



RiverMeadow SaaS: Easing the Migration Path to AWS

Author: John Merryman, Vice President @ RiverMeadow Software

Date: 10/11/16

This whitepaper outlines the architecture, applicability, and use-cases for Lift & Shift migrations to AWS VPC environments. The intended audience includes technical directors, managers, architects, administrators, and leadership interested in streamlining Windows and Linux servers into AWS without risk to your business. Our goal is to accelerate your application and workload transitions into AWS, with predictable and consistent automation.

TABLE OF CONTENTS

EXECUTIVE SUMMARY	2
OBJECTIVE	3
RIVERMEADOW SAAS ARCHITECTURE AND WORKFLOW	3
SYSTEMS ARCHITECTURE	3
NETWORKING AND SECURITY.....	8
SOURCE SERVER READINESS	9
WORKFLOW OVERVIEW.....	13
PERFORMING LIFT & SHIFT MIGRATIONS TO AWS	14
MIGRATION LIFECYCLE.....	14
USE CASE #1: CLASSIC FULL + DIFFERENTIAL MIGRATION	18
USE CASE #2: LEGACY PLATFORM MIGRATIONS.....	18
USE CASE #3: FULL MIGRATION WITH CREATE AMI	19
USE CASE #4: AWS CROSS REGION MIGRATIONS	19
PARTING THOUGHTS	20

LIST OF FIGURES

Figure 1 RiverMeadow SaaS Architecture.....	5
Figure 2 Networking Requirements	9
Figure 3 RiverMeadow SaaS Migration Flow.....	13
Figure 4 Classic Full + Differential Migration	18
Figure 5 Legacy Platform Migrations.....	19
Figure 6 Full Migration with Create AMI	19
Figure 7 AWS Cross Region Migrations	20

LIST OF TABLES

Table 1 RiverMeadow SaaS Components.....	5
Table 2 Windows Server Requirements.....	10
Table 3 Linux Server Requirements.....	11
Table 4 Where RiverMeadow Fits in the Migration Lifecycle	14

Easing the Move into AWS

Executive Summary

Deciding to move into the cloud is easy. But, too often, getting there often proves to be a harrowing process. Typically, a large-scale migration is a combination of brute-force, manual tools, a mixed bag of niche 3rd party tools, and component application/database transformation efforts. And the stakes are high, whether or not the transformation runs smoothly.

A migration rollback impacts business processes and revenue, and all companies face duplicate operating time costs during a transition to the cloud. Amazon® recently published 'A Practical Guide to Cloud Migration' whitepaper¹, illustrating the 'migration bubble' of this cost delta during transition to the cloud.

This cost and business risk exponentially increases should a project run long or off the rails. The classic “migration bubble” problem is reinforced by a general belief that migration is technically difficult, operationally risky and financially expensive, so organizations tend to rely on third party consultants or apply their most expert resources to the task.

The goal of RiverMeadow SaaS is to change this paradigm to one where migration is simple, cheap, and virtually zero risk, enabling organizations to utilize much lower cost resources in a “migration factory” model.

For enterprise customers with an AWS® migration project, RiverMeadow's core business is designed to alleviate the stress of migrating Microsoft® Windows Servers® and Linux® servers, while introducing an AWS focused user-experience to deliver migrations to AWS VPC environments. RiverMeadow Cloud Migration SaaS takes a Lift & Shift approach, by moving the operating system instances and block data storage 'as-is' to the AWS VPC target.

¹ A Practical Guide to Cloud Migration - Migrating Services to AWS
<https://d0.awsstatic.com/whitepapers/the-path-to-the-cloud-dec2015.pdf>

This whitepaper highlights the following technical differentiators for the RiverMeadow SaaS Migration platform:

- No agent installation required on source, no appliance in source data center
- Source server instances remain live and in production throughout the migration process
- We duplicate the source server and spin up a clone in your AWS destination environment
- Data is copied from the source to the clone by running AWS instances
- Sync over the differences, and cutover when you're ready
- Even in a worst-case scenario, you won't lose any sleep. Your source server IS your rollback plan.
 - *Because RiverMeadow SaaS performs a live clone and copy, rollback plan execution is unlikely since target workloads are entirely insulated from source production.*

Objective

This whitepaper outlines the architecture, applicability, and use-cases for Lift & Shift migrations to AWS VPC environments.

The intended audience includes technical directors, managers, architects, administrators, and leadership interested in streamlining Microsoft Windows and Linux servers into AWS without risk to your business.

Our goal is to accelerate your application and workload transitions into AWS, with predictable and consistent automation.

RiverMeadow SaaS Architecture and Workflow

Systems Architecture

The RiverMeadow SaaS design for AWS embraces two fundamental factors:

- 1) Leverage the IP assets, but with a green field design process and
- 2) Deeply engage AWS services and user-experience wherever possible, to generate an AWS 'native' migration experience.

As a result, the design and implementation of this product met these goals yet adhered to the following fundamental platform characteristics:

- No installation of an agent on the source server instance
- No required access to source hypervisor or director/API
- True Lift & Shift of Operating System to target cloud, only changing underlying drivers without source/target dependent OS templates
- Data path is exclusively between the source and target workload instances, without any middle-path hops or staging appliances
- API driven coordination of target instance provisioning and migration
- Expose AWS native features to enrich the migration process, including instance tags, security groups, IAM roles, variable EC2 instance type, AMI creation, and automation with Cloud Formation.
- Provide an ability for the user to extend the migration automation framework with custom scripts/functions

Designed and purpose-built for AWS enterprise customers, RiverMeadow SaaS is a feature-rich and agentless AWS migration tool that delivers high performance, efficient and secure server migrations into AWS.

The following diagram provides an overview of the RiverMeadow SaaS systems architecture.

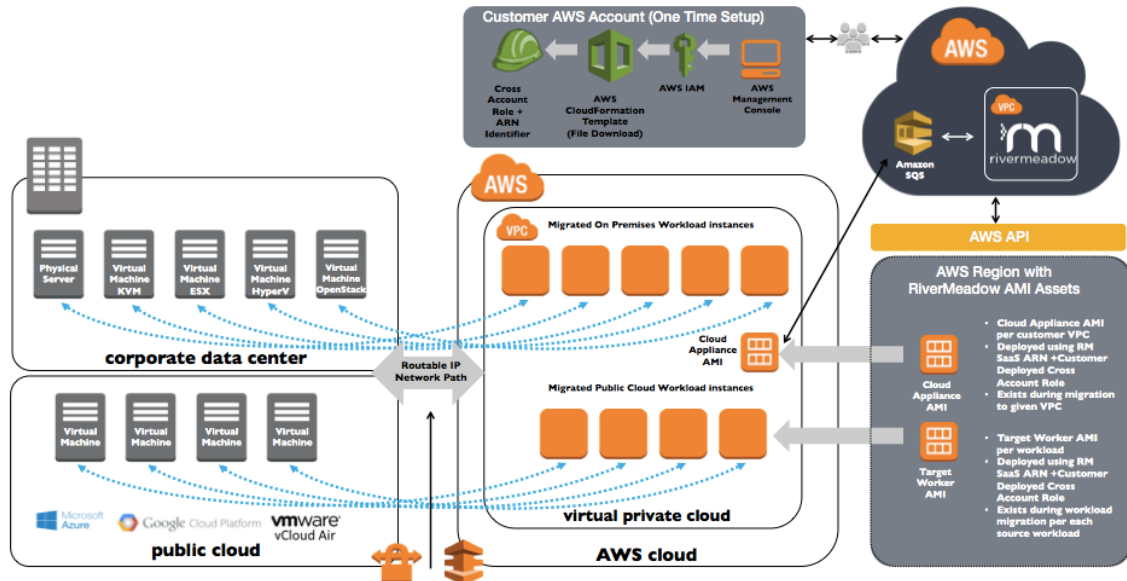


Figure 1 RiverMeadow SaaS Architecture

The components of RiverMeadow SaaS architecture are outlined in the following table:

Table 1 RiverMeadow SaaS Components

Component	Description	Component Interactions
RiverMeadow SaaS	<p>The RiverMeadow SaaS environment is a multi-tier application, hosted in AWS. This global PoP is highly redundant and fault-tolerant, with security design and controls adhering to AWS Well Architected² framework specifications.</p> <p>Its primary function is to authenticate the user onto the RiverMeadow platform, provide</p>	<p>RiverMeadow SaaS provides the web user interface, API services, and communicates to the Cloud Appliance deployed in the customer VPC.</p> <p>Metadata is transmitted to the RiverMeadow SaaS environment, however no application/workload</p>

² AWS Well Architected Framework
https://d0.awsstatic.com/whitepapers/architecture/AWS_Well-Architected_Framework.pdf

	the web UI, and interactively build the rule set for a given migration.	data is transmitted to or stored in this environment.
Cloud Appliance	<p>The Cloud Appliance is a virtual machine instance deployed into the customer VPC environment to control metadata flow between solution components, and to broker metadata between the customer VPC environment and RiverMeadow SaaS.</p> <p>The Cloud Appliance is automatically deployed into the customer account from an AWS AMI (Amazon Machine Image), hosted in each region supported by AWS.</p> <p>Customers are provided an AWS CloudFormation® template to run, which uses IAM to create a cross-account role for RiverMeadow services and an associated ARN identifier used in initial account setup. RiverMeadow is designed with direct compliance with AWS SaaS cross account role standards.³</p> <p>Deployment of the Cloud Appliance into the customer VPC is completed within 5 minutes of applying the IAM cross account role ARN string into the</p>	<p>The Cloud Appliance communicates with each source workload to collect source attributes, coordinate migration orchestration, and broker migration metadata to the RiverMeadow SaaS environment.</p> <p>Note: The Cloud Appliance is only required to run in the VPC environment during migrations. Following migration projects, this instance can be deleted without any impact to the source or target VPC production environments.</p>

³ Securely Accessing Customer AWS Accounts with Cross-Account IAM Roles

<https://aws.amazon.com/blogs/apn/securely-accessing-customer-aws-accounts-with-cross-account-iam-roles/>

	RiverMeadow interface.	
Target Worker	<p>The Target Worker is an EC2 instance deployed from regionally hosted AWS AMIs, into the customer VPC environment, specifically to serve as the landing zone for a specific source instance live clone and copy. The Target Worker image is a stripped down operating system image with an ephemeral mount point, which is also hosted in each region supported by AWS.</p> <p>Like the Cloud Appliance, the Target Worker is deployed into the customer VPC environment using the cross account role created during account setup.</p> <p>Note: RiverMeadow SaaS associates a temporary security group with the Target Worker instance during migration. This is automatically replaced with the user's chosen security groups upon completion of the migration.</p>	<p>The Target Worker communicates with the source workload and is the destination for all data copy from the source. There is also interaction with the Cloud Appliance for metadata control and transfer, however the workload copy data path is exclusively maintained between a given source workload and the Target Worker instance.</p> <p>The best way to think about the work being performed is that RiverMeadow SaaS is distributed systems architecture, with the entirety of the workload migration CPU/Memory/IO 'work' being born between a given source and a given Target Worker Instance. Everything else is orchestration and metadata handling. Ultimately, the target work becomes the clone, so there is no cleanup/removal of this asset required post migration.</p>

Networking and Security

Because RiverMeadow Cloud Migration SaaS is a distributed system that installs no agent on the source, it relies on network paths to deliver migrations. The network and security design tenets include:

1. Keep the customer data, and optionally source instance credentials, behind the perimeter, using routable path between source environment and target VPC
2. Automate security group rules for Cloud Appliance and Target Worker
3. Accommodate variable customer VPC scenarios (private, public facing subnets, peered VPCs, segregated VPC environments)
4. Allow customers to manage source credentials on the Cloud Appliance if required
5. Keep it simple to setup and operate from a networking perspective

RiverMeadow complies with AWS Security Best Practices⁴ and allows customers to retain perimeter controls for source, target, and migration networks. The following diagram illustrates the required network firewall directional rules.

⁴ AWS Security Best Practices

<https://aws.amazon.com/whitepapers/aws-security-best-practices/>

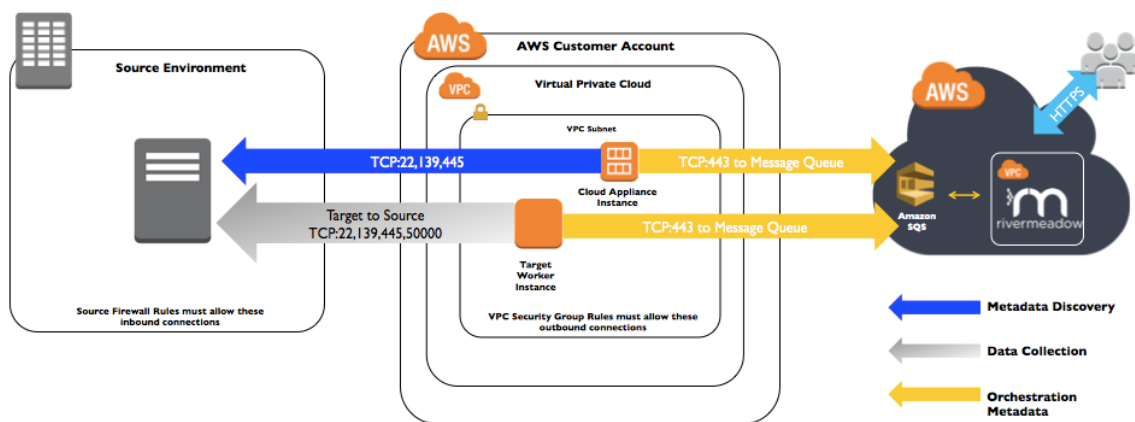


Figure 2 Networking Requirements

Note: The Cloud Appliance outbound Internet traffic is purely for orchestration metadata. The northbound traffic is routed over TCP port 443 (HTTPS) to a proxy IP address in AWS. This allows for whitelisting of a distinct IP for outbound metadata traffic. Custom proxy configurations are also permitted. RiverMeadow SaaS is “air gapped” from the Cloud Appliance and Target Workers in a client VPC – there is no direct access to these assets, thereby ensuring the security of client credentials and data.

Migrations can be run over any routable path source and target VPC following these network rules. We leave the network design, security perimeter, and VPC security group design in the hands of the customer to accommodate the wide-ranging networking scenarios. Whether this is a VPN, MPLS circuit, or AWS Direct Connect® is up to the customer. Even the public Internet can also be used for migrations, but is not advised for production or sensitive workloads.

Source Server Readiness

RiverMeadow SaaS uses native operating system commands and services to perform migrations from the source. This means that for each supported operating system, we have a minimum set of requirements to run full and differential (sync) migrations. RiverMeadow automatically checks these

requirements and will report back with warnings when required dependencies are not fully met. In the event a critical requirement is not met, RiverMeadow will not allow the migration to begin.

Table 2 outlines the requirements for Windows and Linux.

Table 2 Windows Server Requirements

Windows Source Requirements	Details
WMI must be enabled	RiverMeadow uses native Windows capabilities to allow migrations to happen without agent installation. WMI is used to collect metadata about the server.
Administrator Account	User needs local or domain administrator privileges on the source server.
User Account Control (UAC) Settings set to "Never Notify"	UAC permissions are required to prevent Windows from triggering the command approval UI to the console. Our commands & scripts will halt until the UI prompt is answered.
Volume Shadow Service (VSS) must be enabled (only for differential block migrations)	In order to create a consistent replica, RiverMeadow uses VSS snapshots, which must be able to work unimpeded by other configurations and software. VSS error logs must be clear of errors. .
OS must be installed inside Windows folder on the C:\ drive	Complex and custom Windows setups can cause unforeseen issues in various target environments.
Host based Anti-virus & Anti-malware must allow RiverMeadow SaaS to connect and execute processes	We recommend disabling if they are causing problems. If this is not possible, then contact RiverMeadow Support. Getting anti-virus correctly configured to allow our process is sometimes required.
Disk	A minimum of 10% free disk space is needed NTFS, Dynamic Disks Supported
CPU	The migration process utilizes up to 15% CPU usage
RAM	A minimum of 2GB RAM

.NET framework 3.5 (AWS requirement)	.NET 3.5 is required for Windows migrations into AWS
---	--

Table 3 Linux Server Requirements

Linux Source Requirements	Details
ssh server must be installed and PermitRootLogin enabled Note: 1. For migrations into AWS, sftpd <u>must be enabled</u> 2. PermitRootLogin <u>not required</u> when using sudo privileges.	RiverMeadow uses native Linux capabilities to allow migrations to happen without agent install. If direct ssh login has been turned off, it must be enabled before initiating a migration.
Local Root Account (or sudo privileges)	In order for RiverMeadow to have access to all of the data, the account used must also have access to all of the data. RiverMeadow supports non-root credentials. Non-root user must have <u>NOPASSWD sudo access</u> and be allowed to run sudo <u>without a tty</u> .
Native Tools	cat ls uname lsb_release hostname ifconfig which lvs pvs

	vg blkid parted tar sshd sftpd
Host based Anti-virus & Anti-malware must allow RiverMeadow SaaS to connect and execute processes	Your Anti-virus System Administrator should contact RiverMeadow Support if these services are causing migration problems
Disk	1GB free disk space Ext3, Ext4, XFS file systems supported; Servers with LVM layout must have LVM tools installed
CPU	The migration process utilizes up to 15% CPU usage
RAM	2GB minimum
rsync and sshpass (version 1.05+) installed	rsync and sshpass (version 1.05+) must be present on the source (prior to the initial, full migration) and target machine in order to perform a Differential Migration.

Full specifications, including limitations are available in our [support matrix](#). For large-scale migrations, the above prerequisites can be tested and satisfied through systems administration updates across standard administration/automation toolsets (puppet, sccm, Jenkins, chef, salt, ansible, etc.) or manually without invoking source downtime.

Workflow Overview

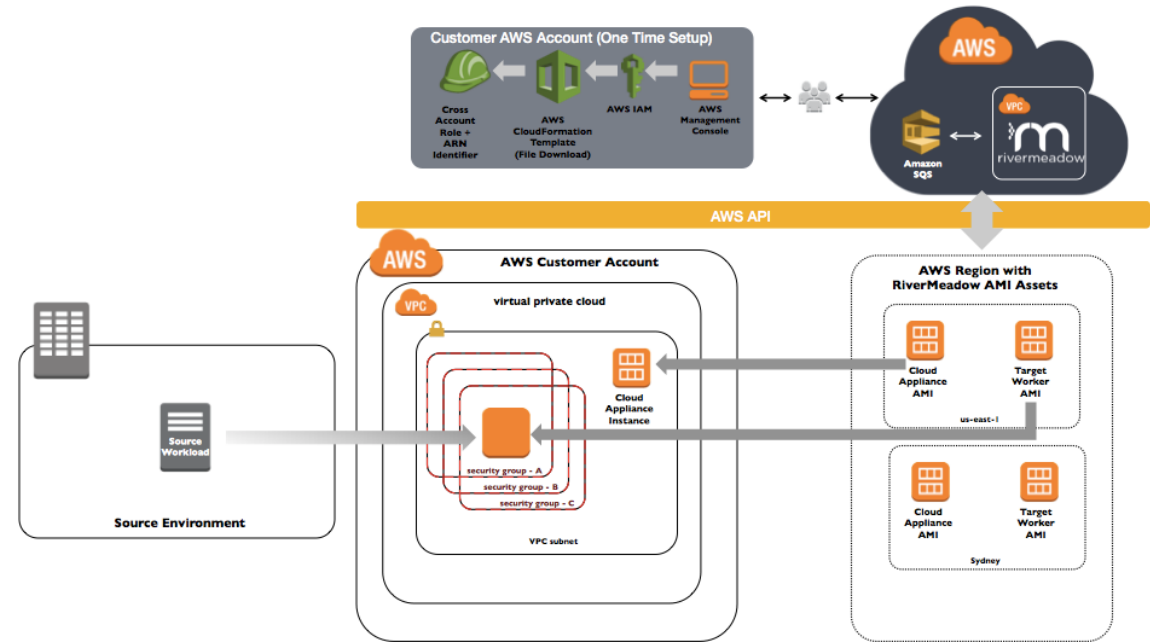


Figure 3 RiverMeadow SaaS Migration Flow

The workflow as depicted in the illustration above:

- 1) Run the Cloud Formation template in your AWS console to create the cross account role and unique ARN string linking RiverMeadow to your account
- 2) Input the ARN string into the RiverMeadow SaaS portal to setup your account and automatically provision the Cloud Appliance into your VPC
- 3) Ensure VPC configuration and source side firewall rules meet specifications
- 4) Ensure source servers satisfy prerequisites outlined above (also programmatically checked during migration pre-flight processing)
- 5) Start migrating source workloads to your VPC environment

RiverMeadow SaaS satisfies technical compliance to AWS Migration Competency Requirements.⁵

⁵ AWS Migration Competency Requirements
https://s3.amazonaws.com/Competencies/Validation+Checklists/AWS_Migration_Consulting_Partner_Validation_Checklist_v.1_June+2016.pdf

Performing Lift & Shift Migrations to AWS

This final section outlines how RiverMeadow Cloud Migration SaaS fits into the migration lifecycle, and how it works across a number of field use-cases.

Migration Lifecycle

RiverMeadow recognizes the complexity of cloud migration; it is an organizationally and technologically intricate series of activities. Often these projects span months to years, and require concentrated efforts from both application and infrastructure teams, not to mention governance and sanctioning from leadership up to the board level in many organizations.

RiverMeadow Cloud Migration SaaS fits into the overall migration lifecycle as an important component, with clear dependencies to be satisfied in advance. The following table provides a generalized view of the migration lifecycle and the relationship to RiverMeadow SaaS technology is *annotated herein*.

Table 4 Where RiverMeadow Fits in the Migration Lifecycle

Lifecycle Component	RiverMeadow SaaS Relationship & Context
Discovery	<p>Migration programs often start with strategy or design, and sometimes revert to a discovery process to fully understand the current applications portfolio and infrastructure estate.</p> <p>Discovery involves a combination of physical, logical, virtual estate discovery, coupled with application portfolio discovery, and interdependency analysis. RiverMeadow SaaS accepts inputs from a range of discovery tools.</p> <p><i>Discovery serves as a direct input to RiverMeadow SaaS for small-scale migrations, and as a programmatic input to migration planning for large-scale migrations.</i></p>
Strategy & Budgeting	<p>Ideally strategy and program budgeting is performed early in the migration lifecycle, however a key input is a valid understanding of the current state and a general concept of the target state design/outcome.</p>

	<p><i>As an indicative costing input to the Strategy and Budgeting process, RiverMeadow SaaS costs can be generally estimated at \$300/workload, not including services or time to manage the migration process. The cost goes lower with consumption on scale, but for budgeting purposes this value works.</i></p>
Migration Planning	<p>Taking inputs from discovery, the migration planning process creates, validates, and manages the application move-groups from concept through to execution. Application move groups are logically grouped applications and associated workloads/server instances, which can be moved to the target VPC at a given time. Classic move-group constraints include application downtime tolerance, change windows, and application/infrastructure team availability and bandwidth to perform testing, validation, rollback, and cutover procedures.</p> <p><i>Migration Planning is a key input for RiverMeadow SaaS migrations. Move groups can be loaded as logical groupings of sources, then manually invoked or scheduled in RiverMeadow SaaS as a migration plan.</i></p> <p><i>While RiverMeadow SaaS does not eliminate the human planning/acceptance factor, a number of natural friction points in a migration effort are reduced by the fact that RiverMeadow SaaS performs a live workload clone and copy, without introducing agents or invasive change procedures to the source production environment.</i></p>
Target VPC Design	<p>The target VPC design and implementation work must be completed prior to performing migrations. While RiverMeadow SaaS can accommodate a range of VPC design and complexity scenarios, this implementation effort must be completed prior to performing migrations.</p> <p><i>Target VPC design and implementation is a key dependency to be satisfied prior to running migrations</i></p>

	<i>with RiverMeadow SaaS.</i>
Migration	<p>The main event. For most large-scale programs, migration includes a mix of transformational, non-transformational, simple to medium complexity, and high-complexity migrations. Some are performed at the application layer, others at the workload and operating system instance layer.</p> <p><i>RiverMeadow SaaS addresses requirements for Windows and Linux 'Lift and Shift' migrations to VPC environments. (Not to be confused with traditional data center Lift and Shift migrations moving physical devices)</i></p>
Testing & Validation	<p>Testing and validation is performed at different levels in the stack following the migration event. Infrastructure teams initially validate the server instance/integrity via systems acceptance testing. Assuming no defects require remediation or rollback, systems and applications are handed over to applications teams for testing and validation. This process can require periodic data synchronization events to update data sources and can vary widely from one application to the next. In addition, infrastructure teams may perform OS level updates to production-ready workloads for the target cloud environment, often including the installation of systems management tools, and numerous manual procedures on the target workload.</p> <p><i>RiverMeadow SaaS introduces a range of options to support migration testing scenarios:</i></p> <ul style="list-style-type: none"> <i>• Performs a live clone and copy; rollback plan execution is unlikely since target workloads are entirely insulated from source production.</i> <i>• Unlimited differential migrations can be performed for a source workload following the initial full migration; For Windows this can be volume/block level leveraging VSS or file-level differential as applied to selected drives/volumes. For Linux this is file-level differential as applied to selected mount</i>

	<p>points/volumes.</p> <ul style="list-style-type: none"> • <i>An optional AMI (Amazon Machine Image) can be created from the RiverMeadow SaaS migration workflow; which can effectively be cloned, copied, deployed as many times as needed for testing purposes.</i> • <i>Instance Operating System level changes can either be preserved prior to cutover. For instance, if you run differential file-level updates to non-OS drives only.</i> • <i>Instance Operating System level changes can also be overwritten prior to cutover for a fully consistent update of the source data changes.</i> • <i>Post-migration routines can be tested and scripted, then uploaded to RiverMeadow SaaS and executed as migration extensions applied to final differential migration processing.</i>
Production Cutover	<p>Production cutover typically involves careful coordination with applications and infrastructure teams to quiesce source production, update DNS, LDAP, and A/D entries, and cutover production to the target VPC environment.</p> <p><i>For testing, any number of options can be exercised via RiverMeadow SaaS prior to cutover. The platform is intentionally designed to leave network-impacting changes (such as DNS updates) in the hands of the customer migration team, to be exercised in concert with final full or differential migration events.</i></p>
Decommissioning	<p>Often an afterthought, work is required to decommission physical and virtual infrastructure in the source environment. This exercise starts with orderly shutdown of VM instances, VM hypervisors, physical machines, cables, networking and storage, all the way to the raised floor.</p> <p><i>While RM SaaS cannot complete the decommissioning process, it does provide a source shutdown option as part of final differential migration processing (to be used with</i></p>

caution).

Use Case #1: Classic Full + Differential Migration

Our standard use-case fits the vast majority of migration patterns for migrating server workloads into AWS VPC environments. Following the initial full migration, block and/or file level differential migrations are run to satisfy testing requirements, prior to production cutover to the target instance.

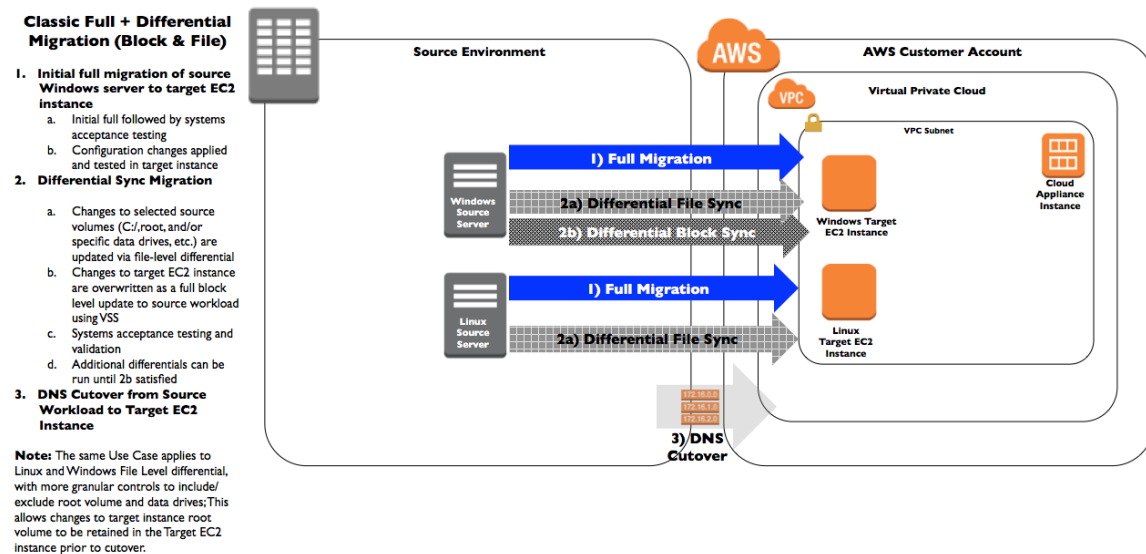


Figure 4 Classic Full + Differential Migration

Use Case #2: Legacy Platform Migrations

Legacy Microsoft Windows server migration to AWS is supported with a variety of options for testing, upgrade, and data synchronization. Following the initial full migration, block and/or file level differential migrations are run to satisfy testing requirements, prior to production cutover to the target instance.

Note: As an option, RiverMeadow has a community sourced Migration Extension to perform an in-place upgrade of the Windows 2003 server instance to Windows 2008 in the target VPC environment. Officially AWS does not provide support for Windows 2003 workloads, so the operational premise is 'run at your own risk' within AWS.

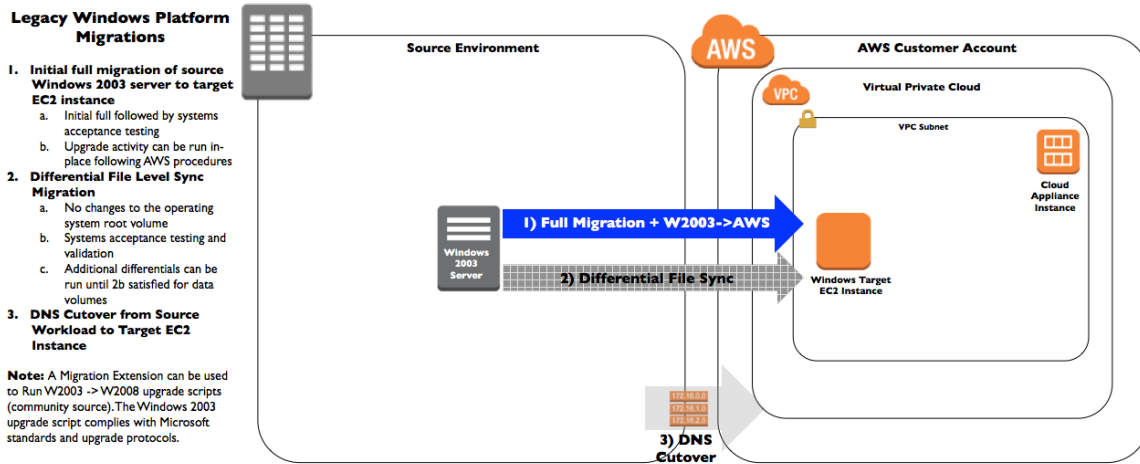


Figure 5 Legacy Platform Migrations

Use Case #3: Full Migration with Create AMI

RiverMeadow migration process creates an AMI instead of a cloned image, if this option is selected for a migration. This results in a static image which can then be used for iterative test copies or scale-out use-cases.

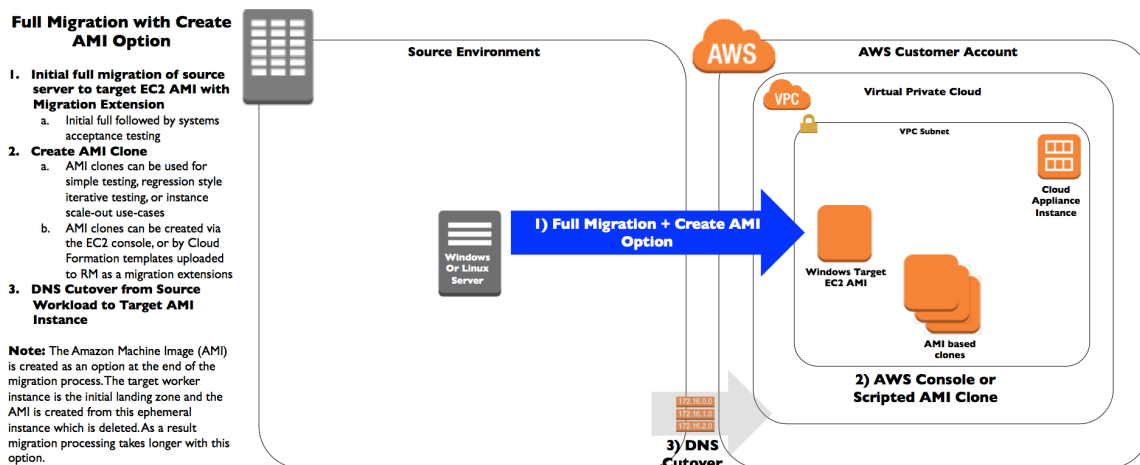


Figure 6 Full Migration with Create AMI

Use Case #4: AWS Cross Region Migrations

EC2® instances can be migrated in serial or batch between AWS regions. Traffic can be routed across VPN or Direct Connect circuits, since VPC peering

is not allowed between regions. This use case satisfies needs to transition applications and workloads due to regional availability, data sovereignty, or application architecture latency decisions.

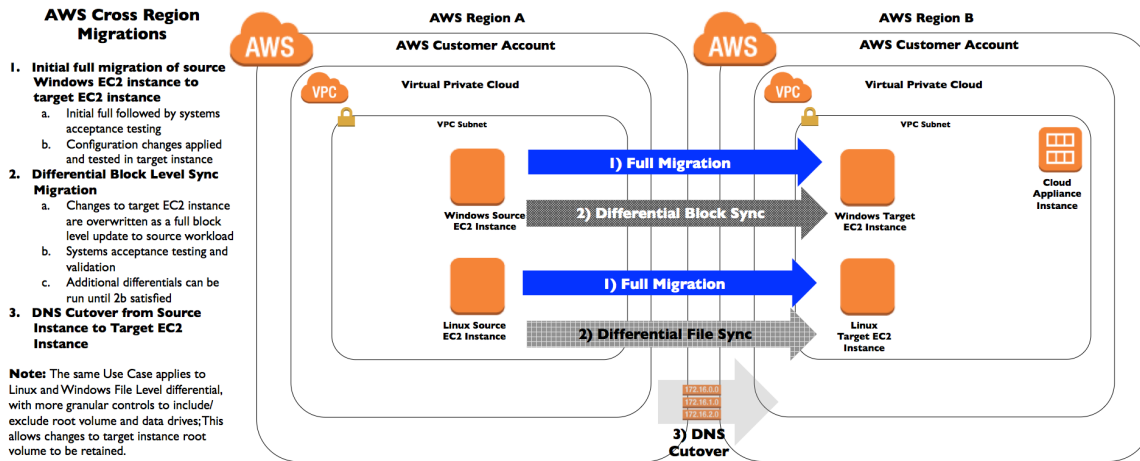


Figure 7 AWS Cross Region Migrations

Parting Thoughts

Leveraging RiverMeadow's Lift & Shift design mitigates the business risks associated with moving your servers. Migration automation plays an important role in transitioning customer applications and workloads to AWS, and we continue to evolve our platform to satisfy our customers' complex migration projects.

We recognize that each migration project introduces a unique range of complex tasks required to move servers from source to target. RiverMeadow SaaS's Lift & Shift migration methodology automates those complex tasks 50% of the time for large enterprise projects and up to 100% for projects that are modular.

This approach can accelerate your transition to AWS, and afford a consistent base from which applications and workloads can be further optimized, transformed, or redesigned to take advantage of AWS services. RiverMeadow offers ease of integration, testing, and innovation by working within the AWS platform.

We are interested in simplifying and accelerating your migration into AWS.
For questions or a technical review, please contact:
awsmigrations@rivermeadow.com.

RiverMeadow Cloud Migration SaaS benefits:

- **Minimal Change Management** requirements for source production
- **SaaS-based** - No software to install on source production environment
- **Lean Systems Architecture** - Native AMI Software Asset Distribution
- **Roles-based Security** - IAM Security, leverages AWS Security Groups
- **Live Server Data Cloning** - allows target workload test/validation before final cut over to AWS
- **Variable Differential Sync** - Block or file-based data synchronization
- **End to End Automation** - RESTful APIs allow automation of migration patterns on scale
- **Post-Migration Automation** - Create AMIs, run executable code, reduce human error

RiverMeadow Software, Inc. 2107 North First Street, Suite 660, San Jose, CA 95131.
Phone: (408) 217-6498 www.rivermeadow.com. Copyright © 2016 RiverMeadow Software Inc. All Rights Reserved. RIVERMEADOW is a trademark of RiverMeadow Software, Inc., Reg. U.S. Pat. & Tm. Off. All other trademarks are the property of their respective owners. For further information visit:
www.rivermeadow.com. This document is RiverMeadow Public. RM/AWS/WP/1016