aws

Planning large VDI migrations to the AWS Cloud

AWS Prescriptive Guidance

AWS Prescriptive Guidance: Planning large VDI migrations to the AWS Cloud

Table of Contents

Planning large VDI migrations to the AWS Cloud

Martin Guenthner, Amazon Web Services (AWS)

December 2023 (document history)

In the migration space, large-scale migrations of applications are a well-established pattern. Many organizations also need to migrate virtual desktop infrastructures (VDIs). However, because VDIs are end-user facing and have special requirements, in terms of availability, accessibility, and provisioning, this type of migration requires special considerations and planning compared to typical large-scale migrations.

This guide describes multiple architectural decision points for planning a large-scale migration of on-premises VDIs to the AWS Cloud. It also includes some additional recommendations for migrating from a Citrix environment, but these recommendations can also apply to other source technology stacks.

This guide helps you plan a large VDI migration to the AWS Cloud by providing a checklist of decisions that you'll need to make during the migration. As you complete the checklist, you'll be making informed decisions about the migration and the target architecture, which includes Amazon Elastic Compute Cloud (Amazon EC2) instances and Amazon FSx file shares.

The technical and non-technical considerations in this guide can also be used to evaluate alternatives to traditional VDI solutions, such as Amazon WorkSpaces or Amazon WorkSpaces Applications. As part of your migration planning, we recommend that you do a feasibility assessment to consider migrating to these services.

Intended audience

This guide is intended for VDI platform owners who are considering a platform migration to AWS, storage architects who support the migration approach, and compliance and risk teams who guide the decision-making process for reliability requirements.

Licensing considerations for large VDI migrations

When planning a large VDI migration to AWS, validate your licensing agreements for the operating system and software that you plan to use on the target VDI.

If you're planning to use your own license for the operating system, give special consideration to the potential impacts this can have. Under standard terms, Microsoft requires that you deploy Windows on hardware that is dedicated to your organization, such as a dedicated Amazon Elastic Compute Cloud (Amazon EC2) instance. Similar licensing agreements can apply to Microsoft Office or Microsoft 365.

If your organization has customized enterprise licensing agreements, get confirmation from the licensing team. Understanding your licensing agreements is vital to the successful planning of the migration.

Although bringing your own license might be more cost effective from a license perspective, it might not be cost-effective from a compute perspective. If dedicated hosts would be required for compliance with your licensing agreements, we recommend that you conduct a cost-benefit analysis.

For Microsoft operating systems and software, we recommend that you take advantage of an AWS Optimization and Licensing Assessment (AWS OLA) to assess on-premises or cloud workloads. AWS OLA makes optimized suggestions for the right EC2 instances for your workloads, but it also reviews your Microsoft licensing position. AWS OLA can help you reduce compute and licensing costs.

For more information about Microsoft licensing requirements and how to request a copy of your Microsoft License Statement (MLS), see Microsoft licensing on AWS in AWS Prescriptive Guidance.

Large migration considerations for VDI operating modes

You can deploy VDIs in either persistent or non-persistent mode:

- In *persistent* mode, a user can connect to the same virtual desktop every time. Changes from a previous session are saved for the next time the user connects. Persistent VDIs are typically faster to deploy, but you must manage patching and updating, and they require at least some storage capacity.
- In *non-persistent* mode, the desktops are generic and identical, and changes are not typically saved for the next session. Non-persistent VDIs can be simpler to manage, more secure, and more cost-effective. Non-persistent VDIs are well suited for those who don't require personalization and who perform a routine, limited set of tasks.

In addition to the general pros and cons for these two deployment options, a few additional considerations apply when moving the VDI solution to the cloud.

The following table describes the benefits and drawbacks of non-persistent VDIs in the cloud.

Benefits	Drawbacks
• *Simplified management* – Use Amazon Machine Images (AMIs) to manage VDIs in order to roll out changes to your entire environment. • *Improved availability* – You can deploy non-persistent VDIs in multiple Availability Zones, which allows users to access a VDI even when a specific Availability Zone is not accessible. • *Storage optimization* – You can reduce the storage requirements for non-persistent VDIs. User personalization layers can be moved to shared storage. Also, there is no requirement to back up VDIs.	• *Complex personalization* – User profiles and applications have to be offloaded to shared storage devices. This can increase implementation effort, operational complexity, and scaling challenges.

Benefits	Drawbacks
• *Flexible capacity* – Non-persistent VDIs do not require a one-to-one mapping of users to instances.	

The following table describes the benefits and drawbacks of persistent VDIs in the cloud.

Benefits	Drawbacks
• *Easy personalization* – Store user profiles and applications on the Amazon Elastic Block Store (Amazon EBS) volume for the instance. There is no requirement for shared storage. • *User customization* – The instance type and storage capacity can be adjusted on a per-user level.	• *Backup complexity* – Store user profiles and applications on the Amazon Elastic Block Store (Amazon EBS) volume for the instance. There is no requirement for shared storage. • *Reduced availability* – Users are one-to-on e mapped to instances and the correspon ding Availability Zone. An Availability Zone failure results in the service being unavailab le for the affected users.

Storage considerations for large VDI migrations

User profiles or virtual hard disk (VHD) files for Citrix user personalization layer (UPL) will require some form of Server Message Block (SMB) file storage, ideally in close proximity to the VDIs. You can use Amazon FSx for Windows File Server to address this requirement.

Careful capacity planning and performance testing is required when designing an Amazon FSx architecture. The design should cater for peak load, and we highly recommend running full-scale performance tests.

Splitting UPL and profiles over multiple Amazon FSx file shares is also recommended. This helps you manage the different usage patterns that both solutions have on a file share. You can group multiple file shares into a common namespace by using Distributed File System (DFS) Namespaces.

Depending on your requirements for performance and storage capacity, distributing UPL and profile servers over multiple Amazon FSx file shares might be required.

When designing and testing a storage solution, consider how the application behaves if you lose access to storage. For example, if Amazon FSx fails over to the secondary host due to a reconfiguration, how will the VDI application react to the temporary loss of connectivity? This can disconnect all sessions, and a reconnect storm can add additional heavy load on the system. If possible, the system should be configured to minimize the effects of this behavior.

Instance capacity planning for large VDI migrations

VDI solutions tend to require the deployment of hundreds or thousands of EC2 instances. To avoid insufficient instance capacity errors, we recommend that you develop a rollout plan and work with AWS Support or your AWS Technical Account Manager to coordinate the capacity requirements with the AWS service team.

Capacity planning is even more important if dedicated instances or dedicated hosts are required. Consider the following when planning instance capacity requirements:

- Concurrent instances required
- Instance families and types
- On-Demand Instances, On-Demand Capacity Reservations (ODCR), or Reserved Instances (RI) capacity
- Cost implications

Capacity reservation considerations

When developing a capacity reservation model, consider and answer the following:

- How much capacity is required for concurrent users?
- Does the solution have to cater for unavailability of a specific Availability Zone or instance types?
- Is reserved capacity required to meet business requirements?

To answer these questions, you can use the following matrix. The matrix includes an example of how you might distribute instance types for a total of 5,000 VDI users.

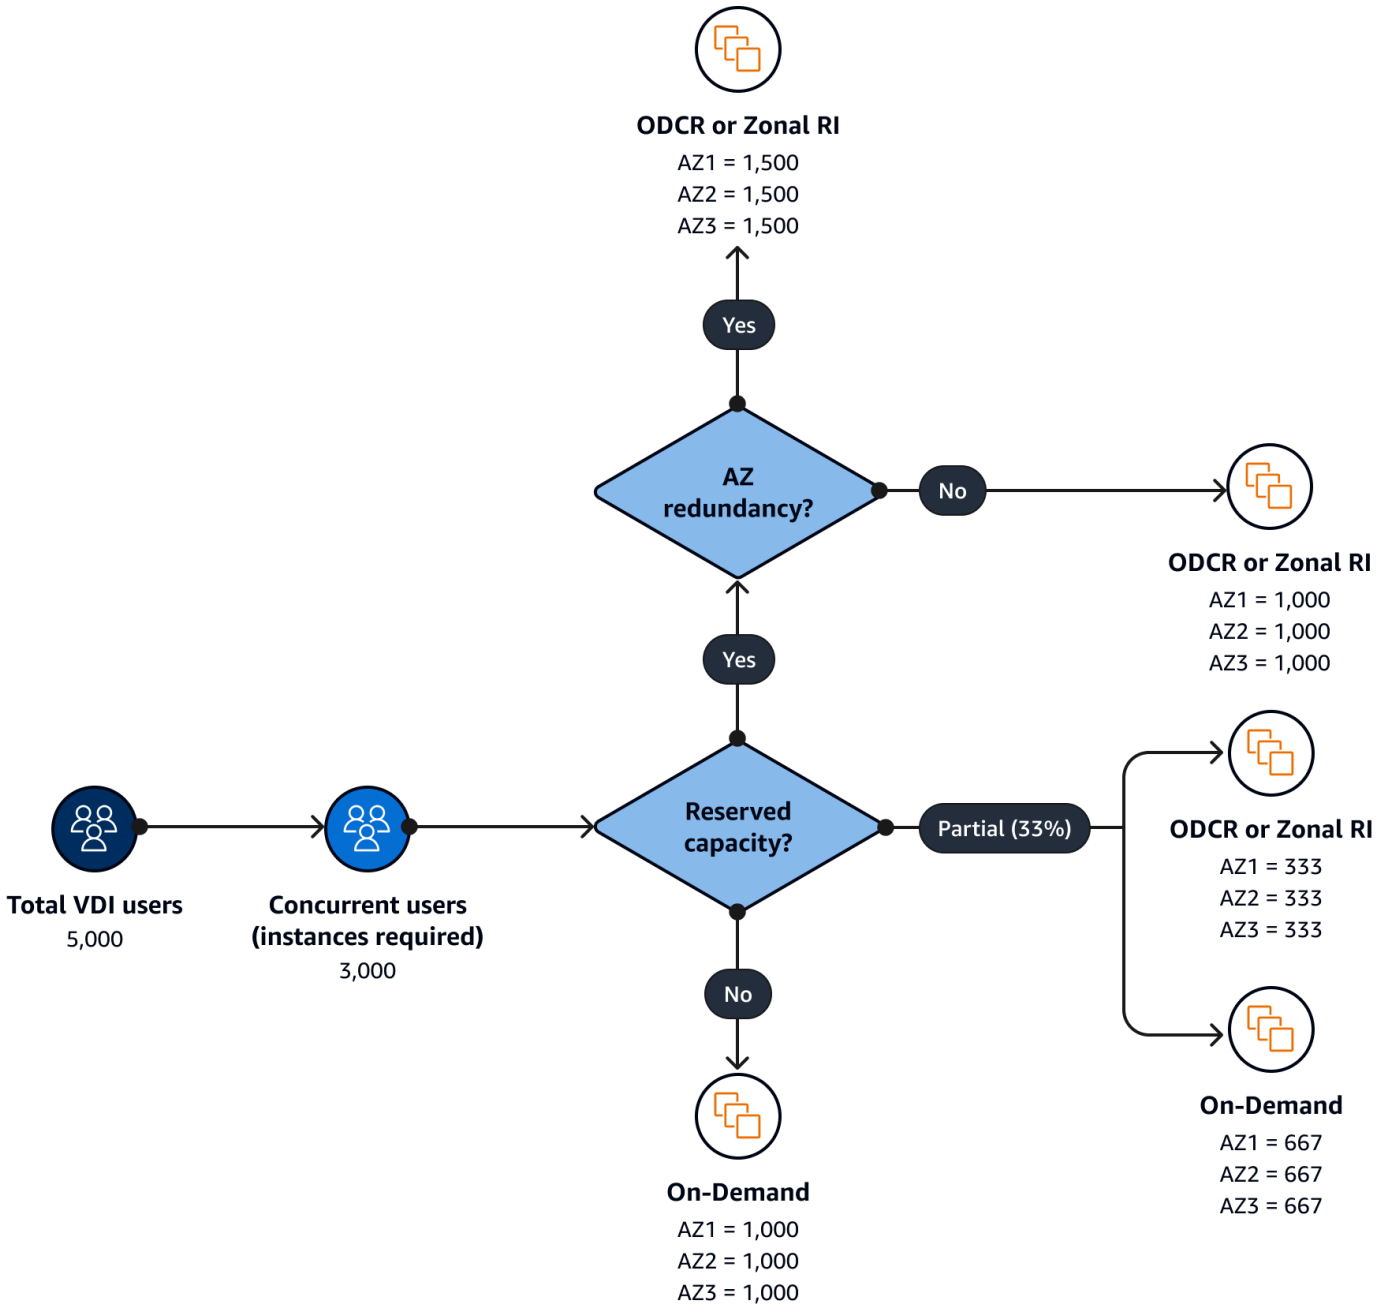
ODCR or Zonal RI
AZ1 = 1,500
AZ2 = 1,500
AZ3 = 1,500
Yes
AZ redundancy?
No
ODCR or Zonal RI
AZ1 = 1,000
AZ2 = 1,000
AZ3 = 1,000
Yes
Total VDI users
5,000
Concurrent users (instances required)
3,000
Reserved capacity?
Partial (33%)
ODCR or Zonal RI
AZ1 = 333
AZ2 = 333
AZ3 = 333
No
On-Demand
AZ1 = 667
AZ2 = 667
AZ3 = 667
On-Demand
AZ1 = 1,000
AZ2 = 1,000
AZ3 = 1,000

Performance testing and metrics for large VDI migrations

Performance testing an application is an important task for any deployment, and the same concept applies to VDI deployments. To understand the performance requirements, you must plan and define metrics and acceptance criteria. These metrics can either be captured from the current VDI environment, or you can define new service-level agreements (SLAs) for your cloud environment.

Consider the following metrics:

- Login time
- Profile loading time
- Network latency to on premises
- Local disk performance, such as to support local files or an application cache
- Network throughput, such as making sure that the maximum network capacity is defined by the chosen instance size

When conducting performance tests, we recommend that you conduct full-scale tests and replicate the size of an average production deployment. Performing performance tests on a small scale can lead to inaccurate results because the impact on shared resources, such as file servers for user profiles or malware signature updates, can be significantly lower.

Network capacity planning for large VDI migrations

Large-scale VDI deployments can add a significant amount of network traffic to resources deployed in the AWS Cloud and on premises. This includes traffic generated by:

- Microsoft Active Directory
- Logging and monitoring
- Application traffic
- Operating system (OS) and software patching
- Access to file servers
- Remote connections from VDI to other resources
- Email and internet activities from VDI users

Depending on your network configuration, some or all of this traffic can route back to on premises. This adds networking load to AWS Direct Connect or VPN connections, firewalls, and other network equipment in the path.

Define a base-load network capacity profile for a VDI. Work with the networking team to develop a capacity plan. This helps you avoid any bottlenecks or other negative effects to network traffic to other applications over the same cloud connection.

Repeating spike events, such as patch day or malware signature updates, should be taken into consideration. If possible, evaluate a rolling deployment for these events.

Resources

AWS resources

- Citrix Desktop as a Service (DaaS) on AWS
- Microsoft Licensing on AWS
- Large migrations to the AWS Cloud

Citrix resources

- Citrix - User personalization layer

Microsoft resources

- What is FSLogix
- Configuration examples for FSLogix

AWS Partner resources

- Accenture

Document history

The following table describes significant changes to this guide. If you want to be notified about future updates, you can subscribe to an RSS feed.

Change	Description	Date
Initial publication	—	December 14, 2023