

# CD Pipeline: Speeding Up the Software Life Cycle

In the two-part series, this second and concluding blog highlights the CD pipeline that orchestrates the provisioning of infrastructure - Virtual Machines (VMs), network and storage, test automation for applications with databases during staging, release automation, etc. There are many other components like security, configuration management, service discovery, etc. that are relevant to the CD pipeline and not covered in this blog.

Pure Storage® FlashArray™ provides a standard data platform for all the phases in the CD pipeline that directly improves the deployment, infrastructure provisioning process and shortening the release cycle. A FlashArray appliance provides better performance for low latency databases used during the CD process with data reduction from inline deduplication and compression.

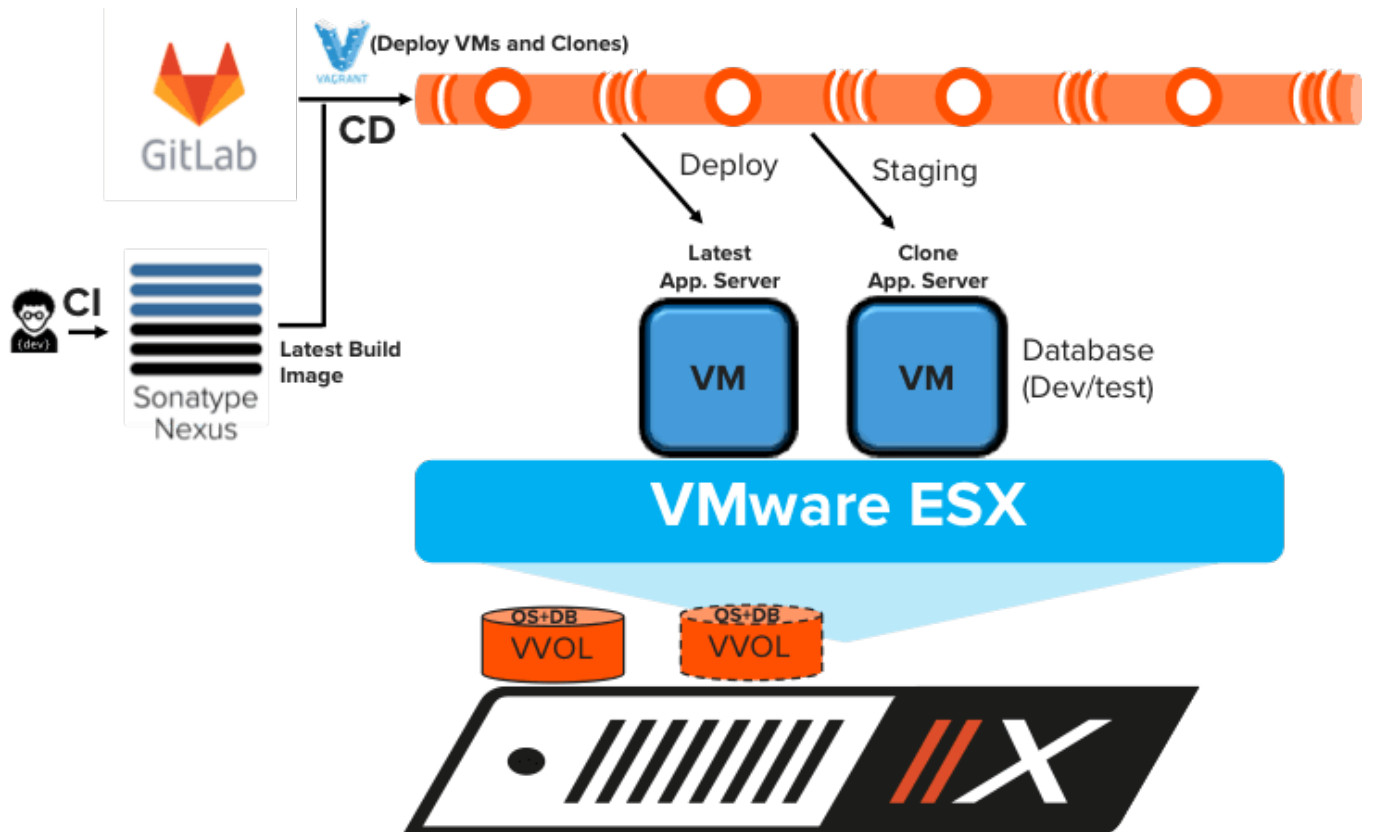
FlashArray also has a tighter integration with VMware vSphere using VMware vSphere APIs for Storage Awareness (VASA). This provides simple and easy data management capabilities like snapshot and cloning at the VM layer that extends seamlessly to the data layer in VVOLS that improves the end-user experience and the CD process. For more information, refer to the white paper [A Modern CI/Cd Pipeline on Pure Continuous Integration/Continuous Delivery \(Ci/Cd\) on Pure Storage](#).

Even though there is a lot of testing done in the critical phases (Shift-left) of the development cycle in the CI pipeline; there are still some roadblocks in the CD process. The following are some of the major challenges that business always wants to overcome during this process -

1. Test automation for new versions of applications with databases.
  - Sensitive data that is not relevant to the testing of the new application needs to be masked.
  - Provisioning a copy of the production database using migration scripts takes time.
  - Reverting from test failures takes longer time.
2. Promoting and demoting copies of the newer builds in and out of production requires additional configuration management to consume infrastructure. The traditional way of provisioning storage for data is not cost-efficient and perform well when the application starts to scale and operate in production.

The first part of the blog highlighted how Kubernetes and Docker containers with Pure Service Orchestrator™ (PSO) was used at scale for different phases of the Continuous Integration (CI) pipeline on a FlashBlade™ array. The CI process concluded after pushing the new version of the “wordpress” binary package into Nexus. The new build replaces the “wordpress” welcome screen with the Pure Storage logo and also includes additional comments for every blog post.

At the start of the CD pipeline, a new VM with an Ubuntu image is instantly spun up with a corresponding VVOLS on a FlashArray using a Vagrant file. The CD pipeline is automated to set up the new “wordpress”. “Wordpress” uses a MySQL database. The sensitive data in the production MySQL database can be masked and the production copy can be cloned instantly to test 1000 blog entries and 1000 comments for each blog. The VM clones instantly mount copies of the DB VMs with zero additional storage consumption by the VVOLS on FlashArray.



While testing new additional records in the MySQL DB adds space to the respective VVOLs on data platform, the data reduction was 24:1 on the FlashArray. This reduces the data footprint for testing other applications in parallel and provides a higher cost-efficient data platform in the software development life cycle. During the test failures, new copies of the database can be provisioned quickly for further testing.

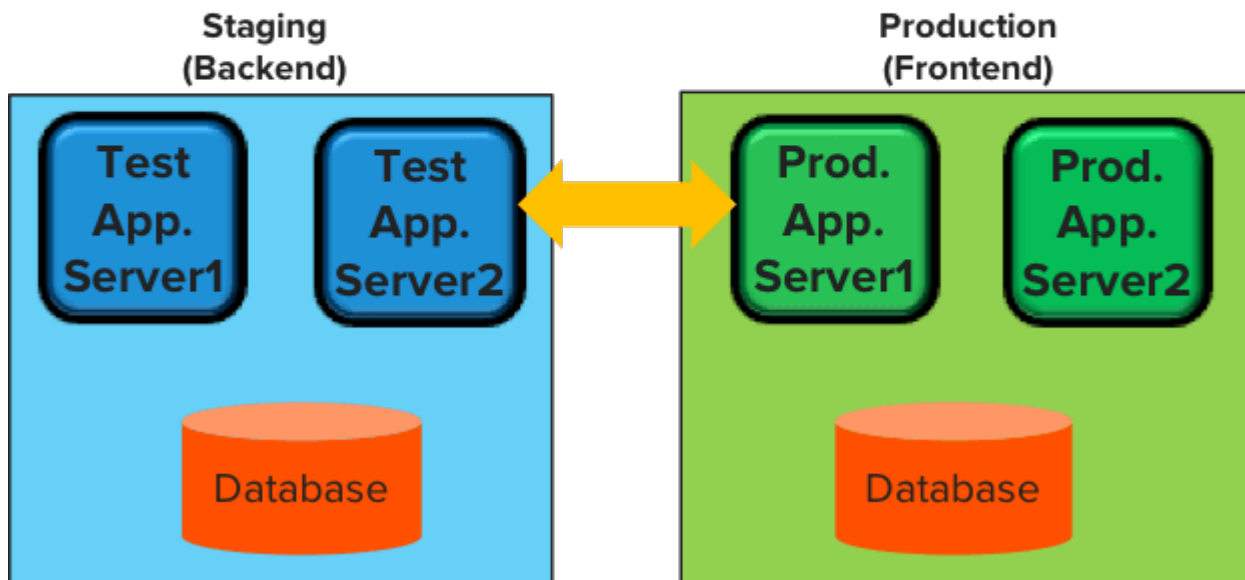
While there are many Container Platform as a Service (CPaaS) tools like OpenShift is available in the market, Blue-Green (B/G) deployment with HA-Proxy for load balancing provides a simple and easy way to release applications into production on VMs. The B/G deployments are also helpful for smaller and medium-sized applications.

Clients accessing  
App. Server



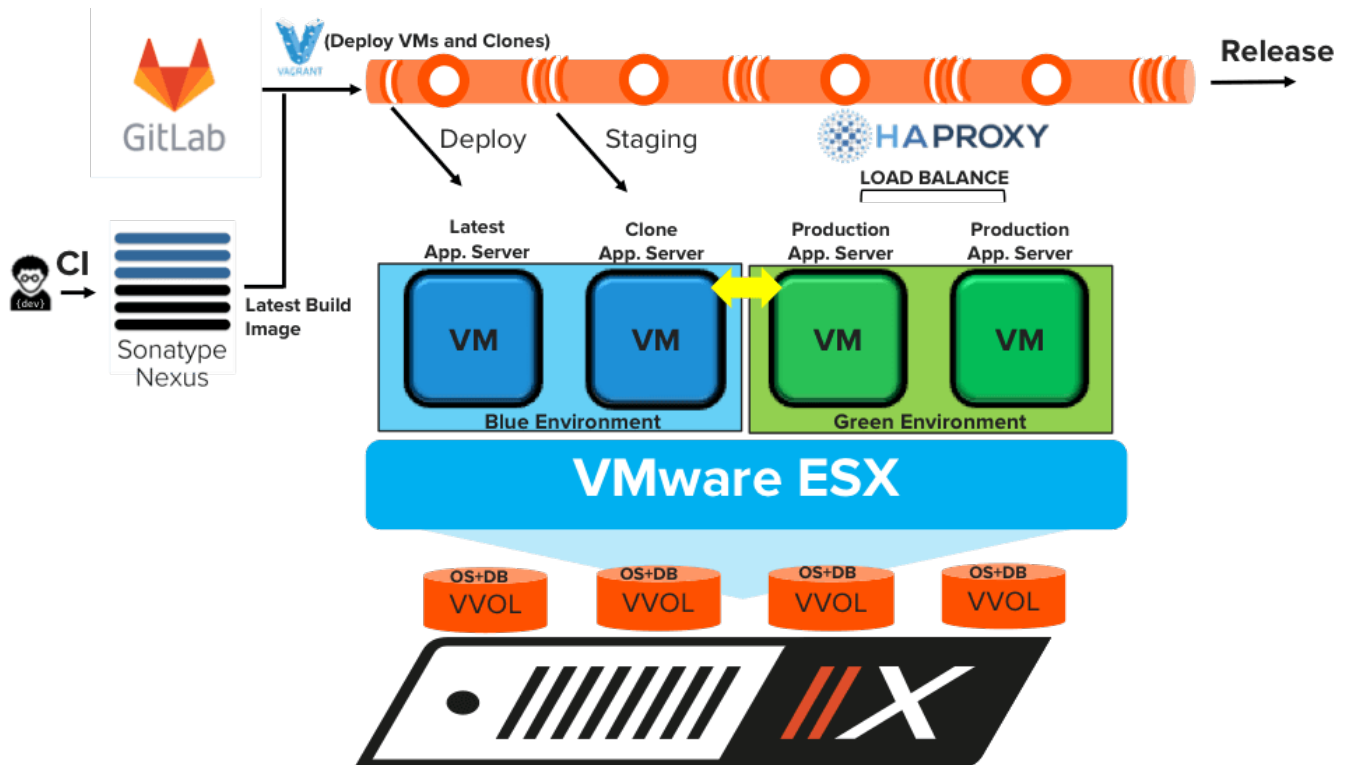
# HAProxy

REVERSE PROXY  
LOAD BALANCER



Setting up a B/G deployment with HA-Proxy on FlashArray allows to automate the CD pipeline to promote and demote new versions of the application ("wordpress") into production without any additional downtime. The application, VM or the data does not need to be ported to another platform for production use.

While HA-proxy can handle the load-balancing of the requests coming from the end-users to the VMs; the VVOLs corresponding to the cloned VMs take no additional space on FlashArray. The space accounting on the newly provisioned VVOLs will include the new changed blocks. At any given time, the latest and the current version of the application can co-exist.



Reverting back to the older release of the application can be accomplished quickly by reversing the proxy. While the VMs functioning as the application server pivot from the current to new versions of the applications, the dataset in the VVOLs for each of those VMs remain anchored in FlashArray.

FlashArray is no longer an incumbent to just store data but has evolved as a data platform in recent times to improve the CD pipeline efficiency and performance for the various phases mentioned in this blog. FlashArray not only provides high storage efficiency but also the ability to “fail fast; fix fast and release fast” during testing and releasing applications into production.

While this blog concludes with a couple of major use cases on FlashArray in the CD process on-premise, there can be extensions in the CD pipeline where applications are deployed or ported to AWS along with the data using Snap-to-Cloud technology in a [hybrid cloud](#) scenario.