

# Migrating to Platform9 Private Cloud Director

Assessing the process for migrating from  
VMware to Platform9 Private Cloud  
Director.

Justin Warren

March 2025

v1.0

Sponsored by  
Platform9



PivotNine

# Contents

1	Executive Summary	1
2	Overview	2
	Assessment Scenarios . . . . .	3
3	The Migration Process	4
	Automated Migration Overview . . . . .	4
	Discovery and Planning . . . . .	5
	Migration Preparation . . . . .	7
	Migration and Progress Tracking . . . . .	8
	Cutover . . . . .	12
	Backing Out . . . . .	13
	Cleaning Up . . . . .	14
4	Conclusion	15
	Assessing Suitability . . . . .	15
	Roadmap . . . . .	18
	About	20
	About the Author . . . . .	20
	About Platform9 . . . . .	20
	About PivotNine . . . . .	20

# Executive Summary

Migrating from VMware to Platform9 Private Cloud Director is a straightforward process that overcomes much of the complexity enterprises might otherwise fear. Instead of a complex and tedious migration, customers can use Platform9's tools to automate the bulk of the work, freeing staff to concentrate on the few complex or high risk migration activities that remain.

The freely available [vJailbreak](#) tool provides a user-friendly graphical interface for operators to migrate VMs en masse. Scaling to hundreds, if not thousands, of VM migrations at once, the tool supports even the largest customers to exit VMware more rapidly than some estimates might suggest.

While there are some workloads that will prove to be incompatible with automated migration, the list is smaller than one might imagine. It is also shrinking rapidly. The authors were surprised by just how quickly the development of features is progressing. The Platform9 team is improving the migration experience so rapidly that many of our early criticisms had already been addressed before this report had finished its initial drafting. Other roadmap items we expected would be six to twelve months away are slated for delivery within a quarter or two.

The migration process allows for a targeted infrastructure refresh, not a mere like-for-like replacement. Customers can choose to change very little, of course, but they should consider taking the opportunity to improve things as they go. Platform9's approach supports customer choice, both to keep things as they are, and to change where change is desirable. It is a welcome contrast to a landscape of vendors that are increasingly fond of imposing their will on customers who once welcomed them as strategic partners.

Moving VM-based workloads to Private Cloud Director could be justifiable on strategic multi-sourcing grounds alone. The credible threat of an active Platform9 presence opens up much potential for future negotiations. Coupled with the strategic flexibility of first-class container support, migrating could set customers up well for whatever their future IT direction might be. Workloads that make sense on VMs can stay there, but they don't have to.

For customers looking to migrate even some workloads away from VMware, Platform9 Private Cloud Director should be evaluated as a candidate. The effort and risk have been reduced so substantially that a migration business case can be justified for a much wider range of environments than other alternatives. Customers will be pleasantly surprised by how easy and flexible the migration process can be.

# Overview

We evaluated Platform9's migration process by working with Platform9's engineering team to run through a number of migration scenarios. We used lab equipment supplied by Platform9 and operated by Platform9 engineers under our direction, mirroring the experience of an enterprise using line staff or outsourced contract labour. We also worked under a compressed timeframe of three weeks for lab work, including scenario planning and environment setup.

The migrations we performed, while simplified for testing purposes, were nonetheless real. We encountered real bugs and configuration mistakes that needed to be addressed, as would be expected in any real migration. The fact that we experienced issues (all of which were overcome) reassured us that our approach to verification was sound.











Migrations <span>Q Search...</span>		
Name	Status	Progress
migration-mig-test-cbt-bak-clone-0-running-eg1	Running	 Running: Migration validated successfully
migration-mig-test-cbt-bak-clone-0-running-eg2	Running	 Running: Copying disk 0
migration-mig-test-cbt-bak-clone-0-running-eg3	Running	 Running: Copying disk 1
migration-mig-test-cbt-bak-clone-0-running-eg4	Running	 Running: Migrating VM from VMware to Openstack
migration-mig-test-cbt-bak-clone-0-succeeded-eg1	Succeeded	 Succeeded
migration-mig-test-cbt-bak-clone-1-failed-eg1	Failed	 Failed
migration-mig-test-cbt-bak-clone-1-failed-eg2	Failed	 Failed: Migration validated successfully
migration-mig-test-cbt-bak-clone-1-failed-eg3	Failed	 Failed: Copying disk 0
migration-mig-test-cbt-bak-clone-1-failed-eg4	Failed	 Failed: Copying disk 1
migration-mig-test-cbt-bak-clone-1-failed-eg5	Failed	 Failed: Migrating VM from VMware to Openstack
Rows per page: 25 1-10 of 10 < >		

Figure 1: An example of the vJailbreak migration progress screen.

We are confident that we have verified what actual customers will experience when using Platform9's tools and processes to migrate VMware workloads to Private Cloud Director. We do not pretend to have conducted exhaustive tests, designed to uncover every corner-case. That was not our goal.

Our goal was to reassure customers that Platform9's migration approach is indeed suitable for the majority of their needs. Most workloads can be migrated with low effort. Operators should be able to focus on the genuine and necessary complexity that exists within their environments, not what has been rendered mundane by Platform9's automated approach. Migrating need not be a tedious process of endless drudgery. It can, and should, be a straightforward task that is achieved without fanfare or fuss.

# Assessment Scenarios

We considered the following scenarios in our assessment of the migration process.

1. Self-contained application VMs.
2. 3-tier applications with web, app, and database servers.
3. Cold and warm migrations.

For each phase of the migration process described below, we consider each of these scenarios and its particular needs. We highlight features of the migration process that assist customers, as well as things to look out for. Customers should pay special attention to situations that may cause problems during a migration. We describe how customers can avoid problems with appropriate preparation, where possible.

## Self-Contained Applications

The simplest of migration scenarios, self-contained application VMs have few, if any, external dependencies. They provide very few, usually just one, point of external access on an outward facing network interface.

Migration is straightforward: Move the VM from VMware to Private Cloud Director, and ensure that external access is available in the new location. A small outage will be required during the cutover but, assuming all goes according to plan, the application could be running on VMware one day, and on Private Cloud Director the next, with application end-users none the wiser.

## 3-Tier Applications

A common deployment pattern in VMware environments, 3-tier applications have a web-frontend, an application tier containing business logic, and a database tier that stores the application data. Each tier is linked to the tier above and/or below, but operates somewhat independently.

The migration process will handle each tier separately, using the pre-existing modularity of the system design to assist the migration process. Customers may prefer to tackle the database tier first, or may prefer to leave it until the end, depending on how their environment is configured and the risks involved.

An assessment of a 3-tier architecture serves as a reasonable approximation for other architectures. An application with only two tiers can be considered a subset of a 3-tier application. Those with more tiers can be considered an extended version of this architecture. We can readily extrapolate our findings for 3-tier applications to many other situations.

## Cold and Warm Migration

We considered both cold migration and warm migration scenarios. In a cold migration, the workload is taken offline for the entire migration process. In a warm migration, data is migrated to the target environment while the source system remains online and in active use.

# The Migration Process

For the majority of common workloads, migrating from VMware to Private Cloud Director can be achieved successfully with a low-effort, automated process.

Every migration project tends to comprise the following phases:

1. Discovery and planning
2. Migration preparation
3. The migration itself
4. Cutover
5. Backing out
6. Post-migration activities

We have investigated the tools and processes used by Platform9 to assist customers in each phase of a migration to Private Cloud Director. In the sections that follow, we have highlighted particular aspects of these tools and processes that support our chosen migration scenarios in each phase.

We have assumed that customers have an existing VMware environment and have built a Private Cloud Director environment that is the target for the migration. We do not describe the configuration of these environments in detail, though we do highlight particular aspects of, for example, the storage or network configuration, where it is relevant to our discussion of the migration process.

## Automated Migration Overview

Before looking at the migration process in detail, it is useful to have an overview of how the automated migration process works in general. It works like this:

1. Customers install vJailbreak in the target Private Cloud Director environment and connect to its GUI console.
2. vJailbreak connects to the source VMware cluster using provided admin-level credentials. It discovers VMs on the source system and presents a list of VMs to be migrated.
3. An operator selects VMs from the discovered list to migrate as a batch. For each batch, the operator selects common parameters such as a mapping of source storage pools to those on the target, the migration method (cold or warm), whether to schedule the migration or to start immediately, etc.
4. The automated migration process copies data from the source VM datastores to the target datastores. If a cold migration was selected, VMs are first taken offline. If a warm migration was chosen, VMs stay online while data is copied in the background.
5. At the cutover point, source VMs are taken offline (if they are not already offline for a cold migration). Target VMs are brought online with the same settings as the original source. They appear, to all external systems, as if they were the original VMs, but now running on Private Cloud Director rather than VMware.
6. Customers verify that the workload is functioning as expected.

How this general process is adapted to suit specific circumstances is discussed in more detail below. While we do not cover every possible wrinkle or nuance, we have attempted to verify that Platform9's process covers the situations customers are likely to encounter most often. For situations outside those covered below, we encourage customers to contact Platform9 to discuss their specific needs. PivotNine consultants would also be happy to assist.

## Discovery and Planning

The vJailbreak tool automatically discovers an inventory of VMs running on the source VMware cluster. VMs that are offline in the source environment cannot be selected for migration. vJailbreak requires that the source VMs are brought online in order to correctly discover some configuration items, such as IP addresses, that can be dynamically allocated at boot.

The tool will provide a list of discovered VMs that can be selected for migration. The automated discovery process also detects storage pools in use and configured network interfaces. These will need to be mapped to the appropriate resources at the migration target. A list is presented to the operator to choose from, allowing customers to remap VMs to new networks if desired, or to maintain existing IPs if preferred.

### Migration Form

1

Source and Destination Environments

Source VMware

vcenter-phx

+ ADD NEW

Destination Platform

openstack-pcd

+ ADD NEW

2

Select Virtual Machines to Migrate

VM Name

Status ↑

Current IP

Network Interface(s)

OS

☐

estore-db01

Running

10.9.14.192

VM Network, VM Net...

Linux

☐

estore-web01

Running

10.9.0.80

VM Network, VM Net...

Linux

CANCEL

START MIGRATION

Figure 2: Discovered VMs are selected from a list in the migration GUI.

Unfortunately, the tool does not perform compatibility checking prior to migration. If a VM has incompatible settings that will prevent a migration from succeeding, this will not be detected until the attempted migration fails. While a source VM can be brought back online after a failed migration without ill effect, it would be frustrating to discover a preventable error midway through a lengthy migration window. We suggest that customers perform some rudimentary checking of their own, particularly of network configurations at the source and target prior to scheduling a large batch of migrations.

## Elevated Privileges

Platform9 provides a list of permissions, authority levels, and configuration settings that should be set before attempting a migration. The automated tool requires elevated privileges in both the source VMware environment and the target Private Cloud Director environment. This is inherent to the nature of a migration.



vJailbreak will be unable to discover the source set of VMs without appropriate permissions being granted. Customers should consider granting a temporary set of credentials to be used by operators during the migration and then revoked. Given the significant scope of permissions needed, the risk of elevated access falling into the wrong hands should be minimised as much as possible.

## Migration Preparation

After candidate VMs have been discovered, customers will need to answer a few critical questions:

1. Should the migration result in a like-for-like copy of the source system into the target, or is there an opportunity for workload consolidation, optimisation, or other improvements during the process?
2. How should source datastores be mapped to the target?
3. How should source networks be mapped to the target?
4. Are workloads good candidates for automated migration, or will they likely require some manual intervention or other special attention from operators?
5. Will migrations need to be performed in a particular order? Which VMs can be migrated in a batch? How large should each batch be, given the customer's risk profile and degree of comfort with the process?

## Workload Consolidation

The migration process does not require that VMs replicate all of the configuration decisions that were made on the source. For example, VMs currently using different storage pools could be consolidated into a single target VM pool to increase utilisation. Similarly, VMs could be moved onto storage with a higher (or lower) performance profile, perhaps splitting them out from peer workloads on the source.

Platform9 does not presume to make these decisions on behalf of a customer. Each environment is different, and customers should consider their desired outcome as part of migration planning. The tool supports both maintaining the status quo or making desirable changes. It is up to customers to decide which they prefer. Should customers change their mind later on, VMs can be readily migrated online using standard operational processes, including dynamic resource rebalancing. These migration-enabled decisions are not binding, merely helpful.

The migration process *does* perform some known-good optimizations at the individual VM level. Platform9 has identified these optimizations as universally applicable so they can be safely applied without customer action.

## Storage Settings

VMware datastores must be mapped to volume types in the target Private Cloud Director system. Customers will need to consider how best to map the old to the new.

The migration provides an opportunity to improve the environment here. Customers should not feel bound to past decisions about storage. In migrating, datastores could be consolidated to improve utilisation, or moved to higher-performance storage. Customers should embrace the opportunity to make beneficial changes that will improve the experience of end-users after the migration.

## External Storage

Note that network-attached storage, such as iSCSI, CIFS, or NFS storage hosted on external storage arrays, is not migrated via the vJailbreak tool. Only storage used to host the VM is migrated. Provided the target Private Cloud Director environment is configured with access to the same external storage arrays, VMs will reconnect to their external storage after cutover.

This simplifies the migration process substantially. Workloads with large amounts of storage, provided it is hosted on external arrays, can be migrated just as quickly as self-contained VMs with modest amounts of local storage. Customers should take advantage of external storage features, such as snapshots, to assist with establishing recovery points to be used if the migration process fails for some reason.

## Network Settings

For ease of migration, the target environment should be configured with networks that align with those on the source system. Connectivity to both the source VMware cluster and the target Private Cloud Director environment should be functional. End-user systems should be able to connect to VMs after migration without re-addressing or DNS changes.<sup>1</sup>

In some circumstances, customers may prefer that migrated VMs do not maintain the same network settings as the original. This approach is also supported by the automated process.

In our testing, the lab environment had different networks with different IP ranges in the source and target environments. VMs were configured with multiple networks: one public network with dynamically assigned addresses, and one private network with a static IP address. For our three-tier eCommerce scenario, we were able to simulate a configuration where application servers bridge between a public-facing web environment and an internal network connecting to database servers.

Reconfiguring external systems with the new network address of migrated workloads will require customer intervention and is not handled by the automated migration process. If customers do plan to re-address their workloads, some consideration of the post-migration effort involved should form part of the migration planning. For our three-tier eCommerce scenario, we knew the external URL of the site at both the source and target and did not reconfigure external load-balancers or DNS.

## Migration and Progress Tracking

For each batch of VMs selected for migration, the following information must be entered into the tool:

- Source VMware environment
- Destination Private Cloud Director platform
- The set of VMs to be migrated
- A mapping of source networks to target networks
- A mapping of source storage pools to target storage volume types

---

<sup>1</sup>Virtual machine MAC addresses do not change when the VMs are migrated from VMware to Private Cloud Director.

- The data copy method
- The cutover method

This information should have been determined as part of migration planning and be readily available. If operators are unclear on what settings to use at this stage, customers should pause the migration and re-assess if migration planning has been properly completed.

Migrations <span>Q Search...</span>		
Name	Status	Progress
migration-mig-test-cbt-bak-clone-0-running-eg1	Running	Running: Migration validated successfully
migration-mig-test-cbt-bak-clone-0-running-eg2	Running	Running: Copying disk 0
migration-mig-test-cbt-bak-clone-0-running-eg3	Running	Running: Copying disk 1
migration-mig-test-cbt-bak-clone-0-running-eg4	Running	Running: Migrating VM from VMware to Openstack
migration-mig-test-cbt-bak-clone-0-succeeded-eg1	Succeeded	Migration validated successfully
migration-mig-test-cbt-bak-clone-1-failed-eg1	Failed	Copying disk 1
migration-mig-test-cbt-bak-clone-1-failed-eg2	Failed	Migrating VM from VMware to Openstack
migration-mig-test-cbt-bak-clone-1-failed-eg3	Failed	
migration-mig-test-cbt-bak-clone-1-failed-eg4	Failed	Failed: Copying disk 1
migration-mig-test-cbt-bak-clone-1-failed-eg5	Failed	Failed: Migrating VM from VMware to Openstack
Rows per page: 25 1-10 of 10 < >		

Figure 3: Additional progress detail is available for individual VM migrations.

## Choosing VMs

Discovered VMs on the source can be filtered with a keyword search to locate a batch of VMs to migrate. Basic information including VM Name, status, OS and current IP address are listed. Operators select the VMs with checkboxes, either individually or as a group.

Groups of VMs thus selected are migrated as a batch. Note that all VMs in the batch are migrated together. There is not currently any mechanism for migrating VMs in an operator-specified order. If operators wish to migrate some VMs before others, the VMs must be migrated as separate batches.

In our testing, we were able to separate the migration of related VMs by tier: migrating the web tier separately from the app tier and database tier. Common network access in source and target environments allowed app tier VMs migrated to Private Cloud Director to communicate with VMs that had not yet been migrated.

Customers can preserve application functionality even as they stage migrations to align with outage windows; applications can have half of their VMs in each environment, provided they are configured appropriately. This smooths the path to migrations, particularly for large migrations coupled with limited outage windows. Migrations can also be grouped by organisational partitions, such as by business unit, helping to align the migration process with existing organizational change management processes.

## Mapping Networks and Storage

After selecting VMs, operators map source networks to target networks using pulldown lists of discovered networks. The mappings entered should match those decided on during migration planning. Similarly, the discovered source storage datastores will be mapped to target storage volume types.

The tool presumes that operators know what they want and will follow the configured instructions. There is potential for failure if operators select an incorrect mapping, as the tool has no way of knowing what the operator's true intentions are. Operators will need to take some care that they are selecting the correct mappings.

## Choosing Migration Options

Customers can choose a fully automated migration using the default options or exert greater control over how and when the migration occurs. The fully automated option copies VM data while VMs remain online, then shuts down the source VMs, finalises the data copy, converts the VM format, and brings the VM online in the target environment. VMs can also be taken offline before data is copied from source to target. This approach can reduce the time taken to migrate, at a cost of a longer outage window.

Operators can choose to migrate VMs immediately, or schedule the entire migration process to occur at a later time. We suggest that customers become familiar with the migration process by performing a small number of manually triggered migrations before moving to scheduled automatic migrations.

## Migration Form

VMware Datastore

OpenStack VolumeType

+

4 Migration Options (Optional)

^

☒ Data Copy Method

☐ Data copy start time

Copy live VMs, then power off

Data Copy Start Time

☒ Cutover Options

☐ Post Migration Script

Cutover immediately after data copy

Post Migration Script

Cutover immediately after data copy

Admin initiated cutover

Cutover during time window

☐ Retry On Failure

Select this option to retry the migration incase of failure

CANCEL

START MIGRATION

Figure 4: Multiple data copy and cutover options can be selected in the GUI.

Operators can cutover VMs automatically once the data has been copied to the target environment, or wait for a cutover trigger. Cutovers can be triggered to occur within a specific time window, or to wait for operators to initiate the cutover manually. These options help customers to align the short outage required for cutovers with operational needs.

## Migration Verification Checks

The system checks for error conditions throughout the process. If operators opted to retry on errors, the process will abandon the migration if an unrecoverable error is detected and restart the migration from the beginning. Only a single retry will be attempted; if a second unrecoverable error is detected, the migration will be aborted completely.

If a migration is unsuccessful, the source VM will not be brought back online automatically. If customers wish to minimise downtime, we suggest they choose the “operator initiated cutover” option and a warm migration. Cutovers can then be scheduled within an agreed window and, if unsuccessful, source VMs can be manually brought online by operators.

## Cutover

The cutover point is the most critical in any migration effort. This is the moment of truth. The moment where we determine if the newly migrated system is viable in its new home, or if something has gone wrong and we will need to go back to the way things were.

For many VMs, a fully automated cutover will be appropriate. Both cold and warm migrations are suitable, particularly for self-contained workloads. Allowing the system to trigger the cutover as soon as the data has been copied helps to minimise the outage window. In our testing, self-contained VMs and tiers of our three-tier app successfully cut over without external intervention. No time was lost waiting for external events or operator intervention.

For some workloads, the cutover time is more sensitive and needs greater operator control. The warm migration method is generally more suitable here, as it minimises VM downtime. Once the data has been copied, the system keeps the target synchronised with changes at the source while it waits for the cutover to be triggered.

The manual cutover method is useful for preparing the migration in advance, perhaps by starting the data copy operation hours or even days before the planned cutover time. For bulk migrations, this helps to minimise operator effort, allowing the automatic processes to handle as much as possible. Once the appropriate time arrives, triggering the cutover is a straightforward process of taking the source VMs offline. The cutover proceeds automatically from that point, and the VMs are brought online at the target once the cutover is complete.

## Verifying Migration Success

The migration process performs some rudimentary automatic verification tests that a migration has concluded successfully. The tests are aimed at verifying the migrated VMs are online and minimally functional in the target environment. These tests are performed at the infrastructure level, and should be used as a precursor to other testing, not a replacement.

Customer-specific acceptance testing is the primary method of verifying the system works as expected after migration. This testing is workload specific, and will need to be developed by the customer for each of their workloads. Ideally, a robust set of well-maintained acceptance tests—including integration, performance, and user acceptance tests—already exists.

Operators can select a post-migration script to run once the migration has completed. This script can be used to integrate with other tools, such as automated test suites, to provide additional automated assurance that migrations have completed successfully. Customers can leverage their existing investments in automation to support the migration process with minimal additional effort.

Note that the post-migration script feature is not designed to perform additional migration-specific tasks during the migration of VMs. For example, if customers want to reconfigure VMs before they are brought fully online in the target environment, the reconfiguration will need to be manually performed after the automated migration process completes. Manually triggered cutover will likely prove the best option here.

The goal of the automated migration process is to move VMs from the source environment into the target environment with as little change as possible. Adding change to the migration process is a delicate process that is highly workload specific. Customers should consult with Platform9 regarding their needs if they anticipate making in-flight changes during migrations.

## Dealing with Failure

If the migration fails, source VMs are left offline; they are not automatically brought back online. Platform9 has chosen to err on the side of caution and allow customers to verify that VMs can be safely returned to service before they are brought online.

In some cases, migration failure will be due to small and easily-corrected errors. Customers will not need to abort the migration entirely. After troubleshooting the error, customers can bring VMs online manually in the target environment, completing the migration.

Some migration failures can occur due to uncommon, transient conditions. Operators can choose to retry a migration on failure to automatically overcome such transient issues. If the retry option is selected, and the first migration attempt fails, the system will attempt to perform the migration one more time using the same migration options as the first attempt. If this second attempt also fails, the system will not make a third attempt.

However, occasionally more substantial errors will occur and the migration cannot be completed successfully. Customers need a mechanism to abort the migration and go back to how things were.

## Backing Out

If the migration does not go according to plan, and errors cannot be corrected, backing out of the migration is fairly straightforward. The original VM data is left intact and unchanged, ready to be reinstated if needed.

If the migration failed, and the VM was unable to be migrated or brought online at the target successfully, backing out involves bringing the original VM back online. Operators can restart the VM using their existing tools and processes. The migration can be attempted again once the errors have been investigated and remedied.

Backing out is more complex if the migration succeeds at the infrastructure level but further verification determines the migration was not properly successful and must be cancelled. In this scenario, the backout and recovery process is workload specific.

Ideally, no important transactions or data were completed at the target between the infrastructure migration completing successfully but the post-migration testing failed. In that case, backing out would involve shutting down the VM at the target and restarting the original, source VM. The migration can be re-attempted later.

If data has been written to the target, re-applying the same changes to the source system may be required. If customers believe this situation is likely, they should consult with Platform9 on their specific circumstances and devise a migration plan that minimises their risk. The best approach is to ensure that the migration has indeed been successful before allowing irreversible changes to be made to the target system.

## Cleaning Up

Minimal infrastructure is required for the migration, and cleaning up is fairly straightforward. vJailbreak workers used for scaling the migration effort can be scaled down to zero using the built-in interface. This scaling process is a regular feature of the Private Cloud Director platform and presents no special difficulties.

The vJailbreak installation can be completely removed once it is no longer required. Alternately, customers can choose to leave the minimal installation in place, but not active, just in case. It has a minimal resource footprint and will not interfere with regular operations.

If customers heeded our advice about using specially-provisioned credentials for the migration, removing access is also straightforward. If further migrations are discovered to be necessary, re-instating access with new credentials is preferable to leaving privileged access intact but unused. Cleaning up is made much easier if access is always provided on an as-needed basis for the minimal time necessary.

Decommissioning the source environment is beyond the scope of this report.



# Conclusion

The migration process offered by Platform9 is robust and covers the majority of virtualised workloads. While coverage is not universal and some situations will require manual effort by customers, these situations are relatively rare. The automated process reduces the operational burden of a migration substantially, freeing up resources to concentrate on these rare situations.

The inherent complexity of migrations not suited to automated migration means that most customers will choose to manually shepherd these systems through a migration anyway. Generally, these are high-value, complex applications in situations specific to a customer. Customers naturally wish to treat them with special care.

## Assessing Suitability

Every customer environment has a mix of workloads of varying complexity. The easiest systems to migrate are always those that are the best maintained, best documented and least complex. As complexity increases, so does the migration challenge. Migrating from VMware to Private Cloud Director is no different.

We provide some general guidance regarding workloads that are good candidates for an automated migration below. We further outline workloads that may prove challenging to fully automate, but nonetheless may be migrated using Platform9's approach with some additional care and attention from customers.

## Good Candidates

Workloads with the following characteristics are good candidates for automated migration and should prove straightforward.

- Stateless workloads
- Self-contained, decoupled workloads
- Workloads with low data volumes
- Workloads connected to external storage

## Stateless Workloads

Stateless workloads, or those with very little state, are straightforward to migrate. There is little data to move, simplifying the data-copying stage of the migration process. Workloads such as web servers, software load balancers, applications deployed using modern DevOps/GitOps approaches, etc. are prime candidates for automated migration.

The migration process need only convert the source VM forward to the target Private Cloud Director format during cutover, minimising the outage window required. Configuration settings are readily migrated and there are few complexities that may derail the automated process.

Workloads in this category can be grouped in large batches and migrated en masse, limited only by the customer's risk appetite and the capacity constraints of the infrastructure.

## Self-Contained, Decoupled Workloads

Applications that are self-contained or decoupled from other workloads in a modular fashion can generally be migrated easily. The unit of workload is the VM, thus each VM can be migrated without adversely affecting other systems. Platform9's processes and tools are well suited to migrating these workloads automatically in a plurality of circumstances.

The majority of the migration effort for these workloads consists of moving the data associated with the VM, followed by the VM format conversion process. By using a warm-migration approach, the outage window required for cutover can be minimised, much like that for stateless applications.

Self-contained or decoupled workloads with modest data volumes will present no special migration considerations. Migrations can be planned in large batches, taking into consideration the resource constraints of the customer's infrastructure. Scaling the migration will largely be a function of providing additional network bandwidth and compute resources for the cutover conversion process.

## Low Data Volumes

Where data is to be migrated into the new environment, the automated process will automatically copy data from the source system to the target system. If the data volume is large, and the network bandwidth restricted, this could be a time-consuming process. A warm migration, where data is copied in the background and changes replicated from source to target with the workload online, is likely to be the best choice here.

Increasing the network bandwidth available will improve the migration experience. Consideration should also be given to which workloads are batched together. There is merit in combining the migration of VMs with larger data volumes with some stateless or low-state workloads to make optimal use of the capacity available. This will speed up the migration of the environment as a whole, at a cost of some scheduling complexity.

## External Storage

Workloads that use external storage can be good candidates for migration, particularly those using network-attached storage, either SAN or NAS. Provided the external storage is supported by Private Cloud Director, migrating the storage can be as simple as disconnecting the primary storage from the source system and reconnecting the storage to the migrated system as part of cutover.

## Challenging Situations

Systems with the following characteristics will prove the most challenging to migrate to Private Cloud Director.

- Use of Raw Device Mapping (RDM)
- Application-level clusters
- Large data volumes

- No maintenance windows

Migrating these systems is still likely possible, but it will require special handling by dedicated staff. These workloads should be excluded from the automated migration process and handled separately.

## Raw Device Mapping

Raw Device Mappings (RDMs) are not currently supported by vJailbreak and Private Cloud Director. If the use of RDMs is vital for the workload in question, they may not be viable candidates for automated migration. However, RDM support is on the roadmap and is likely to be delivered within the next two calendar quarters.

Customers should not rule out migrations completely if they make use of RDMs. Instead, these workloads could be left until later while other workloads are migrated first. Once RDM support is made available, those final few RDM-enabled workloads could be addressed at the later stages of the migration exercise.

## Application-Level Clusters

Application-level clustering, such as for databases like SQLServer or PostgreSQL, tightly couples the application state of multiple VMs. Ensuring this application state is migrated consistently requires careful handling that is application specific. Databases in particular have their own specific mechanisms for ensuring consistency across the cluster.

Platform9's automated migration tools do not currently support these multi-VM consistency synchronisation mechanisms. Customers will need to manually intervene to ensure clusters of VMs are consistent prior to triggering a migration. Most likely the clustered application will need to be shut down, while the VM itself remains online. Once the VMs have been migrated, the application can then be restarted manually, and checked to ensure that it remains in a consistent state.

## Large Data Volumes

If very large data volumes exist, and network bandwidth is constrained, copying the data during migration may prove challenging. There must be sufficient space in the target Private Cloud Director environment to accommodate the data, and the change rate must be below the network bandwidth that is available. Rather than constraining individual VM migrations, this is mostly a constraint on how many VMs can be migrated in a batch.

The capacity of the target environment should be sized appropriately to accommodate the storage used in the source environment. If the target environment is accidentally undersized, migrations can be paused while additional storage is added to the target.

The network bandwidth available should be sufficient to synchronise changes from the source and target during the data replication period. If bandwidth is constrained, fewer VMs will be able to be migrated in a given batch. Customers should ensure they have appropriate infrastructure telemetry in place to monitor data change rates on the source systems, and the network utilization during migrations. Scheduling migrations for periods of low change rate can also assist.

## No Maintenance Windows

For VMs where no downtime can be scheduled, migrating to Private Cloud Director will likely prove challenging. The cutover processes supported by Private Cloud Director all require an outage to be taken. They involve converting the VM data format from VMware to the open format supported by Private Cloud Director, which must be performed offline. Fully online, live migration of VMs from VMware to Private Cloud Director, with no downtime for the VM, is not currently possible and is unlikely to be in the short or medium term.

In practice, such workloads are rare. Very high availability is usually achieved with clustering mechanisms that involve more than one VM. Individual VMs can then be taken offline for maintenance from time to time without interrupting the active workload. Without such a design, outages are forced upon the application when faults occur, completely undermining the high availability goal.

The migration process for VMs in this situation is highly dependent on the application workload, the application clustering mechanism used, and the specifics of the customer environment. For assistance with these environments, we advise speaking to Platform9, technology partners, or PivotNine consultants to assess how an online migration might be achieved.

## Roadmap

The automated migration process is under active development and there are always opportunities for improvement. Platform9 has demonstrated it is responsive to customer suggestions and delivers improvements quickly.

Platform9 has said that it intends to add support for Raw Device Mappings within the next 3-6 months. This is more rapid than we were expecting, and should allow many more workloads to be migrated automatically with little or no intervention from operators.

Rolling migrations are also under active development. This will progressively convert a cluster of VMware ESXi nodes to Private Cloud Director. VMs on an ESXi node will first be automatically migrated to a minimal Private Cloud Director environment. Once all VMs on the node have been migrated, the node will be converted into a Private Cloud Director node and join the Private Cloud Director cluster. The process can then repeat until all nodes have been converted to Private Cloud Director. There are some nuances in the full procedure yet to be clarified, such as automated rollback on failure, but this feature should permit conversion of compatible clusters essentially in-place with very little manual effort.

Platform9 is also working on a health check mechanism for verifying that services running in a VM have been migrated successfully. It will use API endpoints, a method that is in common use in cloud-native environments and familiar to developers of container-based microservices. Health checks should augment the existing post-migration script feature with a more flexible approach. It will help to more fully automate the end-to-end migration process and give customers additional confidence for large-scale migrations.

## Additional Opportunities

Beyond the existing roadmap, there are several improvements and enhancements that we have identified that would benefit customers looking to migrate from VMware to Private Cloud Director. We have not listed every possible improvement, but highlight some of the more obvious ones that customers are likely to encounter.

Some rudimentary checking for known or likely error conditions prior to migration would help customers avoid attempting migrations that will not succeed. Discovering a migration was never possible after waiting some time for data to be copied is frustrating. Having to troubleshoot a predictable error that could have been avoided is a waste. The more Platform9 can assist operators to prepare themselves for success, the better the migration experience will be.

It should be possible for source VMs to be automatically brought back online if a migration fails. While there are some risks associated, an automated migration should either succeed fully or automatically revert to the situation prior to the attempt. Customers should be able to trust that the automated process will not result in lengthy downtimes where operators must always manually intervene if something goes wrong. While this option should not be the default position, it should be a choice customers are able to make for themselves.

Support for database cluster consistency mechanisms would assist with migrating database clusters where the cluster spans multiple VMs. It would be helpful if popular databases like PostgreSQL, SQLServer, MySQL, or Oracle could be detected and made consistent prior to VM shutdown during cutovers. This would remove the need for careful manual handling by operators.

As mentioned in the executive summary above, some improvements we identified early in the validation process were implemented before we finished drafting this report. For example, at the beginning of our research, migration credentials were not saved between migrations and had to be re-entered each time. Credentials are now saved and can be reused for multiple migrations. Platform9 deserves credit for being able to respond so quickly to improve the customer experience.

# About

## About the Author

Justin is the founder and principal analyst at PivotNine. An accomplished IT leader with global experience, he works with clients on enterprise data center, cloud, and information security technologies. His writing has appeared in a variety of well-known outlets including *Forbes.com*, *The Register*, *iNews.com.au*, and *CRN Australia*.

Justin holds an MBA from Melbourne Business School and his preferred programming language is Python.

## About Platform9

Platform9 is the leader in simplifying enterprise private clouds. The company's flagship product, Private Cloud Director, has all of VMware's enterprise-grade features today along with private cloud features for the future. Platform9 was founded by a team of VMware cloud pioneers and has tens of thousands of nodes in production at some of the world's largest enterprises. Platform9 is an inclusive, globally distributed company backed by prominent investors, committed to driving private cloud innovation and efficiency.

All opinions expressed in this report are the genuine opinions of the author and PivotNine. While we speak to and listen to sponsors, we arrive at our conclusions independently and maintain editorial control over all published materials. Any errors or omissions are our own.

## About PivotNine

PivotNine Pty Ltd is a boutique analysis and consulting firm specialising in enterprise and developer technologies.

We assist clients with the evaluation and selection of enterprise and developer technologies for datacenters, cloud, and information security. We conduct custom research on behalf of our clients to better inform their decision-making, and to ensure they are making an informed choice both now and into the future. We assist vendors with product marketing, marketing strategy, and communicating their product benefits to enterprise buyers.

## Become a Client

PivotNine welcomes new clients. To learn more about how we can help you to buy, sell, and use enterprise technology, contact us today.

[sales@pivotnine.com](mailto:sales@pivotnine.com) · [pivotnine.com/contact](https://pivotnine.com/contact)

---

Reproduction and distribution of this document in any form without prior written permission is prohibited. The information in this document is from sources believed to be reliable. PivotNine disclaims all warranties as to the accuracy, completeness, or adequacy of such information. If any legal issues are discussed in this document, let's be clear that it's not legal advice; PivotNine is not a law firm, and we make no claims that we provide anything even resembling legal advice, so if you want actual legal advice, go and hire a lawyer. PivotNine is not a financial adviser. You should consider seeking independent legal, financial, taxation or other advice to check how information in this document relates to your unique circumstances. PivotNine shall have no liability for errors, omissions, or inadequacies in the information contained in this document. PivotNine is not liable for any loss caused, whether due to negligence or otherwise arising from the use of, or reliance on, the information provided directly or indirectly in this document. Do not tumble dry. This document may, or may not, set fire to your datacentre. The opinions expressed in this document were made at a particular point in time based on information available at that time, and we reserve the right to change our minds as new information comes to hand. If you actually bothered to read this far, send us a note at [ilovelegalese@pivotnine.com](mailto:ilovelegalese@pivotnine.com). PivotNine, the Corner P Device, and the PivotNine Ribbon are trademarks of PivotNine Pty Ltd. All other trademarks are the property of their respective owners.