

How to Migrate Without Downtime

An Executive Overview



Accuracy is the minimal criterion for measuring the success of the migration, but that alone is not adequate.

Businesses grow. Technologies evolve. System migrations are inevitable, but business downtime isn't.

Migrating applications and data to new platforms involves risk and takes time. Possibly worse, when using traditional migration methods and technologies, operations have to be halted during the migration. And, the larger the storage footprint, the longer the process will take. For many large systems, migrations routinely span hours or days. Consequently, application, system and data migrations have always been a challenge for IT staff; in fact, most have simply gotten used to the pain and suffering they engender.

The IT department usually cobbles together a range of tools to perform migrations, but, even with those internally developed tools, successful migration is far from quick or assured. Beyond moving data and applications, manually configuring the new platform to match the original system as close as possible, along with the testing, testing and more testing required to ensure that the applications are functioning correctly on the new platform, adds to the time, expense and hassle of migrations.

Accuracy is the minimal criterion for measuring the success of the migration, but that alone is not adequate. Businesses utterly depend on their information systems. What's more, because of the Internet, globalization and competitive pressures, many organizations now operate around the clock. With the old "maintenance windows" now reduced or eliminated, there is no longer any time for system downtime. Therein lies the challenge. How can you migrate a production system without incurring business downtime?

This fundamental quandary—the seemingly impossible goal of successfully migrating systems without business downtime—and the solutions to that quandary are the subject of this executive overview.

Common Migration Triggers

The forces necessitating server or storage migrations are many, including the following:

Cloud – The Cloud has become a popular way to host business applications because it can save companies considerable amounts of money. While the Cloud has a lot to offer, the migration process is not without risk.

Virtualization – Physical to virtual migration (P2V) garners considerable media attention these days, but application virtualization isn't a sure thing. Consider the following questions: What happens when your production applications don't perform well—or don't perform at all—after you've virtualized them? How can you keep your applications running when you change virtualization vendors? What if you don't want to take your production applications offline while the entire machine is virtualized?

Hardware Refresh – Hardware usually has to be refreshed when it has exceeded its warranty, lease period or useful lifetime. Regardless of the reason, you need to move your data to the new hardware as quickly as possible and reduce the amount of interruption to the business processes that depend on the affected systems.

Hardware Upgrade – When an application outgrows the resources allocated to it, you have few choices but to give that application more computing and/or storage resources. If the system that the application is running on is already at full capacity and the application has exclusive use of the system, then your only choice is to migrate to bigger hardware.

Operating System Upgrade – Operating system vendors apply considerable research and development effort to enhance their products by adding new features, improving security, boosting performance and accommodating hardware innovations. In addition, new applications and application versions, which may also offer significant advantages over legacy applications and versions, may not be compatible with earlier operating system versions. Thus, remaining competitive may require migrating to the latest operating system version to take advantage of these benefits.

Changing Vendors – The process of migrating data assets between two arrays from the same manufacturer might be time-consuming, but fairly straightforward nonetheless. In contrast, disparate platforms seldom work together seamlessly. Thus migrating to a different model or manufacturer of hardware may entail risky and manual migration processes that require considerable application downtime.

Whatever the reasons for their migrations, organizations need migration processes and technologies that eliminate user impact and reduce risk. Read on to find out how you can greatly reduce the risks and downtime of migrations, and take back your holidays and weekends at the same time.

Migration Requirements

Migration has different impacts depending on what you're migrating (data only, whole applications, configurations and system state or multi-tiered applications). On the surface, moving data between systems would seem to be the easiest type of migration because it appears that you would just have to copy data to a new environment. However, as file servers and databases have grown over the years, the time required to copy all of those bytes has grown as well. It can easily take hours to copy one terabyte of data, even when using the fastest networks and the fastest storage.

This raises some critical questions for IT departments:

- How are these requirements affected if you need to migrate an entire server, rather than only the related data?
- How can you migrate servers that end-users need to access 24x7?
- How do disparate platforms and geographic distance requirements affect your migration plans?

To answer these questions, you first have to consider the nature of the migration that you are going to undertake, as discussed below.

Data Migration

When migrating data from a file server using traditional data copying tool's you must ensure that users can't change any data while the migration takes place. Otherwise, changes to files that have already been copied will be lost. Locking out users can result in a significant and intolerable amount of application downtime.

You may also need to create a plan for recreating the file shares, permissions, compression, encryption and other settings on the new storage. This can translate into further hours of downtime even if everything works as planned the first time—and much longer if it doesn't. Unfortunately, data migrations using traditional tools often do fail the first time.

Unforeseen obstacles, such as an application or anti-virus processes that lock a file, can cause the copy or synchronization process to fail. When this happens, you usually must restart the copy process from the beginning after finding and stopping the process that's causing the problem.

IT departments need a way to overcome these challenges. They need a way to perform data migrations while production applications continue to function. They also need a way to automatically avoid or correct migration errors.

Application and Full Server Migration

Applications usually consist of more than just a few executable files on disk. Performing traditional migrations of whole applications requires specialized in-depth knowledge of their inner workings. For example, in Windows environments applications can store vital information in configuration files or centralized locations like the Windows registry. Applications can also store vital objects like DLLs outside of their installation directory in shared locations like the Windows System32 directories.

Migration failure is a chief concern for any IT organization and it can happen for any number of reasons. Regardless of the cause, the impact is extensive and it results in additional downtime while the IT staff tries to track down the problem and correct it.

In addition, some Windows applications can embed themselves into the operating system for licensing restrictions. Also, the people that built a custom application system may not be available to determine its requirements. Yet, you must fully understand the application and its requirements in order to create a solid plan to perform a migration using traditional methods. Finally, application migration typically involves lengthy acceptance testing to ensure that the application was migrated completely before it is approved for production use.

The technical details differ, but other environments, such as IBM i and AIX, often place similar obstacles in the way of successful application migrations. There are no hard rules that application developers are required to follow when developing software. This means they have an open canvas and great flexibility, but it also creates a nightmare for administrators who have to migrate between applications.

Operating System Version Migrations

Migrating from one version of an operating system to a newer version is usually not a problem if you can afford the required downtime. However, that is often not the case. In a 24/7 environment, such as a high-volume Web store, downtime is not an option. In these scenarios, migrations must occur while production operations continue to function unimpaired.

Migration between Disparate Platforms

Migration is tough enough when your production and target systems are fully compatible (such as migrating data between two Windows servers). Now imagine the challenges of migrating between systems that aren't fully compatible. Maybe you want to migrate from direct attached storage (DAS) to a storage area network (SAN), or from an older storage product line to a newer product line or change vendors entirely. Maybe you want to change storage technology from fiber channel to iSCSI or change storage geometry from SATA RAID 5 to SCSI RAID 10. Or maybe you are replacing an old application with a new one from another vendor and the two applications run on different platforms.

In addition, as problematic as it can be to migrate between disparate physical platforms, migrating between different servers and virtualization platforms is even more difficult because vendor specific settings can be deeply embedded into the platforms.

If all you have available are traditional data copy tools, it's unlikely that you'll be able to migrate your data and applications between incompatible platforms without significant effort and lengthy production impact.

Unlimited Distance Migration

Migrating to new systems routinely takes place within the same data center. However, there is an additional set of challenges when performing a migration over long distances, which is often the case when moving a data center or consolidating multiple data centers.

Traditionally, there were two generic choices when performing geographic migrations: either suffer extensive application downtime while transmitting the bytes over low bandwidth, high latency WAN circuits; or shut everything down, load up the trucks and start driving.

Solutions that are designed to provide high availability (HA) or that are based on HA principles can help you to eliminate downtime during migrations, while also helping you to avoid many of the sources of migration failures.

Effects of Migration Failure

Migration failure is a chief concern for any IT organization and it can happen for any number of reasons. Regardless of the cause, the impact is extensive and it results in additional downtime while the IT staff tries to track down the problem and correct it.

Beyond outright failure, there might be performance issues that must be dealt with after a migration. Some migration technologies are one-way (no return). This becomes a major problem when you find that, after extensive trial runs, your applications simply don't perform well enough (or at all) under production load. In that case, you might want to split the workload between the old and new servers or revert entirely to the old server until the problem is resolved. Without two-way synchronization capabilities, this could be a problem as there would be no way to copy the changes made on the new server back to the old one.

You can also experience major migration setbacks when the tools that you're using fail to work as expected because of locked files, network disruptions or site outages. It's important that your migration tools are able to perform under adverse conditions, without production impact and without sacrificing data integrity. Otherwise, you'll be forced to start the migration over from the beginning when the initial attempt—or attempts—fails.

Using Technology to Overcome Migration Challenges

Solutions that are designed to provide high availability (HA) or that are based on HA principles can help you to eliminate downtime during migrations, while also helping you to avoid many of the sources of migration failures. How? Consider how HA solutions work. They monitor your production system and replicate data and object changes to a second system that serves as a hot-backup server. The software also typically provides an easy and fast way to switch users to the backup system when your production server becomes unavailable or you need to take it offline to perform maintenance.

HA and HA-like solutions can maintain system availability during most types of migrations because the replication processes that underlie those solutions usually don't require that the production and backup servers run the same version of an operating system or that the servers be of the same size. This means that you can place an upgraded server into the HA topology without needing to upgrade all servers in the topology simultaneously.

Because the replication software keeps the old and new servers continuously synchronized, without impacting production operations, there is no need to shut down operations during the migration process. As a result, migrations can be performed during normal business hours, meaning that it is no longer necessary to force IT staff to give up their nights, weekends and holidays to perform those tasks.

Migration Topologies

Many options are available these days to support your datacenter requirements. Migrating to newer technology can save your organization money and make you more competitive.

Physical-to-Virtual (P2V) and Virtual-to-Virtual (V2V) Migration

Replication-based migration tools provide migration services from physical or virtual production systems to virtual systems. You can perform migrations to virtual machines, such as VMware® ESX to Microsoft® Hyper-V™ and vice versa, that are normally impossible when using free virtualization tools. Replication-based migration tools move data, applications and system state information, including the operating system, registry, event logs and SID of the production machine. This provides a full server migration that doesn't require a lot of manual reconfiguration or acceptance testing.

You can choose the configuration for the new virtual server including CPUs, memory, storage, and network settings. When you perform your migration, you can replace the hardware dependencies to match the new virtual platform within the image, create the virtual machine and start the new virtual environment.

Virtual-to-Physical (V2P) Migration

Migrating from virtual to physical systems has always placed a strain on IT departments, especially when production applications don't perform as expected. Migration tools make this task as easy as P2V migrations. Simply select the virtual machine that you're interested in migrating and the new physical target machine, then have the tool move the data, applications and system state. Once the migration is finished, you can transition to the server by gracefully stopping the production virtual machine and completing the migration.

Physical-to-Physical (P2P) Migration

Whether you're performing a hardware refresh because your systems are coming off lease, out of warranty or they've simply outgrown their existing hardware, replication-based migration tools simplify P2P migrations. You can easily transition between different hardware vendors and complete migrations with only seconds of manual effort. Migration tools eliminate hardware dependencies when migrating servers and performs complete system state migration in record time.

Migrate to the Cloud

Migrating to the cloud may be a quick alternative when you find you need to expand your computing capacity quickly. Whether you are migrating a group of test servers to quickly expand computing capacity, using a cloud environment to quickly spin up some new applications without impacting your own data center, or just establishing a place to store virtual machine configurations for later use, replication-based migration tools make it simple and easy to migrate to the cloud.

Migration Techniques

There are three general techniques for using HA in a migration scenario:

- Switch Method
- Cascade Method
- Parallel Method

Switch Method

Most HA products allow you to halt the transmission portion of the replication processes so that any changes to data will still be captured on the primary system, but they will not be sent to the backup. When transmission is restarted, the HA software resynchronizes the two systems by sending all of the captured changes to the second system. The switch method, as illustrated in Figure 1, takes advantage of this feature to facilitate system migrations with minimal downtime. The process is as follows:

1. Stop the replication transmission processes, but allow users to remain on the primary system and continue to do their work as usual. The HA data stores will continue to collect all of the transactions.
 - On Windows systems these transactions are collected in a replication queue.
 - On IBM i, transactions are collected in journal receivers.
 - On AIX, log files are used as a repository

Replication

2. Upgrade the secondary system (hardware, operating system, or both) by saving and restoring the data from the old secondary system to the new one. This can be done anytime, as it does not affect users.

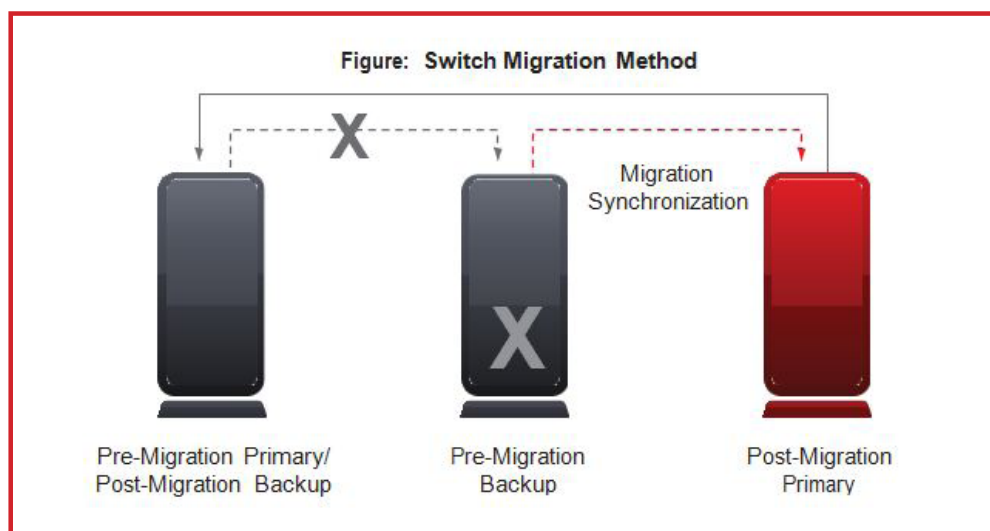


Figure 1: Switch method migration synchronization process

3. Turn the replication transmission processes back on, allowing the HA software to replicate all of the transactions that have transpired on the primary system since the connection to the old secondary system was severed.
4. When the two systems are resynchronized, perform a controlled switch of the users to the new system. Easy switching functionality is standard in most advanced HA software.

If both the primary and backup servers are being replaced, you can then repeat the above steps to swap out the old primary system and replace it with a new one. It is possible to omit the second switch if the primary and backup systems are both able to act equally well as the production platform. This is possible when the servers have similar performance and storage capacity and are in the same location or in equally advantageous locations.

Advantages

Simplicity: This is the most straightforward of the migration methods. If you are already using an HA solution, the existing software performs most of the work.

Cost: This is also typically the least expensive of the migration methods. You don't require any additional software licenses and, if the old primary or backup system will be adequate to serve as the backup system after the migration is complete, only one new server—the one that will be used as the primary server after the upgrade—needs to be purchased.

Challenges

Vulnerability: While the server on one side of the HA topology is being upgraded, the other server is the only one that is available to run production processing. This leaves your data and applications vulnerable if a disaster strikes during the migration process. It should be noted that, because the data and objects on the old secondary system are complete and up-to-date as of the point when you halted the replication processes, there are really only two threats in this case:

1. In the event of a disaster or other unplanned downtime on the active system, business applications will not be available until one of the two systems can be brought back online.
2. Transactions entered after the start of the migration process may be lost if the disaster destroys the active system's disks, including the HA stores, because replication is not occurring during the upgrade of the other server.

The probability of unplanned downtime events is small, but the threat is still too great for some businesses. If yours is one of them, you will have to use one of the other two migration methods, which will be described later.

Performance Constraints: To reduce costs, some companies use a backup system that has a significantly lower performance rating than their primary system. Their thinking is that planned downtime can be scheduled for “off-hours,” when many administrative systems are not being used; and, during unplanned downtime events, which are typically rare, the company can accept the curtailment of some non-critical systems until the primary system can be brought back online. In this case, the upgrade of the primary system will have to be performed during those “off-hours” to avoid seriously impeding business operations.

Disk Usage: While replication transmission is turned off, data updates build up in the active system’s HA data stores. If the upgrade takes considerable time and is performed while the business is operating at peak capacity, those HA data stores may consume considerable space. Before beginning the upgrade process, you must ensure that sufficient disk space is available to accommodate these requirements.

Synchronicity: Depending on the operating system and the replication product used, you may have the option of using either asynchronous or synchronous replication. In asynchronous mode, the writing of an update to the local system and the completion of the user transaction are independent of replication to the remote system. Only asynchronous replication is compatible with the switch method, since the backup server will be down for a period of time during the migration.

Synchronous replication writes updates to the remote system before writing them to the local system. During normal processing, the user’s transaction is not considered complete until replication is complete. This synchronous process may slow down transactions, some companies still insist on it because otherwise there is a very small chance that some transactions may be lost during a disaster.

When using the switch method, the remote server is unavailable. Therefore, synchronous replication will not allow the completion of any transactions. Thus, if the use of synchronous replication is mandatory in your organization and you can’t afford any downtime, you will have to use one of the other migration methods.

Cascade Method

The cascade method ensures that there is always a hot-standby backup server ready to take over operation should the primary system become unavailable for any reason.

The cascade method is appropriate in environments where two new servers (for production and backup) are replacing one or more old servers. Figure 2 illustrates the topology required for the cascade migration method.

Using this method, after bringing in the two new systems and configuring them, you replicate data and objects from the old backup system (labeled System “B” in 2) to the new primary system (System “C”) and from there to the new backup (System “D”).

After the systems are full replicas of all of the others you can keep this cascading replication running while you test the new systems. Any updates made on System “A” will be replicated in real-time, or near real-time, down through the entire chain. Consequently, the new primary and secondary systems will always contain an up-to-date replica.

When you are ready to go live with the new servers, simply switch users to the new primary server. You can then decommission the old primary and secondary systems at your leisure.

Advantages

Risk Reduction. The cascade method ensures that there is always a hot-standby backup server ready to take over operation should the primary system become unavailable for any reason.

The cascade method is advantageous even for companies that use a smaller backup server (System “B”). Unlike with the switch method, where the backup system must assume the production role while the primary system is being upgraded, when using the cascade method, the only time the backup system would need to take over operations is if a disaster strikes the production system during the migration process. The migration may be delayed, but the business can continue. Despite processing capacity possibly being reduced until the primary server can be brought back on line, this is still a better outcome than if a disaster strikes the active system during a switch-method migration.

Fewer Switches. Another benefit of this approach is that only a single switch—from the old to the new primary system—is required. Using the switch method, two switches are usually necessary when upgrading both the primary and backup servers—from the primary to the backup server and then back to the primary.




Figure 2: Cascade Migration Method

Challenge

Configuration Vulnerability. The cascade method assumes that all the necessary data is being replicated correctly between the original primary and secondary systems. If there were any mistakes made when configuring the replication from the old primary system to the old backup system (System “A” to “B”), the data being moved to the new servers will be incomplete or incorrect as well. The new servers will only be as complete as the old backup server. This should not be a concern for a company that regularly audits its HA environment, but it does mean that extra care should be taken to ensure that the original replication setup is correct and complete before starting the migration.

Parallel Method

As illustrated in Figure 3, the parallel migration method shares two benefits with the cascade method. First, there is always a hot-backup system available throughout the migration process. And, second, only a single switch—from the old to the new primary system—is required to upgrade both the primary and backup systems.

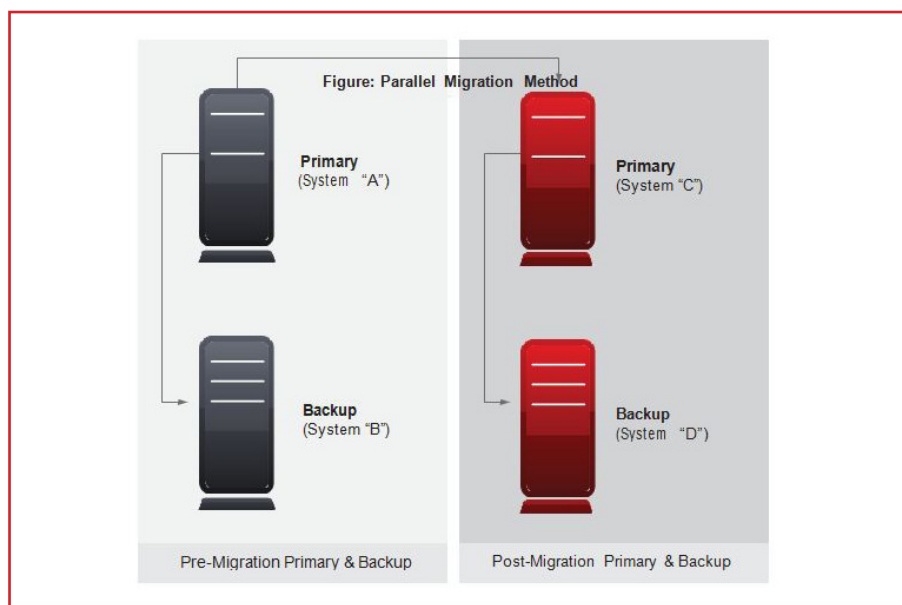


Figure 3: The parallel migration method ensures the availability to a hot-backup system.

The difference between the cascade and parallel methods is in the replication topology during the migration. Rather than replicating from the old backup system to the new primary system, the old primary system (System “A” in figure 3) replicates to the new primary system (System “C”). From there, the new primary replicates to the new backup (System “D”).

In addition to replicating data to the new primary system, the old primary system continues to replicate to the old backup system (System “B”). This ensures that the old backup system remains fully synchronized and ready to take over operations until the new primary and backup are ready to go live.

Migrating to the cloud may be a quick alternative when you find you need to expand your computing capacity quickly.

Once the replication processes have fully loaded the new primary and backup systems and those new systems have been fully tested, users can be switched to the new primary system and the old systems can be decommissioned.

Advantage

Data Integrity: The parallel method reduces the risk of migrating missing or incorrect data to the new servers. Because the migration is done using a new “A” to “C” replication configuration that was carefully planned and designed specifically for the migration process, even if the old “A” to “B” HA configuration was incorrect, the new systems will not be affected by this error.

Challenges

Bandwidth: Using the parallel method, during the migration process the old primary system will have to replicate data and objects to two systems simultaneously. You must be certain that there is adequate bandwidth both on the network and on the channels out of the old primary system to handle this.

Processing workload: Using the parallel method, there will also be some additional processing required on the primary system to manage the dual simultaneous replication streams. If your system is already near capacity (which might be why you are upgrading in the first place), this may result in unacceptable application response times during the upgrade process.

Complexity: The configuration required for the parallel method is slightly more complex than for the cascade method, but it results in the least downtime for users.

Migration Solutions with Vision

Windows Environments

Double-Take Move for Windows gracefully migrates your data while it's still accessible to users, without adverse impact on production systems. You select which data that you want to migrate through an easy Explorer-like interface and Double-Take Move does the hard work, migrating data and metadata such as permissions and attributes.

You can also mix and match storage configurations (Fiber RAID 10 to SATA RAID 5) and replicate just the actual data and not the entire LUN. This further helps you to flexibly migrate data and storage by eliminating configuration boundary constraints imposed by other migration methods. Once the baseline migration is complete, Double-Take Move continues to keep the new target storage synchronized using real-time replication, allowing you to choose when you want to complete the migration based on your schedule – without forcing a cutover when you desire it least, such as during the business day.

IBM Power Systems Environments

MIMIX Move, which migrates IBM i servers, and Double-Take Move for AIX virtually eliminate the downtime traditionally associated with system upgrades and migrations by using Vision's advanced HA technologies to rapidly copy business-critical data to your new IBMi or AIX server and keep it synchronized with your production system in real-time—without downtime. After validating your new environment, you can move users and processes to the new environment in just minutes.

Double-Take Move for AIX and MIMIX Move can keep your old and new systems synchronized indefinitely, allowing you to take your time to validate the new system while users work normally on the old one. While you're performing expert and careful validation, Double-Take Move for AIX or MIMIX Move keeps the new system current by migrating changes from the old system as they happen. Finally, when you're ready to flip the switch, operations will pick up on the new system exactly where they left off on the old one.

When using manual methods or traditional data copy tools, data and objects that are not transferred to the new system are common points of failure for migrations. MIMIX Move eliminates this risk by mirroring IBM i data, user profiles, authorities, data areas, data queues, IFS files, programs, spool files and all other objects needed for a successful Power Systems upgrade or migration. Double-Take Move for AIX achieves risk mitigation by replicating the database, user home directories, third-party programs, and user file systems along with file attributes needed for a successful system upgrade or migration.

Because a virtual machine looks the same as a physical machine to both Double-Take Move for AIX and MIMIX Move, migrating from physical systems to partitioned virtual servers, or vice versa, is as easy as migrating between physical systems or between virtual systems. Migrations can be trying and complex, but with Double-Take Move for AIX and MIMIX Move, you're not in it alone. From start to finish, Vision personnel work closely with you to install and configure the software, enable and monitor synchronization, and perform the final switch.

Cross-Platform Migrations

Occasionally, organizations decide that the current platform is no longer the best one to support a particular business function. Perhaps a new best-of-breed application offers enough advantage over the legacy application to justify migrating to it, but the new and old applications don't run on the same platforms. Or perhaps the organization has adopted a new technology standard and is moving applications to the new platform one at a time as they come up for redevelopment.

Whatever the reason for a cross-platform migration, the old applications will likely have to be replaced with new ones, but the same is definitely not true for the data associated with those applications. That data needs to be migrated to the new platform in its entirety and without jeopardizing any data entered into the production applications while the migration is being performed.

Double-Take Share can meet this objective. It replicates data between either homogeneous or heterogeneous databases, including between different database management systems, running on different hardware and operating system platforms, and using different database schemas. Thus, it is easy to map between the data structures and field names of the old and new databases, without losing any of the data content.

With Double-Take Share in place, you can install your new application, take your time to make sure it is properly configured and tested, while continuing to use the old applications and ensuring that the new and old databases stay synchronized the whole time.

What if you are performing a staged platform migration, moving applications individually over a period that may span years? In that case, the applications on the old and new platforms will likely have to share at least some data across platforms until all applications have been migrated to the new platform. Double-Take Share can help with that too. Simply leave it in place until you finish the last of the migrations. Double-Take Share will keep the databases—or just the required portions of them—on the old and new platforms synchronized for as long as necessary, even if the old and new systems are separated by great distances.

About Vision Solutions

With over 25,000 customers globally, Vision Solutions is one of the industry's largest providers of business continuity and information availability solutions for Windows, Linux, IBM i, AIX, and Cloud environments. Vision's trusted Double-Take, MIMIX and iTERA brands keep business-critical information continuously protected and available. With an emphasis on affordability and ease-of-use, Vision products and services help customers achieve their IT protection and recovery goals, which in-turn improves profitability, productivity, regulation compliance, customer satisfaction and quality of life.

Vision Solutions also offers the tools and competency needed to migrate complex, multi-layered computing environments. We can eliminate the strain on your resources, dramatically reduce server downtime, and offset the risks associated with migrations. Regardless of OS or hypervisor, Vision Solutions offers the technology needed to make every migration a success.

Vision Solutions oversees a global partner network that includes IBM, HP, Microsoft, VMware, Dell and hundreds of resellers and system integrators. Privately held by Thoma Bravo, Inc., Vision Solutions is headquartered in Irvine, California with development, support and sales offices worldwide.

For more information call: 1.800.957.4511 (toll-free U.S. and Canada)
1.801.799.0300 or visit visionsolutions.com.

Also find us on:

Facebook: <http://www.facebook.com/VisionDoubleTake>
Twitter: http://twitter.com/#!/VSI_DoubleTake
YouTube: <http://www.youtube.com/VisionDoubleTake>
Vision Solutions Blog: <http://blog.visionsolutions.com/>



15300 Barranca Parkway
Irvine, CA 92618
800.957.4511
888.674.9495
visionsolutions.com

Double-Take® **MIMIX®** **iTERA**

© Copyright 2013, Vision Solutions, Inc. All rights reserved. IBM and Power Systems are trade- marks of International Business Machines Corporation. Windows is a registered trademark of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds. WP_Migration_woDowntime_EO-121112