



## **NewgenONE OmniDocs RMS**

# **Newgen Enterprise Products Containerization for Azure**

**Version: 4.0 SP1**

**[Newgen Software Technologies Ltd.](#)**

This document contains propriety information of NSTL. No part of this document may be reproduced, stored, copied, or transmitted in any form or by any means of electronic, mechanical, photocopying, or otherwise, without the consent of NSTL.

# Table of contents

|          |   |           |
|----------|---|-----------|
| <b>1</b> | <b>Preface</b>  | <b>2</b>  |
| 1.1      | Revision history  | 2         |
| 1.2      | Intended audience   | 2         |
| 1.3      | Documentation feedback                                      | 2         |
| <b>2</b> | <b>Containerization of Newgen Enterprise Products</b>       | <b>3</b>  |
| 2.1      | Key advantages  | 3         |
| 2.2      | Deliverables – Product docker images for initial deployment | 4         |
| 2.2.1    | ECM Suite - OmniDocs  | 4         |
| 2.3      | Deployment architecture                                     | 5         |
| 2.3.1    | OmniDocs Vanilla deployment architecture                    | 5         |
| 2.3.2    | OmniDocs Multi-Tenant Deployment                            | 7         |
| 2.4      | Used tools and technologies                                 | 8         |
| 2.5      | Kubernetes Autoscaling                                      | 9         |
| 2.6      | Logging architecture  | 10        |
| 2.7      | Container security  | 12        |
| 2.8      | CI/CD Pipeline  | 13        |
| 2.8.1    | CI/CD Pipeline of Base Products                             | 13        |
| 2.8.2    | CI/CD Pipeline of Custom Code                               | 15        |
| 2.8.3    | CI/CD Pipeline of Product’s Hotfix                          | 16        |
| 2.8.4    | Deployment Changes Against File Types                       | 17        |
| 2.8.5    | Deployment Downtime   | 17        |
| 2.8.6    | Docker Image Management for Rollbacks                       | 18        |
|          | <b>Appendix</b>   | <b>18</b> |

# 1 Preface

This guide describes the containerization approach for Newgen Enterprise products, which thus details out the container-based deployment architecture and container-based DevOps pipeline for NewgenONE flagship products OmniDocs & RMS on Microsoft Azure.

## 1.1 Revision history

| Revision Date | Description         |
|---------------|---------------------|
| April 2024    | Initial publication |

## 1.2 Intended audience

This guide is intended for Cloud Administrators, System Administrators, developers, and any other users seeking information about OmniDocs & RMS Container Architecture and DevOps Pipeline to manage deployments. The reader must be comfortable to understand the computer terminology.

## 1.3 Documentation feedback

To provide feedback or any improvement suggestions on technical documentation, you can write an email to [docs.feedback@newgensoft.com](mailto:docs.feedback@newgensoft.com).

To help capture your feedback effectively, request you to share the following information in your email.

- Document Name:
- Version:
- Chapter, Topic, or Section:
- Feedback or Suggestions:

## 2 Containerization of Newgen Enterprise Products

The containerization support is enabled for Newgen flagship product OmniDocs & RMS on **Azure AKS (Azure Kubernetes Service)**. Newgen has isolated the product suite's application modules that must be deployed in a separate docker container to enable independent scalability for each of these components or modules. Newgen product suite's web components and EJB components are isolated for deployment in separate Docker Container Instances. Web components is deployed on the underlying web server JBoss WebServer 5.7.x. EJB components is deployed on the underlying application server JBoss EAP 7.4.x.

### 2.1 Key advantages

The key advantages are as follows:

- **Reduce time and effort in environment configuration:** With Docker containers, you can reduce the time and effort in environment configuration, you don't need to bother about the installation parts. Whenever you want to use these products, pull the docker images and run them along with some configuration. Also, Containers are immutable by design so no need to worry about corrupted VMs.
- **Consistency across the environments** (Dev, UAT, Production, and so on): Docker container provides the environment consistency that means it never face such issues that this application is working on a Dev environment, but it is not working on UAT and stage environment due to O.S updates or patches.
- **Auto Scaling of Containers:** Scaling of containers is far faster and easier than deploying additional VMs. Docker containers are auto-scaled up and scaled down depending on the CPU utilization and memory consumption. Not only containers but worker node instances can also scale up and scale down automatically in Kubernetes.
- **Easy distribution and portability:** As our application and their dependencies all are bundled inside the Docker image then it becomes easy to distribute our applications using the container registry.
- **Maintain the Liveness of applications running inside a container:** Docker containers have healing power, if an application running inside the container gets down due to any reason or becomes unresponsive then Kubernetes auto-restart that application inside the container.
- **Seamless deployment of updates (Rolling Deployment):** Deploying updates as a Docker image are far faster and very efficient. Whenever we push new images to the container registry, Kubernetes 1st of all creates new containers to make it up and running once new containers become fully functional after that Kubernetes start deleting the existing containers. In such a case, the end-user never faces any downtime even the user session also not gets expired, and the latest updates gets deployed.

## 2.2 Deliverables – Product docker images for initial deployment

Newgen has isolated the product suite into multiple Docker containers to enable the independent scalability of each Docker container. This separation is done based on the product's usability. At a broad level, Web components and EJB components are isolated for deployment in separate container Instances. Web components are deployed on the underlying web server JBoss WebServer 5.7.x. EJB components are deployed on the underlying application server JBoss EAP 7.4.x. Newgen releases multiple Docker images for the different product suites along with some configuration files for data persistence, YAML files for deployment, and some documentation for end-to-end configurations and deployments.

### 2.2.1 ECM Suite – OmniDocs & RMS

In the case of ECM Suite – OmniDocs & RMS, Newgen delivers the following Docker images for initial product deployment:

- OmniDocs+RMS Web Components
- OmniDocs Web Service Components
- OmniDocs+RMS EJB Components
- OmniDocs Add-on Services (Wrapper, Alarm Mailer, Scheduler, Thumbnail Manager and LDAP)
- EasySearch (Apache Manifold only).
- Text Extraction Manager or Full-Text Search (TEM or FTS)
- OmniScan Web Components

OmniDocs Services container consists of seven different add-on services of OmniDocs & RMS. If there is a requirement, you can split them into separate containers.

EasySearch container contains the Apache Manifold which is freeware software.

If OmniScan Web components are required for scanning services, then it also constitute a separate Docker instance.

| Docker Image Name                 | Docker Image Size |
|-----------------------------------|-------------------|
| OmniDocs Web                      | 771 MB            |
| OmniDocs Web Service              | 482 MB            |
| OmniDocs EJB                      | 706 MB            |
| OmniDocs Services                 | 1.02 GB           |
| EasySearch (Apache Manifold only) | 695 MB            |
| TEM or FTS                        | 251 MB            |
| OmniScan Web                      | 439 MB            |
| SharePointAdapter                 | 176 MB            |
| OmniDocs Wopi                     | 315 MB            |

## 2.3 Deployment architecture

This section describes the deployment architecture of OmniDocs & RMS Docker Containers.

### 2.3.1 OmniDocs & RMS Vanilla deployment architecture

This section describes the deployment of the vanilla flavor of OmniDocs & RMS.

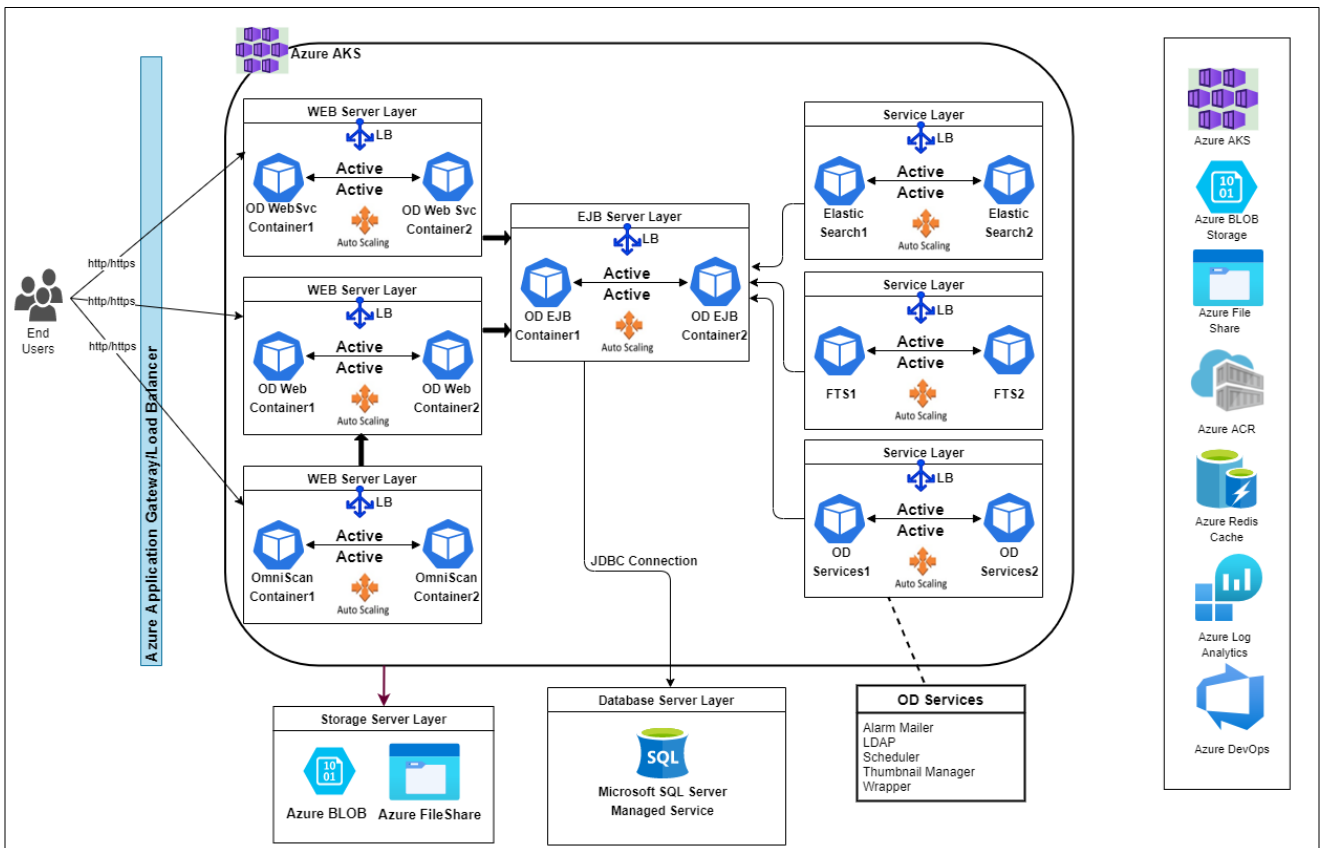


Figure 2.1

- In this architecture, all the user requests come to Azure Application Gateway or Load Balancer, which is created and managed by the Application Ingress controller. An ingress controller is an object running inside the Kubernetes cluster that is used to manage the host-based routing rules.

For example, you can set the host-based routing rules like if the URL is

*omnidocs.newgendocker.com* then the ingress controller redirects the user request to

OmniDocs & RMS WEB containers. If the URL is *omnidocswebservice.newgendocker.com* then the ingress controller redirects the user request to the OmniDocs WEB container.

- Now, there is a separate container for web services that can directly connect with the EJB's present in another container. In case, users want to connect their business or productivity applications with OmniDocs, they can call out these web services directly (which are now deployed in separate container). Web server components and App server components are still deployed in separate containers for accessing OmniDocs through the GUI.
- Here, each Docker image is running in HA mode, one container in each availability zone, and both the containers is running in Active-Active mode so that they can serve the user requests at the same time and manage the load properly. All the containers are also running with auto-scaling enabled so that whenever user requests increase drastically, new pods can auto-scaled up and start serving the user requests.
- Here, a separate Docker container is created for OmniDocs Addon services like ElasticSearch, FTS, Alarm Mailer, and all these services is connected to the OD & RMS EJB container. Additionally, OD & RMS EJB is connected to the database servers.
- In this architecture, for the database layer, the MSSQL works as a managed service.
- Here, Azure FileShare works for config file persistence and log file persistence.
- Azure BLOB storage kept the PN files. PN files are encrypted files that contain all the added, uploaded, or scanned documents by Newgen products.
- On the right panel, there are several Azure services that configures and manage the Docker containers.

For example, Azure AKS (Azure Kubernetes Service) works as a container orchestration, Azure FileShare for the configuration files, and log files persistence. Docker images gets stored to the private container repository like Azure ACR (Azure Container Registry). Azure Cache for Redis persists the user's web session. Azure Monitor and Azure Security Center monitors the container's health. Also, the CI/CD pipeline implements the end-to-end deployment.

## 2.3.2 OmniDocs & RMS Multi-Tenant Deployment

This section describes the deployment of the OmniDocs & RMS Multi-Tenant.

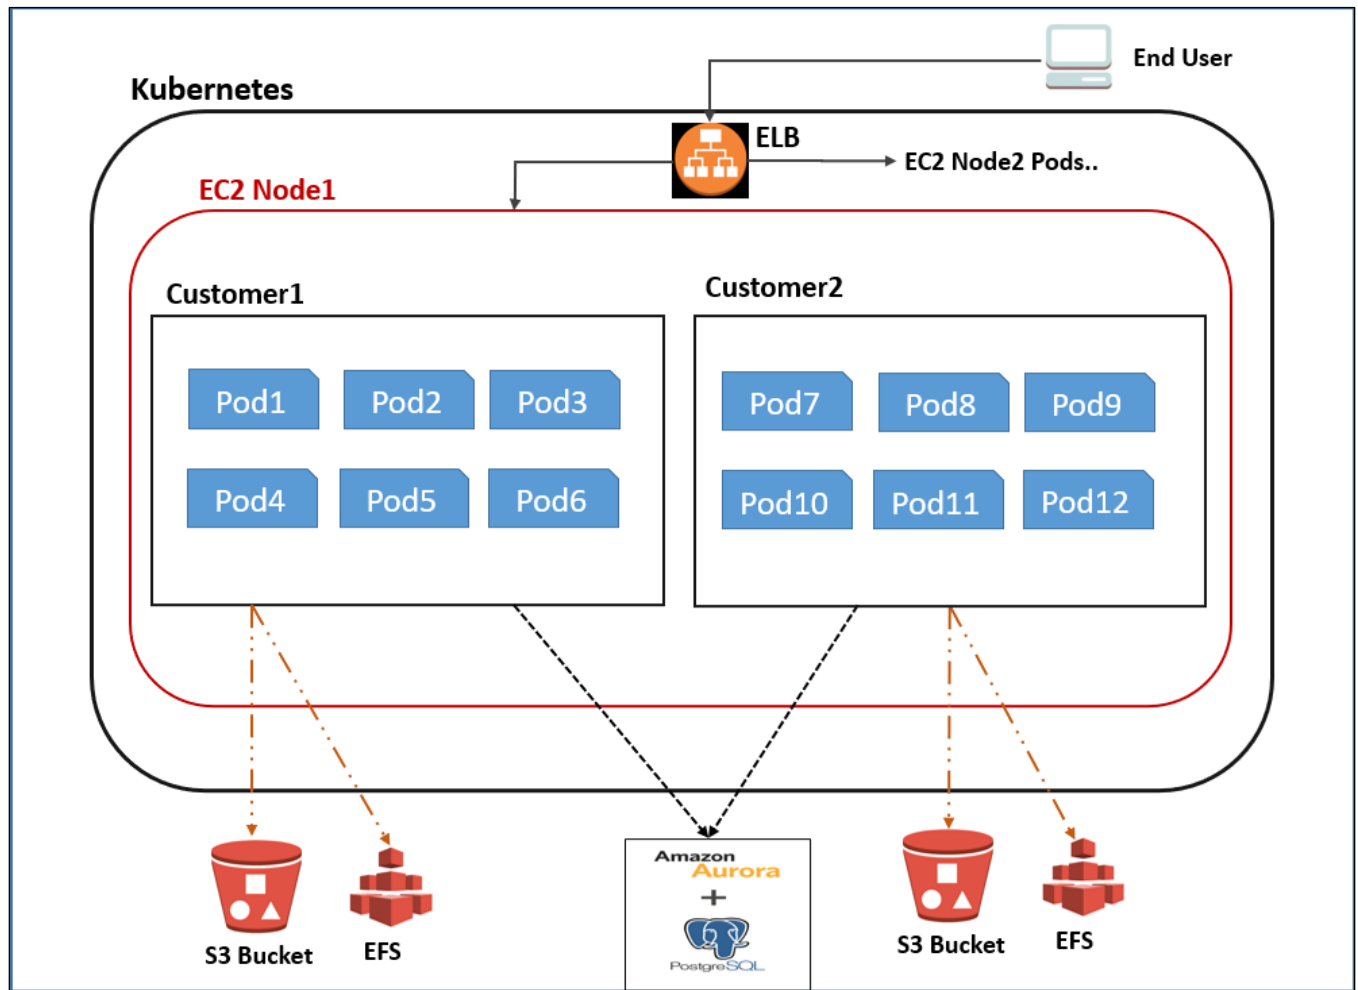


Figure 2.2

- In the case of Multi-Tenant Deployment, containers of multiple customers are running on the same EC2 nodes.  
For example, in the above architecture diagram, POD1 to POD6 are running for customer1 and from POD7 to POD12 are running for customer2. Also, all the containers are running on the same EC2 node.
- In container orchestration, define the resource limit to each container so that one container must not consume all the resources available on the EC2 node due to which other container's performance remains intact.
- In this architecture, end-users' requests come to the Azure Application Gateway or Load Balancer which is managed by the Application Ingress controller in Kubernetes. After that



Ingress controller routes the incoming requests to the target Kubernetes services for different customers according to host-based routing rules. Host-based routing is a capability of ALB redirects the user requests to the right service based on the request-host header.

For example, you can set the rules as described below:

- IF URL is *customer1dev.aws.co.in* THEN redirect to POD1.
- IF URL is *customer2dev.aws.co.in* THEN redirect to the POD7.

Where POD1 and POD7 are OmniDocs Web containers for customer1 and customer2 respectively.

- In this architecture, the same Load Balancer and worker nodes can be used for multiple customers.
- Containers can be different for each customer depending on the changes in custom code.
- Azure BLOB storage and Azure File Share are different for each customer to maintain the data security.
- The database server can be either Microsoft SQL, Oracle, PostgreSQL, or AWS Aurora PostgreSQL. You can use either the same database server for all customers or different database servers for each customer. It depends on the business requirement.

## 2.4 Used tools and technologies

The following tools and technologies are used for this containerization approach:

- **Container Orchestration: Azure Elastic Kubernetes Service**

Azure Kubernetes Service (Azure AKS) is a fully managed Kubernetes service. EKS runs the Kubernetes management infrastructure across multiple Availability Zones, automatically detects and replaces unhealthy control plane nodes, and provides on-demand, zero downtime upgrades, and patching. AKS automatically applies the latest security patches to your cluster control plane.

- **Worker Node Host OS: Ubuntu**

Here, Ubuntu is used as Worker Node host OS. Kubernetes worked node is used to run the containerized applications. Every Kubernetes cluster has at least one worker node.

- **Container OS: Alpine**

Container OS is a lightweight, optimized operating system that is used to run Docker containers and bundled inside the Docker images. It defines the file system structure of a container.

For example, Ubuntu, CentOS, Alpine, Windows, RHEL, and so on. In Newgen, the Alpine is used whose size is about 5 MB only. If there is some special requirement, the Docker images is built with another container OS on-demand basis.

- **Docker Image Registry: Azure Container Registry (ACR)**

Azure Container Registry is used to store the Product's Docker Images and by default Image scanning is enabled.

- **AppServer: JBossEAP 7.4.x**

JBoss EAP 7.4.x is used to deploy the Newgen Product's EJB components.

- **WebServer: JBossWebServer 5.7.x**

JWS 5.7.x is to deploy the Newgen Product's WEB components.

- **JDK: OpenJDK 1.11.x**

Inside the Docker images, OpenJDK 1.11.x is bundled. If there is some special requirement, then you can build the Docker images with Oracle JAVA on-demand basis.

- **Database: Microsoft SQL Server**

Newgen Products support three database servers: Microsoft SQL Server, Oracle, and PostgreSQL. The supported database versions for products are available in the Product Support Matrix.

- **PN Files Storage Server: Azure BLOB Storage**

PN files are encrypted files that contains Newgen products added, uploaded, and scanned documents. Also, Azure BLOB Storage is an object storage service that offers industry-leading scalability, data availability, security, and performance.

- **Log Files Persistence: Azure FileShare**

All the files are created inside a container are stored on a writable container layer. This means that the data doesn't persist when that container no longer exists or is terminated due to any reason. It can be difficult to get the data out of the container if another process needs it. Thus, to persist these types of data, use the Azure FileShare.

- **Configuration Files Persistence: Azure FileShare**

Same as log files and using Azure FileShare to persist the configuration files.

- **Load Balancer: Application Gateway or Azure Load Balancer**

It is not using Load Balancer directly. Instead using the Application Gateway Ingress controller. The Application Gateway Ingress Controller is a controller that triggers the creation of an Azure Load Balancer. The Ingress resource configures the load balancer to route HTTP or HTTPS traffic to different Kubernetes services within the cluster.

- **Session Persistence: Azure Cache for Redis**

Azure Cache for Redis is an easy-to-use, high-performance, in-memory data store. It offers a mature, scalable, open-source solution for delivering sub-millisecond response times making it useful as a cache or session store.

ules can then be applied to this event stream to detect abnormal behavior.

## 2.5 Kubernetes Autoscaling

In Kubernetes, auto-scaling happens at two-level depending on multiple parameters like CPU utilization, memory consumption, and so on.

1. **Pod level:** This is controlled by Horizontal Pod Autoscaling (HPA) or Vertical Pod Autoscaling (VPA).

- a. **Horizontal Pod Autoscaler (HPA)** automatically scales the number of pods depending on CPU utilization and memory consumption. Like if CPU utilization goes more than 80% then a new pod automatically gets launched and is mapped to the load balancer. Below are some points regarding HPA's default behavior:
    - The default HPA check interval is 15 seconds.
    - HPA waits for 3 minutes after the last scale-up events to allow metrics to stabilize.
    - HPA waits for 5 minutes from the last scale-down event to avoid autoscaler thrashing.
    - In HPA, the autoscaler works on the resource's request parameter instead of the resource's limit parameter.
    - HPA can be configured on CPU utilization, memory utilization, or both with an OR condition.
  - b. **Vertical Pod Autoscaler (VPA):** just like HPA, a Vertical pod autoscaler does not add more replicas but increases the number of system resources available to the particular pod. It automatically adjusts the CPU and memory reservations for the pod. VPA gives some recommendations based on the data collected from the matrix server and updates the existing pod with recommended values.
2. **Cluster or Node level:** It is also known as Cluster Autoscaler. It automatically adjusts the number of nodes in your cluster when pods fail to launch or are pending due to lack of resources or when nodes in the cluster are underutilized and their pods can be rescheduled onto other nodes in the cluster. In other words, Cluster Autoscaler adds nodes to a cluster whenever it detects pending pods that cannot be scheduled due to resource shortages and delete nodes from a cluster whenever the utilization of a node falls below a certain threshold defined by the cluster administrator. Thus, it automatically scales up or down depending on the resource utilization.

## 2.6 Logging architecture

This section describes the logging architecture for OmniDocs.

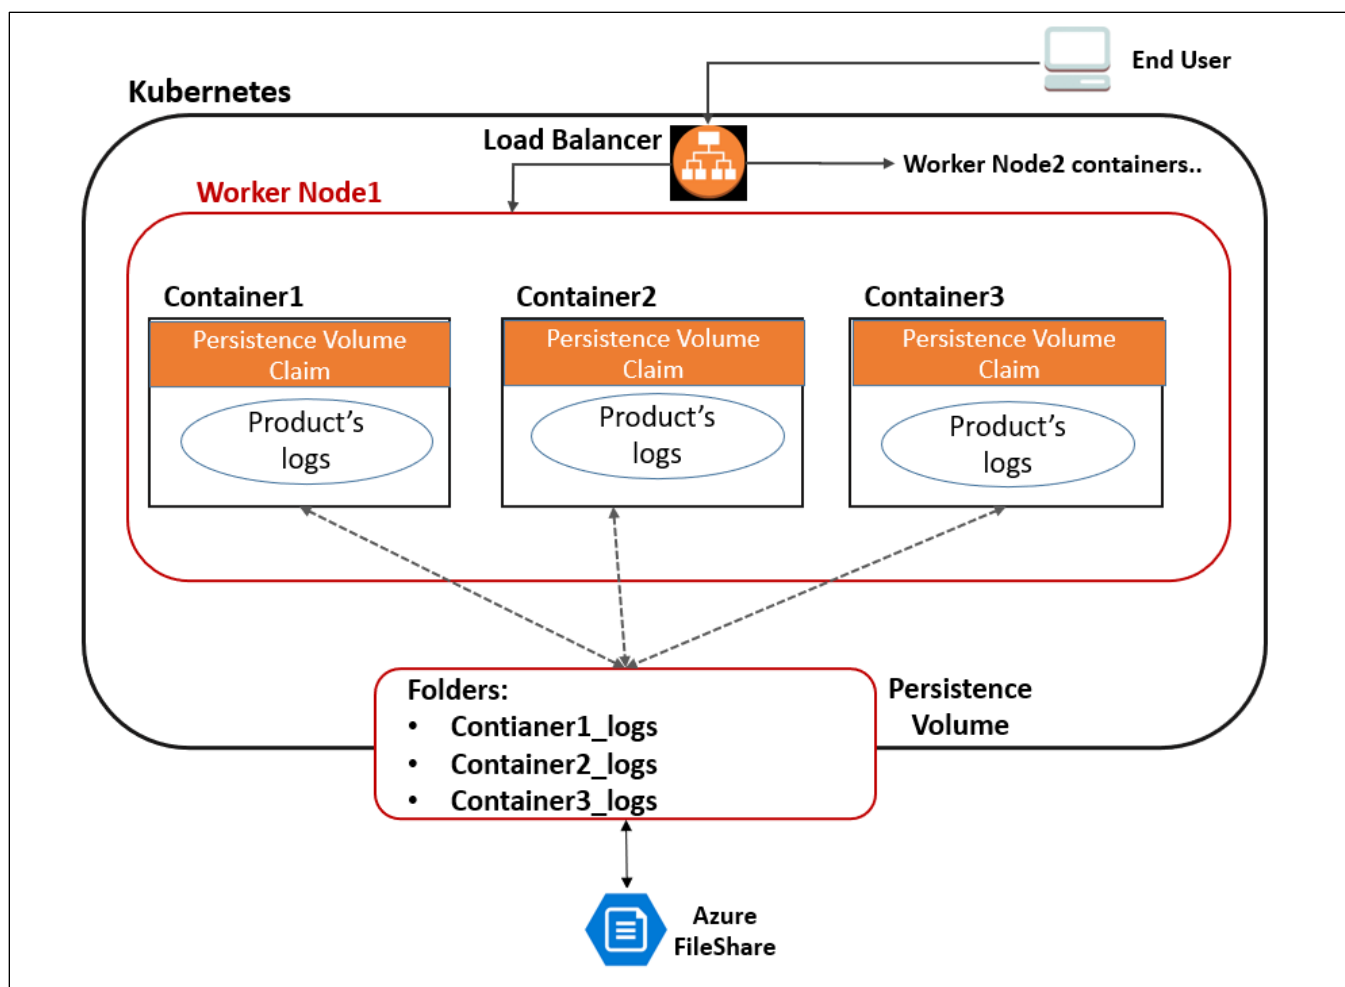


Figure 2.2

- All the files are created inside a container are stored on a writable container layer. This means that the data doesn't persist when that container no longer exists, and it can be difficult to get the data out of the container if you need it later. Thus, to persist these types of data, use Azure FileShare. Azure FileShare provides a simple, scalable, fully managed elastic NFS file system. It helps to share data across multiple nodes and pods.
- Since the product's log files are dynamic in nature and are created at runtime so they must be kept outside the container. Therefore, Azure FileShare is used to persist in the Newgen Product's log files.
- To persist the data in Azure FileShare, the concept of PV/PVC (Persistence Volume/Persistence Volume Claim) is used:
  - Persistence Volume: Persistent Volume (PV) is a piece of storage in the cluster that has provisioned using Storage Classes. It contains the Azure FileShare secretName and shareName that is already created during Azure FileShare creation. It is also used to set the access permission on Azure FileShare using mountOptions attribute.

- Persistence Volume Claim: A PersistentVolumeClaim (PVC) is a request for storage by a user. It is similar to a Pod. Pods consume node resources and PVCs consume PV resources. Pods can request specific levels of resources (CPU and Memory). Claims can request specific size and access modes (for example, they can be mounted ReadWriteOnce, ReadOnlyMany or ReadWriteMany).
- So, all the product's logs created inside the containers can be claimed by the PVC in some predefined folder structure.
  - For example, container1's product logs are stored inside the container1\_logs folder on an EC2 instance, and container2's product logs are stored inside the Container2\_logs folder on an EC2 instance and so on.
- Further, all the log folders created, is mounted to Azure FileShare so that whenever a Pod goes down, our product's logs must always be available in Azure FileShare. Whenever it is required, you can fetch it from there.
- Mounted product's logs are in sync in real-time. So, whenever the product's logs are generated inside the container, these logs are available in Azure FileShare as a persistence volume at the same time. There is no delay.

## 2.7 Container security

The Container Images are as follows:

- **Using Secure Images:** Since you are bundling numerous third-party libraries as dependencies, these dependencies contain some vulnerabilities. Also, Docker images are created using nested Docker images at multiple layers and any parent Docker image may contain vulnerabilities. So, it becomes necessary to scan all the Docker images against vulnerabilities. As all the Docker images are stored inside the Azure Container Registry, Azure ACR image scanning gets enabled by default. Whenever you push an image to the container registry, Security Centre automatically scans it and then checks for known vulnerabilities in packages or dependencies defined in the file.
- **No container should be running with root privileges:** The security context *RunAsNonRoot* is configured for all the containers in the deployment manifest file. This ensures that the application must not be running with administrator or root privileges inside the container.

**Container Runtime Environments are as follows:**

- **Restricted communication between containers:** In a Kubernetes cluster, containers can communicate with any other container running in the same cluster. However, containers need to communicate with each other to accomplish their goals, but they must not communicate with all the running containers. The reason being that if a hacker gets access to one container, then that hacker can communicate with any other container as well. So, allow containers to communicate to only those containers that are absolutely required to minimize the attack surface. To handle this

case, the Network Security Policy is implemented between the containers by defining some ingress and egress rules.

#### **Orchestrators:**

- By using Azure Kubernetes Service for container orchestration, it automates many of the tedious tasks required to deploy containerized apps. It can enforce role-based access control policies and network policies. For example, it can restrict the ability of a container to access resources on the host server.

#### **Container Registry:**

- Container registry providers already provided security on the container registry. You cannot push or pull any Docker images without authorization.
- A private container repository is available on the cloud, but it can be accessible through whitelisted users only.
- Vulnerability scanning is enabled to scan the pushed Docker images.

#### **Persistent Storage:**

- Azure FileShare is used to persist the Newgen product's configuration files and log files. Azure FileShare is the fully managed storage service of Azure. Access to Azure FileShare requires credentials that Azure can use to authenticate your requests. All the data kept in Azure FileShare is highly secure and no one can access the EFS without the required permission.

## **2.8 CI/CD Pipeline**

Use the CI/CD pipeline to manage deployments with Kubernetes orchestration on cloud platforms. Here at Newgen, there is a segregation of the Build Pipeline and Release Pipeline into two parts. The Build Pipeline is done by the Jenkins server which can be installed on an on-premises machine or a cloud machine. The Release Pipeline is managed by Azure DevOps Release Pipeline cloud service.

### **2.8.1 CI/CD Pipeline of Base Products**

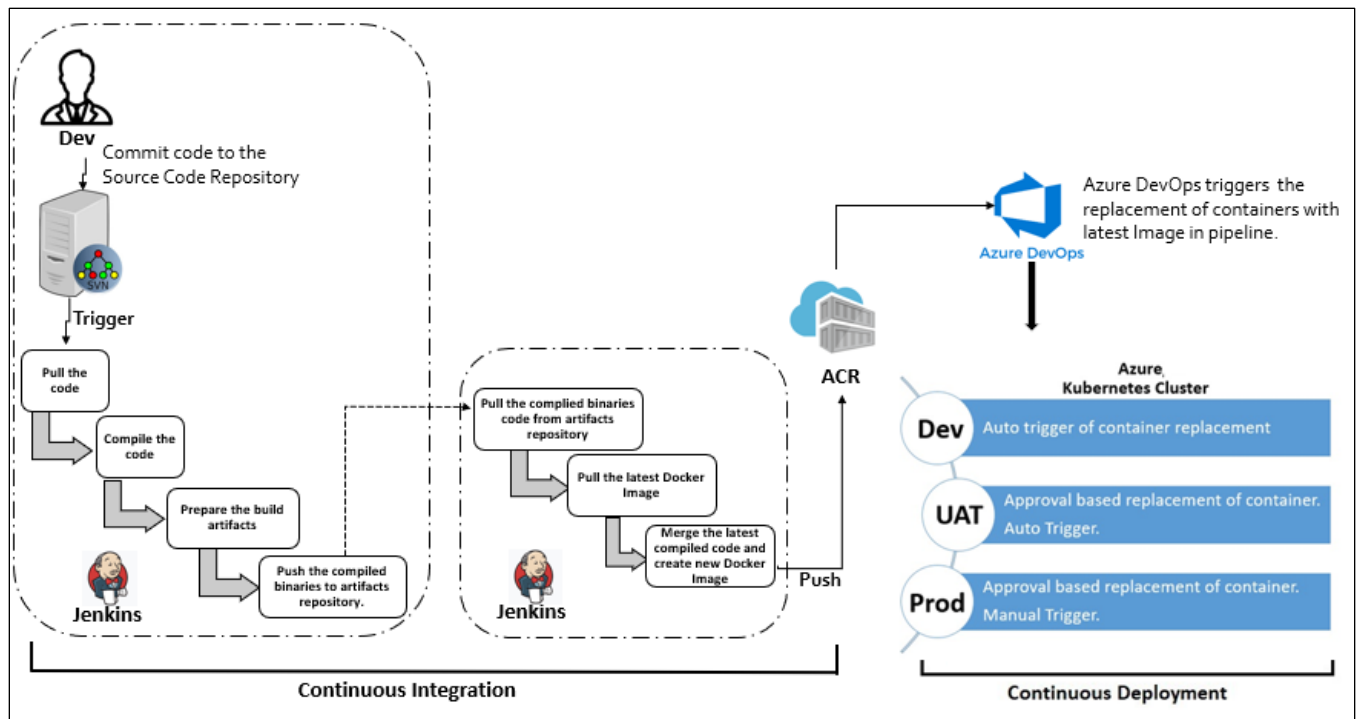


Figure 2.2

- Here, the Build Pipeline and Release Pipeline is segregated into two parts. Build Pipeline is done by Jenkins server which can be installed on an on-premises machine, or a cloud machine and Azure DevOps cloud service is used to manage the Release pipeline. Further, the build pipeline is segregated into two teams: development and the operational team. The development team keep pushing their code to the source code repository. As soon as they commit their code to the source code repository, the 1st Jenkins pipeline gets triggered. It pulls the latest code, compile them and prepare the build artifacts and push them to the artifact's repository like Azure FileShare or some other similar kind of media. The development team doesn't have to bother about Docker's image creation. After that 2nd Jenkins pipeline gets triggered. It pulls the compiled build artifacts from the artifacts repository, pull the latest Docker image and create a new Docker after merging the latest code changes. Last, it pushes the newly created Docker images to Image Repository like Azure Container Registry. After that Azure DevOps appears.
- In this architecture, there are three created stages: Dev, UAT, and Production and on each stage, deployment is quite different. You can have some more stages depending on the requirements.
- As soon as any Docker image is pushed to the Azuer ACR, Azure DevOps triggers the deployment to the Dev environment.
- UAT and Production deployments are approval based and they are called on-demand. Once you are ready to deploy to the UAT environment, trigger the UAT deployment. When that deployment is triggered, an approval mail is sent for the approval. After receiving approval, the UAT deployment starts automatically.

- The production deployment is also approval-based, but it is multi-level approval. To deploy to a production environment, approvals from all stakeholders are required. Most importantly, after receiving approvals from all the stakeholders, deployment to the production environment cannot be triggered automatically. A manual intervention mail is sent to the engineer who is supposed to deploy to production with a checklist. After evaluating that the checklist points are covered or not. If not, then the deployment to the production gets rejected.

### 2.8.2 CI/CD Pipeline of Custom Code

This section describes the CI/CD pipeline of custom code.

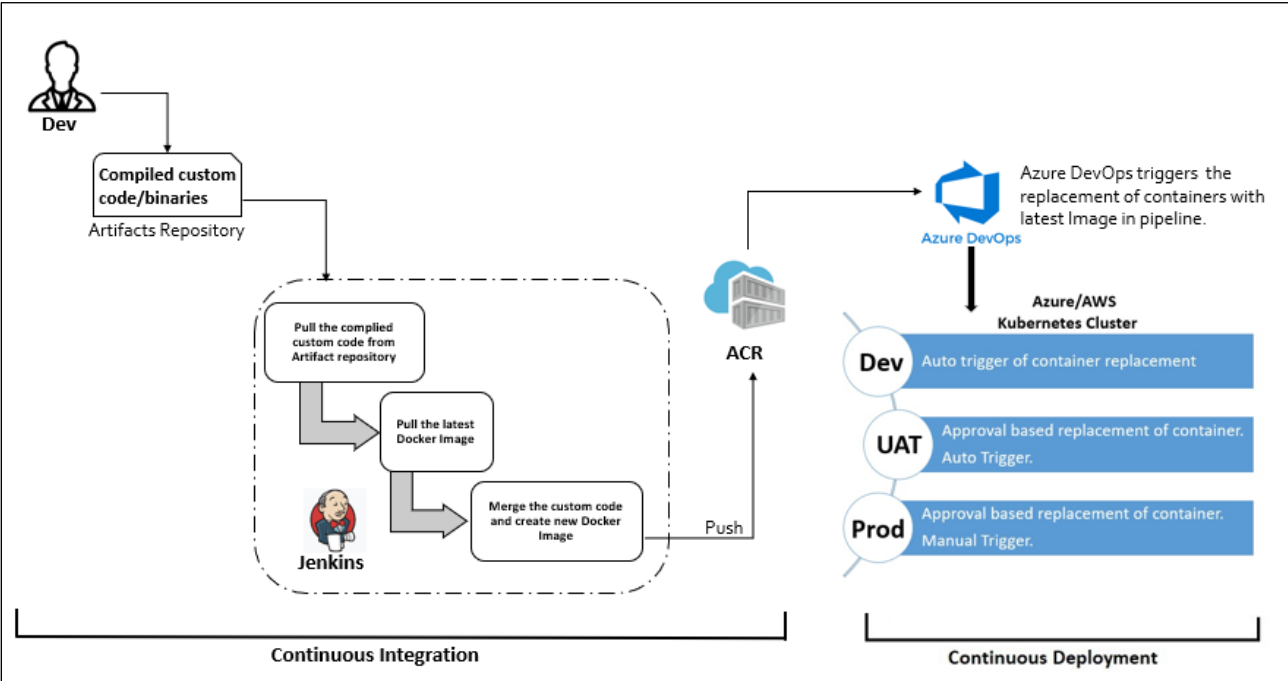


Figure 2.3

- The Custom Code deployment pipeline can also be configured and maintained in Jenkins.
- For initiating custom code deployment, the Implementation team push the compiled custom code to the artifacts repository. An artifacts repository can be anything like Azure FileShare, SVN, FTP, JFrog, and so on.
- After that Jenkins will pull the compiled custom code from the artifacts repository, it pull the latest Docker image and create a new Docker. After merging the custom code changes, it pushes the newly created Docker images to Image Registry.
- In this architecture, Cloud/Infra team have full access to initiate the build pipeline configured on the Jenkins server. The Implementation team having no access to this Jenkins server. However, a common artifacts repository is shared for both teams.



- As soon as any Docker image is pushed to the Azure Container Registry, Azure DevOps triggers the deployment to the Dev environment.
- UAT and Production deployments are approval based and they are called on-demand. To deploy the UAT environment, trigger the UAT deployment. When that deployment is triggered, an approval mail is sent to the project manager or team concerned. After receiving approval, the UAT deployment started automatically.
- The production deployment is also approval based but it is multi-level approval. To deploy to a production environment, approvals from all stakeholders are required. Most importantly, after approvals from all the stakeholders, deployment to the production environment cannot be triggered automatically. A manual intervention mail is sent to the engineer who is supposed to deploy to production with a checklist. After evaluating that all the checklist points are covered or not. If not, then the deployment to the production gets rejected.

### 2.8.3 CI/CD Pipeline of Product’s Hotfix

This section describes the CI/CD pipeline of the product’s hotfix.

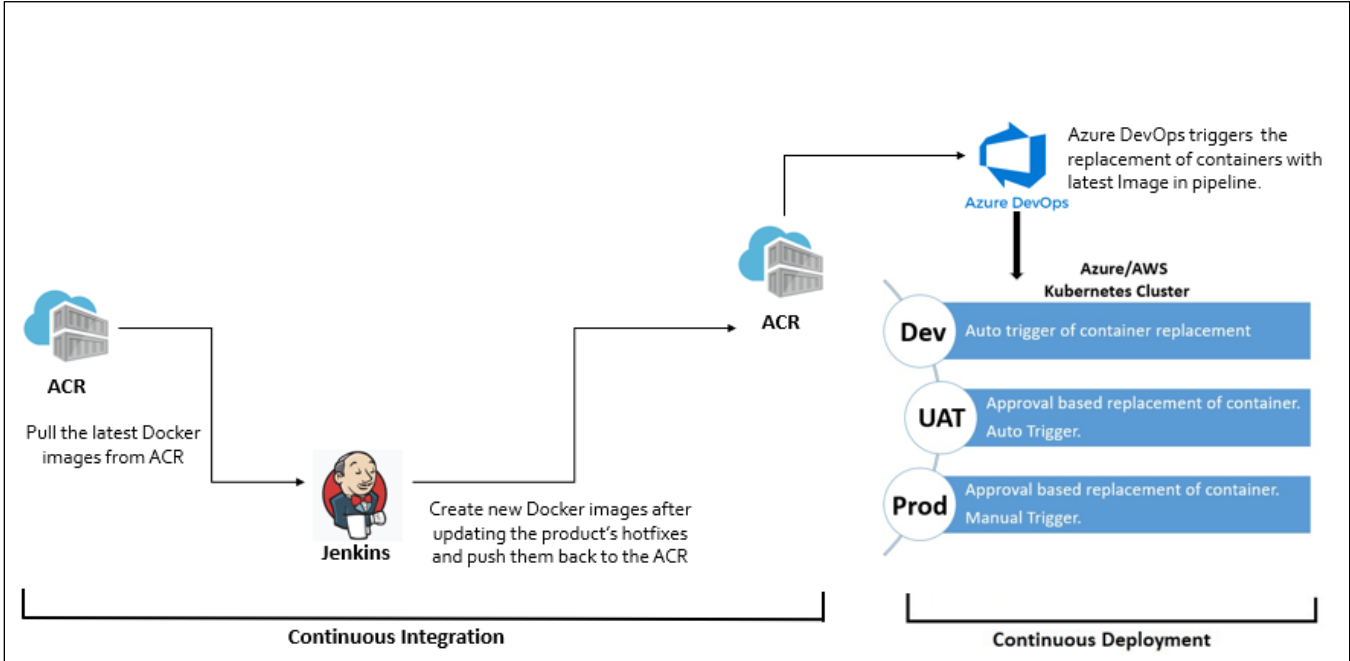


Figure 2.4

To deploy the Newgen product’s hotfix, follow the below steps:

1. Pull the product’s base images or latest images that are already deployed in the current environment from the container registry.

2. Update the hotfix files in the earlier running Docker images and create new Docker images. The deployment structure of these hotfix files (Dockerfile) is shared along with the hotfix files, which indicates how to update the Docker images.
3. Push the newly created images to the container registry.
4. As soon as any Docker image is pushed to the Azure Container Registry, Azure DevOps triggers the deployment to the Dev environment.
5. UAT and Production deployments are approval based and they are called on-demand. To deploy the UAT environment, trigger the UAT deployment. When that deployment is triggered, an approval mail is sent to the project manager or team concerned. After receiving approval, the UAT deployment starts automatically.
6. The production deployment is also approval-based, but it is multi-level approval. Approvals of all the stakeholders are required to deploy to a production environment. Most importantly, after approval from all the stakeholders, deployment to the production environment cannot be triggered automatically. A manual intervention mail is sent to the engineer who is supposed to deploy to production with a checklist. After evaluating that whether all the checklist points are covered or not. If not, the deployment to the production gets rejected.

## 2.8.4 Deployment Changes Against File Types

In Newgen products, majorly there are four types of files:

- Deployable files (\*.war, \*\_ejb.ear)
- Dependent Libraries (\*.jar)
- Configuration Files (\*.ini, \*.xml)
- Database Scripts

Deployable files and libraries can be merged inside the Docker container as there are no dynamic changes in these types of files.

But as configuration files are dynamic in nature so they must be kept outside the container.

For which the volume persistence is used and mapped them to external disk storage like Azure FileShare. So, whenever configuration changes are found in a custom code or product's hotfix, update the configuration files located at external disk storage along with updating Docker images.

For the database scripts, provide the DB scripts, which can be manually executed through Database Client software. However, if this operation requires automation, configure the DB scripts execution in Jenkins. In that case, create DB script execution job types in Jenkins.

## 2.8.5 Deployment Downtime

Deployment downtime is dependent on the types of binaries used in custom code or hotfixes that need to be deployed. If changes are in deployable files or into dependent libraries, then downtime is not required, as Kubernetes provides the rolling deployment feature. In this case, deploy the latest

Docker image, Kubernetes first launch the new container with the latest updated Docker image, make it runnable and when the latest container is ready to serve the requests, it starts deleting the existing container.

But, if changes required are in configuration files and database scripts then some downtime is required to update the configuration files and to execute the DB scripts, in the persistence mapped location. For the Production environment, must take the requisite downtime to ensure that product is working fine. To ensure that recheck the prerequisites and perform the smoke testing. Also, do some regression testing before making it live.

## 2.8.6 Docker Image Management for Rollbacks

For DockerImage management, two tags creation is recommended for each Docker image through Jenkins while building the Docker Images, before pushing them to the container registry.

For Example,

- omnidocs11.0web:latest
- omnidocs11.0web:build-10 (where build-10 is the build pipeline number)

This means that there's always a copy of the Docker image of each build pipeline in the container registry so that whenever you require to rollback the current deployment, you can use the previous build pipeline number tag – Docker Image for rolling back the deployment.

## *Appendix*

This guide contains third-party product information about configuring Microsoft Azure CICD Pipeline for Container Deployment on AKS Azure Kubernetes Cluster. Newgen Software Technologies Ltd does not claim any ownership on such third-party content. This information is shared in this guide only for convenience of our users and could be an excerpt from the Azure documentation. For latest information on configuring the Azure Kubernetes Cluster and Azure DevOps refer to the Azure documentation.