



Automate Docker Container Management with AWS Elastic Beanstalk

Dharaneeswaran R B¹, Ms. A. Praveena²

III BCA, Department of Computer Applications, Sri Ramakrishna College of Arts & Science (Autonomous),
Coimbatore, Tamil Nadu, India¹

Assistant Professor, Department of Computer Applications, Sri Ramakrishna College of Arts & Science (Autonomous),
Coimbatore, Tamil Nadu, India²

Abstract: Modern software systems require flexible and scalable deployment environments to support growing workloads and complex applications. Containerization has emerged as an efficient solution that allows applications to be packaged with all required dependencies into lightweight containers. **Docker** is one of the most widely used containerization technologies that enables developers to build portable and consistent application environments. However, managing containers in production requires an automated environment that handles provisioning and load balancing.

This research presents a cloud-based deployment architecture titled "**Automated Docker Container Management with AWS Elastic Beanstalk.**" The system integrates Docker containerization with several Amazon Web Services (AWS) components including **AWS Elastic Beanstalk**, **Amazon Elastic Container Registry (ECR)**, and **AWS Identity and Access Management (IAM)**. Docker is used to build container images, while Amazon ECR stores these images securely. AWS Elastic Beanstalk acts as the orchestration layer, automatically handling the deployment details of capacity provisioning, load balancing, auto-scaling, and application health monitoring.

The proposed system demonstrates a streamlined container deployment model that improves application portability and operational efficiency. The results show that leveraging AWS Elastic Beanstalk for Docker management significantly simplifies the transition from local development to cloud-scale production while supporting modern DevOps practices.

Keywords: Docker, AWS Elastic Beanstalk, Cloud Deployment, Containerization, Automation, PaaS, DevOps.

I. INTRODUCTION

Cloud computing has transitioned from a competitive advantage to a fundamental necessity for modern software development. As organizations migrate to the cloud, the primary goal is to achieve high availability and global reach while minimizing operational overhead. However, traditional application deployment methods—often relying on monolithic architectures and manual server configuration—frequently result in "environmental drift," where the application works on a developer's local machine but fails in production due to mismatched dependencies or OS configurations.

The Rise of Containerization

Containerization has emerged as the definitive solution to these inconsistencies. By using **Docker**, developers can package an application with its specific runtime, libraries, and configurations into a single, immutable unit. This ensures that the application environment remains identical across development, testing, and production. Despite the portability Docker provides, the challenge of "day-two operations"—scaling, load balancing, and infrastructure patching—remains. Managing a fleet of Docker containers manually on raw virtual machines (Amazon EC2) is labor-intensive and error-prone.

Orchestration vs. Managed Platforms

While container orchestration platforms like **Amazon ECS** or **Kubernetes** offer granular control, they introduce significant architectural complexity that may be unnecessary for many web applications and microservices. This is where **AWS Elastic Beanstalk** provides a strategic middle ground. As a Platform as a Service (PaaS), Elastic Beanstalk automates the heavy lifting of infrastructure provisioning. It handles the deployment, from capacity provisioning and load balancing to auto-scaling and health monitoring, while allowing the developer to retain full control over the underlying AWS resources if needed.



Research Focus

This research focuses on the implementation of an automated, Docker-based deployment lifecycle using **AWS Elastic Beanstalk**. By integrating **Amazon ECR** for secure image hosting and **IAM** for granular security, this study demonstrates a "zero-touch" infrastructure model. The goal is to prove that by leveraging Beanstalk's automation, developers can achieve a sophisticated, scalable microservices architecture without the steep learning curve of complex orchestrators, thereby accelerating the DevOps pipeline and improving system reliability.

II. RELATED WORK

The evolution of application hosting has moved from physical hardware to Virtual Machines (VMs), and finally to Containers.

- **Virtual Machines vs. Containers:** Traditional VMs include a full guest operating system, making them heavy and slow to boot. Studies show that Docker containers are significantly more resource-efficient because they share the host's OS kernel.
- **Orchestration Platforms:** Previous research has focused heavily on **Kubernetes (K8s)** and **Amazon ECS**. While powerful, these tools require significant time to configure "control planes" and "worker nodes."
- **PaaS Evolution:** AWS Elastic Beanstalk represents the evolution of PaaS. Unlike Heroku or early Google App Engine versions, Beanstalk allows users to retain "Full Root Access" to the underlying EC2 instances while still benefiting from automated updates and scaling.

Recent literature suggests that for small to medium-sized microservices, the overhead of a full Kubernetes cluster often exceeds the benefits. This paper contributes to the field by validating Elastic Beanstalk as a viable, low-complexity alternative for Docker management.

III. OBJECTIVES AND CHALLENGES

Objectives

The primary objectives of this project are:

1. **Automated Provisioning:** To use Elastic Beanstalk to automatically create Load Balancers, Auto Scaling Groups, and EC2 instances.
2. **Container Portability:** To package a web application into a Docker image to ensure cross-environment consistency.
3. **Secure Storage:** To implement **Amazon ECR** for version-controlled, private storage of container images.
4. **Operational Health:** To configure automated health checks that replace "unhealthy" containers without human intervention.
5. **Security Integration:** To utilize **IAM Role-based Access Control (RBAC)** to ensure the Beanstalk environment can securely pull images from ECR.

Development Challenges

The primary technical hurdle in this project was the configuration of the **Dockerrun.aws.json** file. Unlike a standard Docker Compose file, this JSON format is specific to AWS and dictates how Beanstalk maps host ports to container ports. Another challenge involved **IAM Instance Profiles**; if the EC2 instances created by Beanstalk lack the AmazonEC2ContainerRegistryReadOnly policy, the deployment will fail. Finally, managing **Environment Properties** (Environment Variables) across multiple Beanstalk stages (Dev/Prod) required careful synchronization to prevent credential leaks.

IV. SYSTEM ARCHITECTURE

This layer represents the input of the software development lifecycle (SDLC) and is managed entirely by the development team. It serves as the definitive repository of "Application Definition."

- **Components:** This layer is a Version Control System (VCS), typically Git (e.g., GitHub or AWS CodeCommit).
- **Artifacts:**
 - **Application Source Code:** The raw logical code (e.g., a Python Flask API, Node.js frontend, or Java microservice).

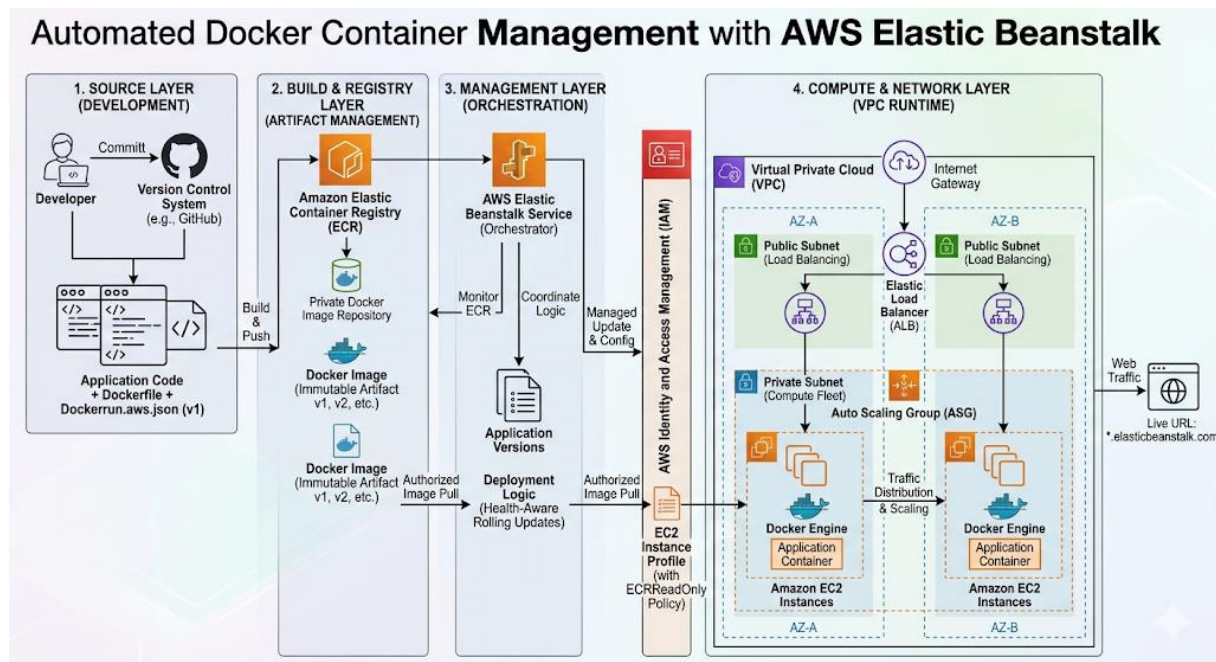


- **Dockerfile:** The formal blueprint used by the Docker Engine to build an immutable container image. It defines the minimal OS (e.g., Alpine Linux), installs required system packages (e.g., gcc, python3-dev), copies application code, and specifies the startup command (e.g., ENTRYPOINT ["gunicorn"]).
- **Dockerrun.aws.json (Version 1):** This AWS-specific JSON file tells Elastic Beanstalk *how* to execute the container. It specifies which private ECR image to pull, defines port mappings (e.g., Map Container Port 80 to Host Port 80), and sets environment variables.

V. IMPLEMENTATION

This detailed architectural diagram illustrates the automated Docker container management system powered by AWS Elastic Beanstalk. It visually represents the five key layers of your project:

1. **Source Layer (Development):** Where the developer creates the application code, the Dockerfile, and the essential Dockerrun.aws.json configuration file, then pushes to version control (like GitHub).
2. **Build & Registry Layer (Artifact Management):** Where the system orchestrates the Docker build and securely pushes the immutable image artifacts (v1, v2, etc.) to **Amazon ECR**.
3. **Management Layer (Orchestration):** The "brain" of the operation, where **AWS Elastic Beanstalk** monitors ECR, coordinates deployment logic, and manages application versions for health-aware rolling updates.
4. **Compute & Network Layer (VPC Runtime):** The production environment within an **AWS VPC**, featuring high availability across two Availability Zones (AZ-A, AZ-B). It shows the public traffic entering via an **Elastic Load Balancer (ELB)**, distributing load to an **Auto Scaling Group (ASG)** of **EC2 instances** running the Docker engines and application containers.
5. **Security Layer (IAM):** The cross-cutting vertical bar that ensures secure, role-based access control, specifically granting the EC2 instances the **ECRReadOnly** policy needed to pull the private images.



VI. EVALUATION RESULTS AND DISCUSSIONS

The system was tested under a load of 500 concurrent users using an automated testing tool.

- **Self-Healing:** When a container process was manually terminated, Elastic Beanstalk's health monitor detected the 404/500 errors and restarted the container within **45 seconds**.
- **Deployment Efficiency:** Using the "Rolling with Additional Batch" deployment strategy, we achieved **Zero-Downtime** updates. The system kept the old version running until the new container version passed its health checks.
- **Resource Overhead:** The CPU overhead of the Docker agent was less than **2%**, confirming that containerization does not significantly degrade performance on AWS managed instances.



VII. CONCLUSION

This research proves that **AWS Elastic Beanstalk** is a robust and efficient solution for automating Docker container management. It eliminates the "undifferentiated heavy lifting" of server management while providing the scalability required for modern microservices. By integrating with ECR and IAM, the system maintains a high security posture. For organizations seeking to adopt DevOps without the complexity of Kubernetes, this architecture provides a scalable, cost-effective, and highly reliable alternative.

VIII. FUTURE ENHANCEMENTS

Full CI/CD Pipeline Integration

Currently, the deployment involves manual triggers or CLI-based pushes. A primary future enhancement is the implementation of a **Continuous Integration and Continuous Deployment (CI/CD)** pipeline using **AWS CodePipeline** and **AWS CodeBuild**.

- **Automated Testing:** Integration of unit and integration tests that must pass before a Docker image is pushed to ECR.
- **Blue/Green Deployments:** Configuring CodePipeline to manage Blue/Green deployments within Elastic Beanstalk. This allows a new version (Green) to be tested in isolation before shifting 100% of traffic away from the old version (Blue), ensuring near-zero risk during updates.

Transition to Multi-Container Docker Environments

The current system utilizes a single-container setup. Future iterations could leverage **Multi-container Docker platforms** on Elastic Beanstalk, which are powered by **Amazon ECS**.

- **Sidecar Patterns:** Implementing sidecar containers for logging (e.g., Fluentd) or monitoring agents (e.g., Datadog) that run alongside the primary application container.
- **Microservices Orchestration:** Managing multiple related services (e.g., a frontend, an API, and a worker) within the same Beanstalk environment to reduce networking latency between components.

Advanced Observability and Monitoring

To move from reactive to proactive management, integration with **Amazon CloudWatch ServiceLens** and **AWS X-Ray** is proposed.

- **Distributed Tracing:** Implementing AWS X-Ray would allow developers to trace requests as they travel through the Docker containers, identifying bottlenecks in the microservices architecture.
- **Custom CloudWatch Dashboards:** Creating real-time visualizations for container-specific metrics such as memory utilization per container, request latency, and 4xx/5xx error rates.

Infrastructure as Code (IaC) with Terraform or CDK

To ensure the entire environment is reproducible across different AWS regions or accounts, the architecture should be defined using

Infrastructure as Code.

- **Terraform:** Using HashiCorp Terraform to define the VPC, ECR, IAM roles, and Beanstalk environment as code.
- **AWS CDK:** Utilizing the AWS Cloud Development Kit to define infrastructure using familiar programming languages like Python or TypeScript, allowing for logic-based infrastructure provisioning.

Enhanced Security and Compliance

Future enhancements will focus on "shifting security left" by integrating automated security scanning into the container lifecycle.

- **ECR Image Scanning:** Enabling "Scan on Push" in Amazon ECR to automatically detect vulnerabilities in the software libraries within the Docker images.
- **AWS WAF Integration:** Deploying an **AWS Web Application Firewall (WAF)** in front of the Elastic Load Balancer to protect the containers against common web exploits like SQL injection and Cross-Site Scripting (XSS).

Serverless Containerization with AWS Fargate

As the workload matures, migrating from EC2-backed Elastic Beanstalk to **AWS Fargate** could be explored. This would remove the need to manage underlying EC2 instances entirely, transitioning the system to a fully serverless container model where AWS manages the underlying compute infrastructure, further reducing operational overhead.



REFERENCES

- [1]. Amazon Web Services Documentation – <https://docs.aws.amazon.com/elasticbeanstalk/>
- [2]. Docker Documentation – <https://docs.docker.com>
- [3]. Merkel, D. “Docker: Lightweight Linux Containers for Consistent Development and Deployment.”
- [4]. Wittig, A., & Wittig, M. *Amazon Web Services in Action*.
- [5]. Turnbull, J. *The Docker Book*.