` function in the Lambda console and set the `send_anonymous_data` environment variable to `No`.

Send Us Feedback

We welcome your questions and comments. Please post your feedback on the AWS Solutions Discussion Forum.

You can visit our GitHub repository to download the templates and scripts for this solution, and to share your customizations with others.

Document Revisions

Date	Change
June 2018	Initial release
December 2019	Updated Regional Deployment and Amazon Inspector Rules Packages information; upgraded the solution to Python 3.8