Data Center Migration & Automation Whitepaper

Authored by: John Merryman

RiverMeadow Software

Table of Contents

1 Overview

Data Center Migration projects represent a massive segment of compute workload migration worldwide. The typical data center today is mostly standardized on x86 compute infrastructure, yet heterogeneous platform, storage, and network infrastructure typically populates the data center landscape. As the number of applications and disparate infrastructure types expand, the number of migration use-cases and associated complexity also increases. As such, many traditional migration methodologies are still highly relevant to data center migration and consolidation, while emerging technologies for migration automation have increasing relevance.

The following whitepaper seeks to address a range of factors associated with data center migration. Starting with business drivers, we investigate the typical triggers for a customer to engage in this activity. This is followed by an overview of relevant tools and technologies, ranging from discovery, planning, governance, to migration automation using RiverMeadow's SaaS solution.

Next we explore in depth each phase of a data center migration from a practitioner's point of view, including a walk through of specific migration methodologies (such as physical lift and shift, versus physical to virtual migrations). Lastly, we explore the relevance and benefits of applying migration automation to a data center migration project, and the associated technical and business benefits.

The content herein provides value to Infrastructure Directors, Managers, and Architects who have remit for execution of a data center migration project. The content applies to mid to large-scale migrations, starting with hundreds to thousands of data center assets in scope for migration. The benefit of reading this paper is to obtain a holistic view of data center migration from a practitioner's viewpoint and to learn the benefits of RiverMeadow migration automation technology for data center migrations.

2 Business Drivers

There are numerous reasons why organizations initiate a data center migration. The business drivers for migration must be identified and understood by all parties, as each individual reason introduces its own set of project priorities, parameters, limitations, time-lines, risks, and costs. Any Data Center migration is a complex and potentially disruptive undertaking that involves thorough planning and meticulous execution of that plan.

Typical business drivers for data center migration include:

- **Data Center Facilities Upgrade.** Where relocating to a new facility may present a better cost benefit ratio than refurbishing an existing facility. The existing facility could simply be out of space or the existing infrastructure too old or complex to replace or refurbish. Energy efficiency gains via physical plant design and operations in many cases warrant a new facility versus retrofitting a legacy site.
- **Data Center Footprint Reduction.** Many organizations are now downsizing their operating model from one of "owner / operator" of their own facilities, to outsourced, or hosting models.
- **Business Continuity Improvements.** The existing facility is a stand-alone facility and so does not provide the level of DR / BC that diverse facilities will provide.
- **Mergers and Acquisitions.** One of the most common reasons where due to M&A, the new company has too many data centers and needs to rationalise systems into a lesser number of facilities.
- **End of Lease.** An operating lease on their current facility may be drawing to a close and a new contract for a new facility may present better economic as well as technical benefits.
- **End of Facility.** Redevelopment of physical premises often triggers the need for data center exits. There's nothing quite like a wrecking-ball as a motivator.
- **Cloud Migration.** Cloud migrations are increasingly relevant to data center migration projects. While all legacy systems will not be candidate for cloud, this type of initiative often triggers the need for data center consolidation and migration.

Any data center migration is an exercise in risk mitigation. The planning process itself requires significant management commitment and resources to identify and address the key issues and risks. Any migration brings with it an increased risk of downtime. Even if the migration is not involving a critical device or service, the risk of collateral damage is always present when infrastructure running undisturbed for lengthy periods of time is suddenly disturbed. It is essential for plans to be in place to deal with unplanned failures and data loss. Therefore, any migration must take into account the wider business operational requirements and plan accordingly. Lastly, cost of migration must be rationalized against the cost of risk mitigation - a delicate balance often requiring careful strategic and budgetary planning.

3 DC Migration Phases

As a general overview, the following phases constitute a data center migration initiative.

Project Initiation: Data Center Migrations typically start with an initial start-up phase, which provide migration strategy and high-level design/plan inputs necessary for project budgeting and go-forward approvals. Often this phase is skipped and budgetary estimates require massive rework in the middle of a data center migration project. When performed, this phase delivers:

- High-level migration strategy & design documents
- High-level project plans including a work breakdown structure
- Preliminary bills of materials
- Preliminary resource requirements for the entire project
- Preliminary budgetary estimates and signoff

Discovery: Detailed Discovery will involve a significant number of customer resources, internal data sources, and 3rd party tools to accurately identify the scale, performance, complexity and inter-dependencies of the current environment from both an infrastructure and an application / system point of view.

Design: In all migrations the target state design effort articulates the detailed systems and infrastructure components required for migrations to the target data center environment. Typically design involves key stakeholders in facilitated sessions to establish the appropriate designs and plans. Appropriate migration methodologies will be assigned to each element during this phase.

Planning: Based on the agreed-upon designs and migration methodologies, detailed migration project plans and timelines are developed. In smaller companies, where in-house resources may be lacking in numbers, skills, or migration experience, 3rd party services often assist in detailed migration planning efforts. Migration planning outputs include detailed migration plan activities, resource assignments

Execution: Lead the program and provide other key project personnel to deliver the migration(s) as per the agreed SOW. The project lead will act as a “Command and Control Center” to coordinate the migration activities, as well as to track and communicate progress, perform problem management and escalation coordination when problems occur. Successes and failures are documented and used as input into the “Lessons Learned” process after the migration to improve the process for subsequent migrations.

Post Execution: Co-ordinate the effort to test the new environments until customer sign-off or as agreed. We will also provide active support by sharing our extensive Data Center Management experience thus contributing to the client’s own knowledge base. In this regard, follow-on opportunities can be developed in the areas of DC Operational Management, Cost/Financial management, Energy management and conservation, Best-Practise implementation

Project Closure: Facilitate the delivery of contracted documentation and to manage remedial activities in the closing down of project related defects.

In totality, a data center migration is a complex organizational and technical exercise, as illustrated in the following flow diagram.

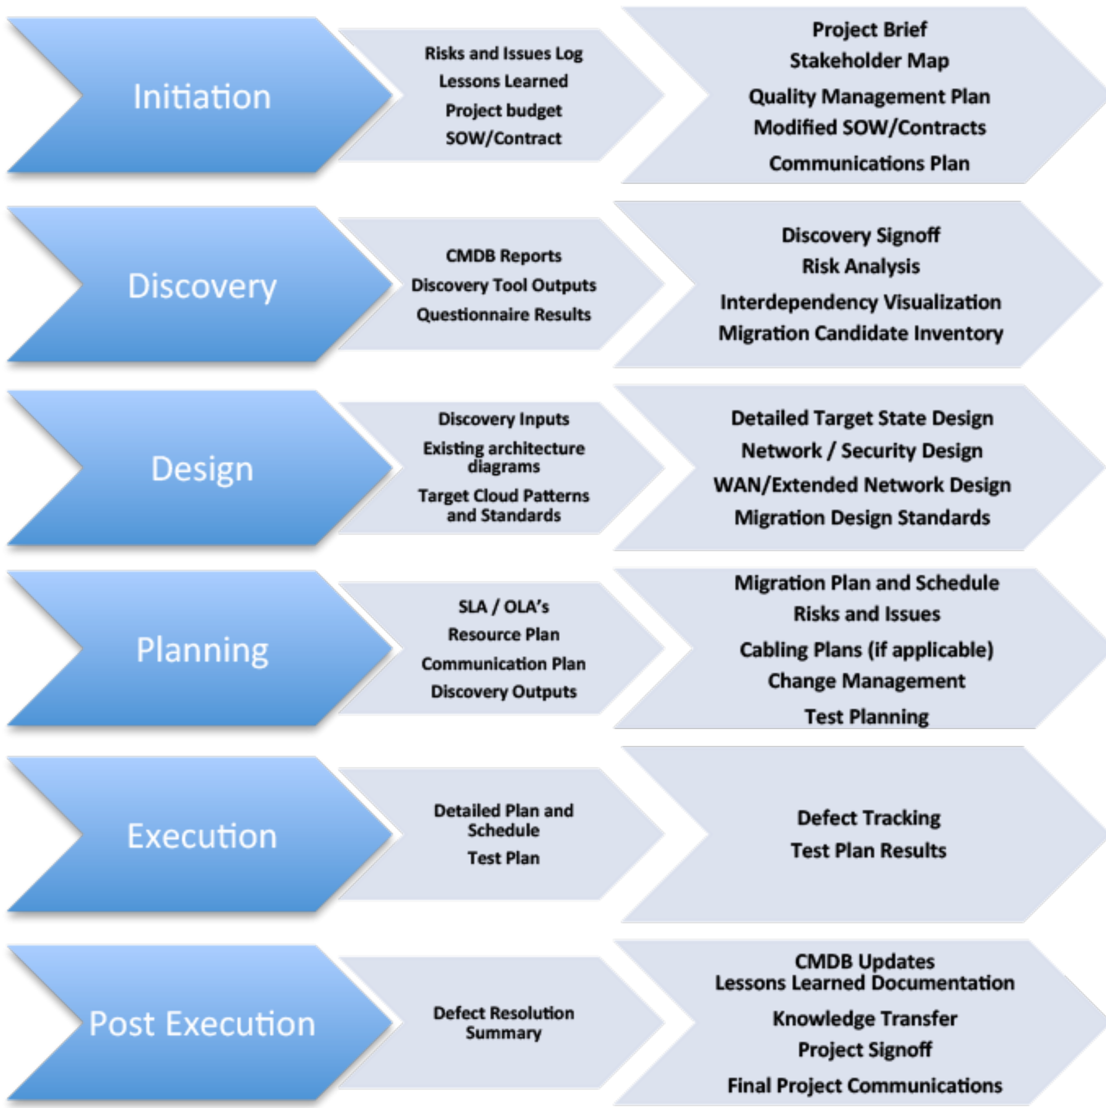

Figure 1 - Flow Diagram of a Data Center Migration Program

The following subsections detail they types of migrations and associated attributes. For purposes of this whitepaper we do not go into related detail for project initiation, design, planning, or post execution phases, instead focusing where migration automation has appropriate overlay with traditional migration methods.

4 Tools & Technologies

Most enterprise migration projects start with the leveraging of incumbent enterprise tools. It's not unlike the majority of home improvement projects where an existing toolkit is adapted to the job at hand, and regardless of the requirements, the toolkit is adapted to meet some version of the requirements. As is often the case, many discovery phases of enterprise migrations are ongoing with no clearly defined endpoint.

The reality is that the majority of mid-sized migrations function on the back of GANTT schedules, spread sheets, and 'grunt'. If you're in the market of moving less than 500 assets (virtual, physical, etc.), you very well may benefit from not using tools to deliver your DC migration project. The time, up skilling, investment, management, and patience to properly learn and implement these tools are prerequisites to obtaining value from their deployment, and the ROI may never be realised.

However, if you're facing an enterprise migration project, where multiple locations, many 100's of applications, and 100's to 1000's of assets, a well-planned and deployed toolkit is essential.

The following is the scope of tools required for an enterprise data center migration project.

4.1 Discovery

Discovery tools function to collect all data regarding the applications and infrastructure within the source environment. Typically, a CMDB is regarded as the 'partial' source of truth, and discovery tools are used to inventory the physical and virtual environment to help compile a complete view of the infrastructure topology. Traditional tools include client agent installs and perform a variety of client side query activity. TADDM (IBM), ADM (EMC), ADDM (BMC) are good examples of these tools designed for ongoing discovery and inventory of a customer source environment. Often, information overload can result from feature-rich tools in this space, so a common challenge includes lengthy integration times followed by massive data analysis and parsing efforts.

Discovery newcomers include Risc Networks, who deploy less invasive network based discovery methods. Both address host interdependency mapping, which is often required to design migration move groups in the migration planning process. We see a lot of merit in this improved approach.

Ultimately, the sources for the discovery data will include operational data, data from deployed tools, and manually input source files such as questionnaires and workshop outputs. The latter are often employed to extract institutional knowledge into a coherent view of 'what's running'. Organising and maintaining this information into an actionable state is another challenge, addressed below with migration management & governance.

4.2 Analysis

Increasingly for enterprise migrations, the business wants to rationalise the application portfolio before moving it into the cloud. For cloud migrations, the common question is 'will my application work in the cloud?'. The combined themes trend towards simplifying the legacy application environment in parallel to modernising the compute stack. Two levels of analysis are required for this activity: business application analysis & technical application analysis.

Business application analysis is traditionally performed with workshops and enterprise architects, a process both expensive and time consuming. An emerging set of tools works to address business application analysis, through online decision trees and workflows. The main benefit of employing these tools is to create a consistent set of outputs that can be used to drive the application modernisation strategy. Examples include Cloud Conveyor, Cloud Genera, & Gravitant. These tools provide a rational view of your application portfolio, at a business and high-level architecture level.

Technical application analysis targets the lower level challenge of mapping application binary characteristics from a source environment to determine if they will work in a target environment. This is a major task and often a significant problem in terms of the migration analysis and decision making for application modernisation. In the desktop virtualisation arena, AppDNA from Citrix has provided considerable value in mapping remediation requirements, providing estimations of effort for application packaging, and has started to show application of this technology to the AWS cloud migration use-case. Tools from various onsulting companies have also entered this space with bold and smart new products.

4.3 Migration Management / Governance

Once the discovery data collection and analysis phase of the migration is underway, keeping it organised and linked to the migration planning process is the next challenge. Again, for small to medium scale projects, it's likely that you would be better off with worksheets and standard project management tools, but on scale, the data volumes and linkage to migration planning becomes unwieldy. Taken to the extreme of what not to do, we've seen programs running with GANTT charts of 100,000 rows and dedicated teams administering the GANTT data.

Migration management and governance tools, which have been mostly built for consultants in migration boutique consultancies or SI's, provide a central repository for migration metadata, along with migration governance and planning functionality. This allows the discovery data to be linked to migration move groups, schedules, and coherently maintained throughout the process of migration planning and execution.

To obtain these tools, you usually have to work with the consulting group who built them, or with a clever boutique consultancy that has setup OEM agreements with the tool's creators. Bell Microsystems (UK), Transitional Data Systems (US), along with the major SI's who have invested in this space (Accenture, Dell, HP, etc.) are firms to look out for in the area of migration management and governance tooling.

4.4 Migration Automation

Today, a cumbersome medley of work-around scripts, manual copy methods, third party tools, and field engineering drive cloud migrations for enterprise application workloads. RiverMeadow's SaaS platform migrates any x86 workload from physical, virtual, or cloud sources, into virtualized, public, private, hybrid cloud environments. Although there are a number of different methods and companies out there, RiverMeadow provides live point-to-point migration automation (Secure Direct Migration), which helps increase migration speeds, decrease migration costs, and lower risks associated with manual migration methods.

5 Types of Cloud Workload Migrations

A data center migration typically includes a mix of workload (server operating system), storage, network, and miscellaneous hardware migrations and relocations. For purposes of this whitepaper, we focus in specifically on workload migrations. A workload may be physical or virtual, and is generally viewed as a single instance of an operating system. All workload migrations include an element of solution design, and one type of migration is applied to a single workload.

Types of migrations in a data center migration initiative may include:

- Virtual-to-Virtual Relocation (V2V)
- Virtual-to-Cloud Relocation (V2C)
- Physical to Virtual Relocation (P2V)
- Physical to Cloud Relocation (P2C)

In any migration, all migration strategies share a common set of principles:

- Migrations should minimise business disruption, by reducing to a minimum the time required for cutover of services between sites. Where ample time has been allocated by the client to complete the migration, use the extra time to minimise risk.
- In the event that a given migration is unsuccessful, the migration strategy should provide a straight forward remedy plan to achieve a stable state. This could be a rollback plan to a previously stable environment, or as sometimes happens, a corporate decision is made to not execute any rollback at all. In this case, a failed migration would see focussed effort on problem resolution with a view of achieving a new stable environment in the target DC,

- The migration strategy should, where possible, provide the ability to test completed migrations prior to service cutover to minimise the likelihood of issues for high-risk business applications

For Lift and Shift migrations, the minimum outage time for each application would be the time taken to power down, de-rack, prepare for transport, transport, prepare for installation, rack, cable and power-on whilst rollback takes the same amount of time to complete but carries the same or greater risk as the physical migration and cannot rapidly be executed. This is why migration automation and workload migration often provides a lower risk solution for data center transitions. The following sections outline key characteristics of various workload migrations seen in a data center migration.

5.1 Virtual-to-Virtual Migrations

V2V is the practice of moving one virtual machine image to a new location without any physical conversion. This entails the migration of an operating system (OS), application programs and data from one virtual machine or disk partition to another virtual machine or disk partition. The relocation of existing Virtual Machines is in many ways the simplest relocation scenario. Virtual Machines are inherently portable thanks to their very nature. The abstraction of the server hardware layer and encapsulation of the server workload into monolithic files resident on the SAN infrastructure means that no reconfiguration is required within the VM following its relocation to the new site, even if the underlying hardware platform is different. A number of traditional Virtual Machine migration options are set out below.

Option	Advantages	Disadvantages
Relocate Virtual Machine files via SAN data migration	Can occur continuously in the background, requiring only a brief synchronisation period prior to cutover. Replication throughput is quite high and provides a straight-forward fall-back scenario	Replication and cutover occurs at the volume level, requiring additional manual process to allow selective cutover of individual VMs.
Relocate Virtual Machine files via Storage Clones	Cutover occurs at the VM level whilst replication occurs via a link and is unaffected by storage replication volumes Minimal pre- and post- work required to support this technique	Clones are comparatively slow – with typical throughput at the VM level of around 2GB/minute.
Import VMs at target site	Eliminates dependency on inter-site links	Significant logistical overhead. Extensive outages required (no data must change between last backup and re-entry to production at the new site)

Table 1: Traditional V2V Migration Options

If a new SAN is to be part of the migration effort, then clearly SAN replication will be the preferred relocation method for Virtual Machine images. A temporary staging volume would be deployed on the target SAN, sized appropriately to accommodate the largest move group. This volume will be mirrored asynchronously following which it will be mounted to the virtualisation infrastructure at both sites. This way, a single synchronised version of the latest data is being used at both the old and the new sites until such time as the synchronization link is broken.

Most traditional V2V migrations are cold migrations, which require that the source system be powered off during the conversion process. It is, however, possible to perform a live migration of a virtual machine using RiverMeadow. An initial assessment needs to be done of the source and destination hardware specification relevant to the migration such as storage space, memory and CPU allocation. Following an assessment typically the procedure will be as follows:

- Begin VM in source Hypervisor and document configuration. These can include VM settings such as CPU, disk type, memory, disk settings. Also document virtual network settings such as VLAN along with network adapter settings including IP address, DNS, subnet, and gateway.
- Ensure OS is up to date with the required hotfixes and software for best practice.
- Uninstall any source hypervisor drivers.
- Shutdown the VM and backup any data.
- Clean up any differencing disks that may cause alignment issues.
- Depending on version of hypervisor, use the appropriate tool to connect to source hypervisor to backup and migrate to destination hypervisor. Typically, for VMware, there VMware Converter. Xen Converter exists for XenServer and Hyper-V. Migration tool may migrate and also install any destination hypervisor drivers.
- If the tool only converts to a file such as .vhd, then copy the newly created converted VM to destination such as the VMM library in Hyper-V.
- Start your VM and ensure hypervisor drivers have installed correctly, network settings have applied and VM specifications are as intended.

Conversion Tool	Support
VMware vConverter	VMware Workstation, VMware Player, VMware Server, VMware Fusion, Parallels Desktop, Microsoft Virtual PC and Microsoft Virtual Server
PlateSpin Migrate	VMware, Hyper-V, XenServer
Quest vConverter	VMware, Hyper-V, XenServer

Table 2: Traditional V2V Tools

5.1.1 Benefits of a V2V Migration

- Can be migrated to new virtualization platforms using conversion tools.
- Source computer is available throughout the conversion.
- No data change during /post conversion in the source server.
- Target virtual server is an exact clone of the source server.
- Application configuration and data are the same.
- Target virtual server can be customized based on performance trends of the source server.
- Drivers for the virtual server are provided by the virtualization platform.
- Source Server can be reverted back in case of application issues in target server.
- Conversion tool can (usually) be automated for post conversion tasks.
- Source server can be converted to an image and tested before the actual conversion.
- Where no connectivity exists between data centers physical servers can be converted into images and then transferred to the new data center using file transfer mechanism for hosting it as virtual machines.
- Virtual machine hardware can be upgraded to the latest hypervisor release pre or post migration.
- Lesser effort when compared to P2V
- In the event that the migration UAT fails, simply powering on the original Virtual Machine and shutting down the original can achieve rollback.

5.2 PHYSICAL TO VIRTUAL MIGRATION

Conversion of existing physical machines to Virtual Machines presents a number of clear benefits.

- Virtual Workloads, as already explained, are inherently portable, simplifying their relocation

- Virtual to virtual relocation is quicker, less complex and less risky than physical to physical relocation
- Physical assets freed up by the virtualisation process can be redeployed as seed/swing assets to aid the relocation of remaining physical assets or returned to the pool to be re-tasked as development or test kit

However, the conversion of existing physical machines to Virtual Machines introduces a greater degree of application risk due to changes at the OS level such as driver changes, new MAC addresses, new CPUID's etc. P2Vmigration should therefore be primarily conducted against low-risk, low-complexity devices.

Typically, a tool will be used to convert from physical to virtual. These P2V products connect to the physical server, copy all data from that physical server into a virtual disk on the virtual server, replace the drivers in the guest operating system with virtual drivers, and start the new virtual machine. In some cases, there is no downtime for end users of that server.

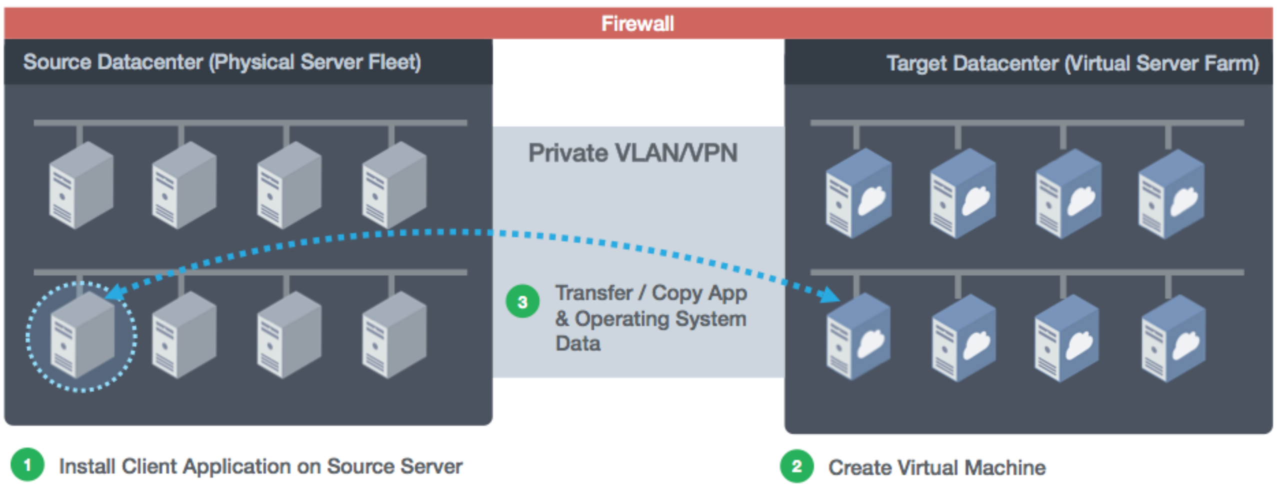

Figure 2 - Traditional Physical to Virtual Conversion Process

- The first phase of a P2V conversion is to survey the hardware configuration of the source computer and make sure the patch cache contains all necessary drivers and system files to support the configuration. If any drivers are missing, you will get specific error messages indicating where to get the necessary drivers.
- Document security credentials for the source machine and target hypervisor host. Identify any services on the source server that needs to be stopped during conversion or disabled on the target server. Resources required on the target, such as number of CPUS, amount of RAM, number of disks and NICs, plus a temporary IP address.
- Consideration to be given to when the migration should run – it can be scheduled to run out of hours if required.
- Ensure there is adequate storage space and VM specifications such as CPU, network, memory and disk type that you would like to assign available.
- Bad sectors on disk cannot be transferred during a P2V conversion. To avoid data loss, run a disk maintenance tool such as Chkdsk on the source computer to detect and correct any file system errors.
- To help minimize the time required for the imaging phase, perform a disk defragmentation on the source computer's hard drives. Also, ensure that you have a fast network connection between the source computer and the host.
- Using vConverter you can convert while the source is turned on. Using this method is simple and easy however it still requires that you run the clean-up tasks and ensure network connectivity on the destination following completion.

- Converting from offline requires other tools such as VMware Cold-Clone which requires you boot into the CD on the source physical machine and run through the conversion process.
- If using Hyper-V, conversion can be done using SCVMM. This process is similar to vConverter and requires the physical machine to be turned on. It is a single step process where it connect to source machine and converts to a virtual machine.
- There is a degree of post-conversion tidying up to do, uninstalling applications no longer required (such as hardware agents), working through Plug & Play detection to install additional drivers and installing VMTools (though many apps do this for you).
- As the original physical server is unmodified during the P2V conversion process, simply disabling the new Virtual Machine and powering up the original physical server can achieve rollback.

Physical to Virtual conversion can be undertaken in one of two ways as outlined below.

Option	Advantages	Disadvantages
Option A: Servers can first be virtualised within the source environment and then relocated using the V2V strategy outlined previously	Allows virtualisation to begin prior to the destination virtualisation infrastructure becoming available. Allows the impact to server performance arising from the virtualisation to be assessed independently of any impacts arising from the migration	Requires double handling of servers as they must later undergo a separate outage in order to relocate them to the new facility.
Option B: Servers can be virtualised directly to the destination Data Center virtualisation infrastructure, eliminating the need for them to then be separately migrated between facilities.	Reduces the number of outages sustained to each application Reduces the number of servers that need to be migrated between sites.	Compounded changes makes it difficult to separate the impact of relocation from the impact arising from virtualisation

Table 3 - P2V Conversion Approaches

There are two key pieces of capacity planning required to support the P2V migration scenario:

- Virtual Machines must be optimally sized to deliver equivalent or better performance to their previous physical server equivalents
- The virtualisation infrastructure must be sized to provide sufficient capacity for the additional virtualised server workloads, whilst still permitting growth.

> **Note:** There is a school of thought, which suggests that including a P2V change into the migration is a bad idea in itself. Compounding changes is always a higher risk. If possible, the P2V should be done as a separate sub-project at the source Data Center prior to the migration and then during the migration window, treat this device as a V2V migration. (As per Option A above)

5.2.1 Benefits of a P2V Migration

- Source computer is available throughout the conversion.
- No data change during /post conversion in the source server.
- Target virtual server is an exact clone of the source server.
- Application configuration and data are the same.
- Target virtual server can be customized based on performance trends of the source server.
- Drivers for the virtual server are provided by the virtualization platform.
- Source Server can be reverted back in case of application issues in target server.
- Conversion tool can be automated for post conversion tasks.
- Can convert source server to an image and tested before the actual conversion.
- In scenarios where no connectivity exists between data centers, physical servers can be converted into images and then transferred to the new data center using file transfer mechanism for hosting it as virtual machines.
- Less effort required
- All the intrinsic benefits of Virtualization.

5.2.2 Risks of a P2V Migration

- The primary risk is the introduction of a changed environment during a move. Compounding changes is never a good idea as it makes it difficult to assess the overall success or failure of the move.
- If the application is not compatible for virtualization it might lead to server issues.
- Clean-up of drivers and non-plug and play devices will require manual effort.
- Servers with hardware limitations cannot be converted.
- Performance issues can arise if the virtual server is not tuned to meet application requirements.
- Hostname changes in the target can lead to application issues.

6 Migration Automation Benefits

6.1 RiverMeadow SaaS Architecture Overview

Migration automation software potentially lowers risk and complexity for a data center migration, and offers promising capability to automate a large number of standard workload migration procedures. The following diagram outlines the general reference architecture for RiverMeadow software, a migration automation tool built for private and public cloud migrations.

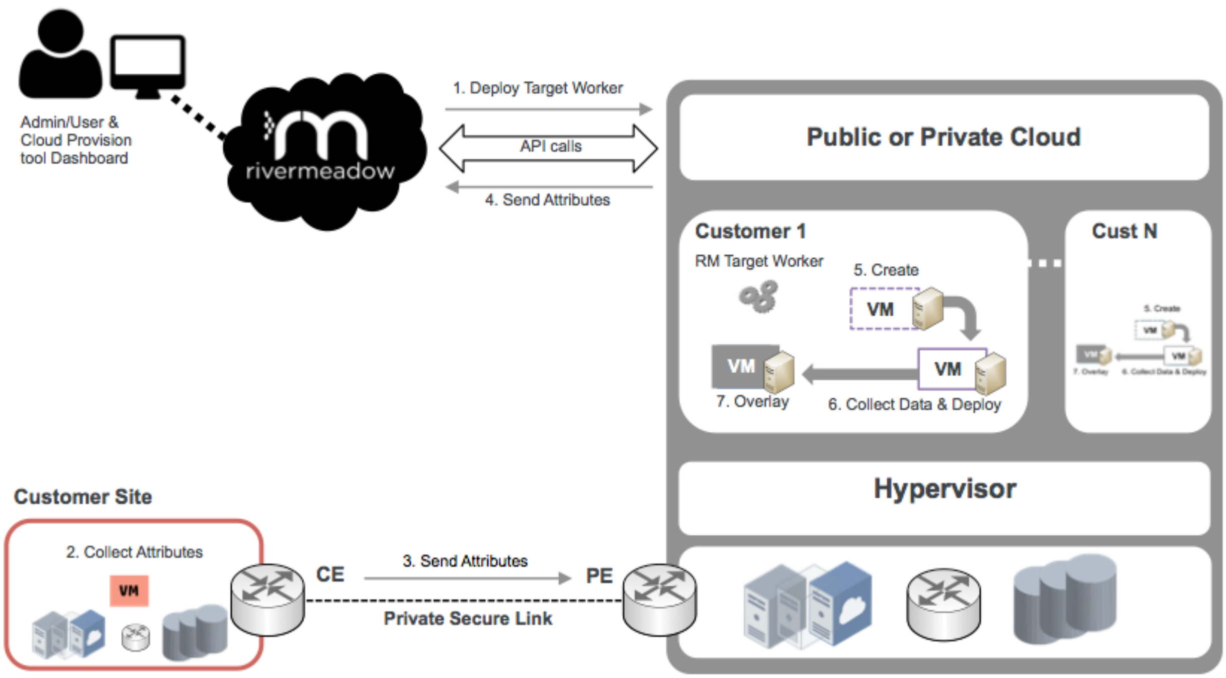

Figure 3 - RiverMeadow Reference Architecture

The following sequence outlines how Secure Direct Migration happens with RiverMeadow:

1. CloudAppliance Deploys to Target Cloud
2. CloudAppliance gets instructions from the Message Bus
3. CloudAppliance collects attributes from source
4. CloudAppliance pushes attributes up to RiverMeadow
5. RiverMeadow SaaS does modelling and creates a TargetWorker
6. TargetWorker is pushed into target cloud
7. TargetWorker collects data
8. TargetWorker customizes for target cloud and reboots
9. Live workload migration is complete

No software is installed in the customer premise on source or target workloads, and the data plane accommodates a private networking traffic and perimeter security boundaries as long as network traffic between the source and target machines is permitted and routable. The following chart outlines the network port requirements for the control plane communicate between RiverMeadow SaaS, and the target cloud environment.

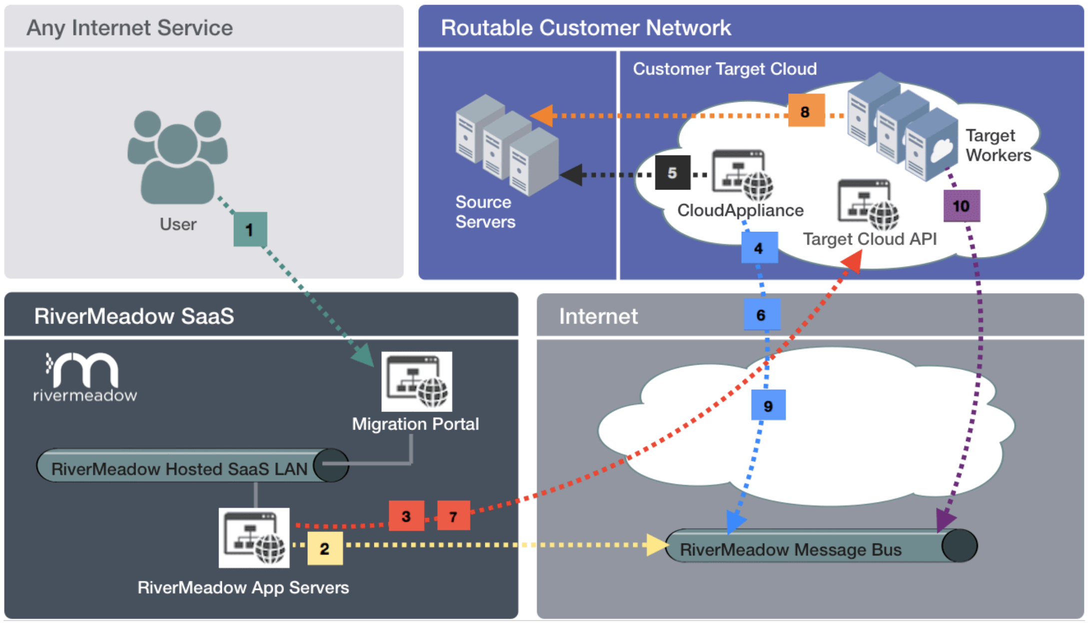

Session	Source Host	Destination Host	Protocol: port(s)	Comments
1	End User Host	RiverMeadow Migration Portal	TCP:80/443	Customer UI
2	RiverMeadow App Servers	RiverMeadow Message Bus	TCP:443	Ongoing status connection
3	Application	Target Cloud API Host	TCP:443	CloudAppliance instantiation
4	CloudAppliance	RiverMeadow Message Bus	TCP:443	CloudAppliance check-in
5	CloudAppliance	Source Server	TCP:22, 139, 445	Attribute collection
6	CloudAppliance	RiverMeadow Message Bus	TCP:443	Attribute notification
7	Application	Target Cloud API Host	TCP:443	TargetWorker instantiation via SSL API
8	TargetWorker	Source Server	TCP:22, 139, 445, 50000	Data migration
9	CloudAppliance	RiverMeadow Message Bus	TCP:443	Success/Failure notification
10	TargetWorker	RiverMeadow Message Bus	TCP:443	Log data

Figure 4 - RiverMeadow SaaS Network Requirements

Given the right mix of migration types, customers can leverage automation tools and frameworks to standardize and automate workload migration processes. As such, the following table outlines the types of data center migrations where migration automation software has particular benefits.

Migration Type	Automation Benefits	Comments
Virtual-to-Virtual Relocation (V2V)	Workload migration automation Integration with private cloud environments	Hypervisor director access is required
Virtual-to-Cloud Relocation (V2C)	Workload migration automation Integration with public cloud environments	Public cloud migration automation typically requires integration with cloud provider API's along with initial design/planning for the target environment
Physical to Virtual Relocation (P2V)	Virtualization automation Reduced # of tools and associated methods	Live migration and subsequent offline target environment testing/validation lowers P2V risk
Physical to Physical Relocation (P2P)	N/A	
Virtual to Physical (V2P)	N/A	
Physical Lift and Shift (L&S)	N/A	
Physical to Cloud (P2C)	Provides API drive migration path for physical workloads to external cloud service provider	Network and security design is a major prerequisite to this form of transition

Table 4 - Migration Methods and Automation Benefits

6.2 RISK MITIGATION

Collectively, increasing the diversity of migration scenarios increases data center migration complexity & risk. With added tools and techniques, the typical customer environment bandwidth and skill-base is often stretched as a result. Normal pressures of off-hours change windows and weekend/night migration windows, also add to the risk and introduce increasing factors for human errors and omissions during the overall migration process. RiverMeadow helps to mitigate this risk through:

- GUI driven migrations - a simple interface for human driven operations
- Live migration & offline testing - no disruption to the host/source environment, and the ability to offline test the target machine prior to cutover
- API driven migrations and bulk automation - ability to automate migration operations on scale via the RiverMeadow API interface and general automation code (less hands, fewer mistakes)

6.3 CONSOLIDATION OF TOOLS & METHODS

The typical data center migration project involves a myriad of internal and 3rd party tools. Often, base operating system utilities, scripts, migration software, and storage device specific tools are applied to platform specific migration requirements.

> **For Instance:** We typically see a mix of Robocopy and storage array copy commands for Windows servers, then Platespin or VMware tools for Windows virtualization requirements, and VMware tools for hypervisor upgrade requirements, all for a homogenous Windows compute environment. As Linux, Unix, and other disparate platform environments are covered customers typically add more and more tools, hacks, and methods to accomplish the migration exercise. As a result, many customers manage 6-12 disparate toolsets and utilities to accomplish workload and data migrations. Even for advanced system administrators, this presents challenge and risk when these tools are not used on a daily basis.

The following table provides an overview of platform specific tools and where migration automation software provides a consolidated toolset option.

Migration Tool/Utility Type	Migration Automation Functionality
Windows Operating System Copy (OS backup)	Yes
Windows data copy / synchronization (Robocopy)	Yes
Linux Operating System Copy (export commands)	Yes
Linux data copy / synchronization (rsync)	Yes
VMware hypervisor upgrade (vConverter)	Yes
Virtualization Automation (P2V Platespin, etc.)	Yes
Unix Operating System Copy (OS export/copy)	No
Unix data copy	No
Storage array copy	No
Network service migration	No

Table 5 - Migration Tools and RiverMeadow SaaS Capability

6.4 Conclusion - Lowered Risk & Increased Efficiency

A number of migration methodologies are supported by RiverMeadow SaaS. Diligence in discovery, planning, and testing is still required, however this technology does afford possibilities to automate bulk migrations and move-group actions. In addition, with API interfaces and live-migration capability, less disruptive migration dress rehearsals are also capable of lower risk for actual migration cutover events.

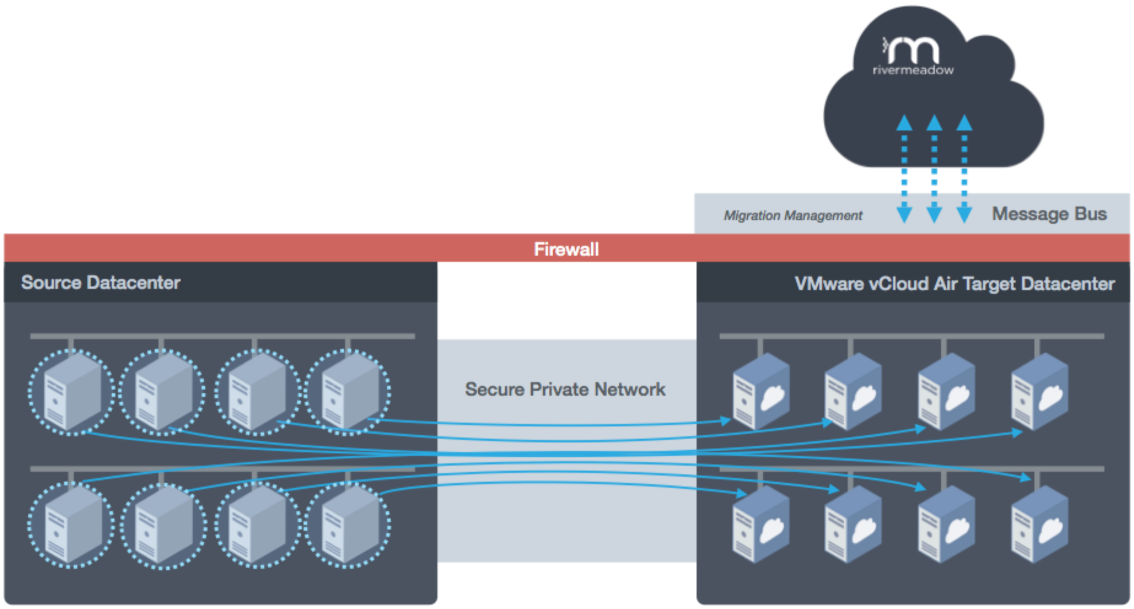

Figure 5 - Migration Automation on Scale

In conclusion, we present a proven methodology and toolset to help address many of the complexities and issue mentioned above. The value proposition for migration automation includes lowered risk, lowered complexity, and increased efficiency for data center migrations. As a result, adopting migration automation within a data centre migration program will provide:

- Fewer hands, less source side disruption, and fewer mistakes during migration move-group events
- Reduction in manual copy methods and adhoc 'work-arounds' - often usually increase risk
- Leverage of a common migration software platform with fewer associated migration methods
- Ability to automate workload migrations with increased efficiency and reduced risk profile per migration event and for the migration program as a whole
- Adoption of devops 'infrastructure as code' for traditional migrations using automation tools and API programming