IBM Global Technology Services
June 2007

Softek TDMF z/OS extended distance migration.

For a data center relocation or consolidation

Contents

Introduction

The purpose of this paper is to give a brief overview of Softek® TDMF™ (Transparent Data Migration Facility) z/OS®, outline some of the benefits a client may derive from using the product and show how, when combined with available channel extension technology, TDMF may be used for migrating data over virtually unlimited distances.

In April of 1997, Softek announced the availability of a software-based data migration solution called TDMF that allows users to transparently migrate data in an IBM Multiple Virtual Storage (MVS) environment to a designated target disk from a source disk that is in use by other applications. With TDMF software, two types of migrations may be performed.

The first type is a swap in which the VOLSER initially associated with the source disk is dynamically switched to the target disk. Applications using this VOLSER may be left running throughout the migration and upon completion of the swap are unaware that the data they are accessing is located on a different physical disk.

The second type of migration is a point-in-time migration. In this scenario, TDMF software may be used to create a copy of a source volume onto a target disk while the source disk is still in use by other applications. A point-in-time migration does not perform dynamic VOLSER switching. Upon completion of the point-in-time copy, applications continue to access data from the initial source disk.

Highlights

When using the right software, data transfer from the local to the remote site can take place without disruption during normal system operations.

When migrating data with TDMF software, the master LPAR does the actual migration of data from a source to a target disk.

The relocation of workloads to a new location involves a disruption in service to end users. To reduce the impact on business, the operational switch from one data center to the other must take place in the shortest possible time with minimal outage to revenue-generating applications. The desired reduced impact cannot be accomplished if operational data has to be backed up to tape, physically transported and then restored onto the remote system at the alternate location.

Using TDMF software, the service disruption is limited to the time required to close the application on the local site, switch network access from the local to the remote site and then re-IPL (initial program load) the system on the remote site. Data transfer from the local to the remote site can take place without disruption during normal system operations.

To migrate data using TDMF software, it is necessary to establish a session. A TDMF session may be created by running a copy of TDMF software in each of the logical partitions (LPARs) that have access to the source and target Direct Access Storage Device (DASD) subsystems involved in the migrations. If multiple LPARs are involved, TDMF software in one of the LPARs will be designated as the master and the other LPARs will be designated as agents. The master and its associated agents communicate and coordinate migration activities using a common SYSCOM data set. The master does the actual migration of data from a source to a target disk. If only one LPAR has access to the source and target volumes, it runs as a master with no associated agents. A session consists of one master and up to 31 agents each running in different LPARs and all sharing the same SYSCOM data set. An interactive system productivity facility (ISPF)-based TDMF monitor may be brought up to view and communicate with TDMF sessions in progress. The monitor may also be used to look at performance data from both current and past sessions via the SYSCOM data set.

Highlights

Softek TDMF software is simple to use, enables access to data during migration and is not vendor specific.

TDMF capabilities

Design

Softek TDMF z/OS was designed as a fully integrated product that is simple to install, simple to use and simple to maintain.

Transparency

One of the most important requirements of data processing today is continuous availability of data. Traditionally, one of the downsides of the backup or migration of data has been that it requires an interruption of the data's availability to the end user. TDMF software's ability to transfer data transparently can greatly reduce and, in many cases, eliminate periods of data unavailability. This results in increased productivity and reduced operational costs. In addition to its transparent use, TDMF software may be installed dynamically, requiring no IPLs.

Non-vendor specific

Because TDMF software is completely MVS-based, it does not require special micro code or hardware features in the DASD control unit. This means that it may be used in a multivendor environment. TDMF software supports all extended count key data (ECKD) control units. Device support includes standard size 3380 and 3390 images as well as hyper volumes.

Tuning parameters

To adjust performance impact on applications during a migration, TDMF software provides the ability to control migration rate using the following parameters:

- *FASTCOPY—copies only allocated tracks on a volume*
- *FULLSPEED—doubles buffer I/O blocks, nearly doubling TDMF performance*
- *CONCURRENT—establishes the number of concurrent volume migrations*

Highlights

TDMF software provides the ability to control the data migration rate.

TDMF software enables logical grouping of volumes for better operational control. And it provides a view into the entire migration process.

- *SYNCgoal—in a swap migration, establishes how long the swap is allowed to take*
- *PACING—establishes dynamic or fixed, increased or decreased size of the I/O TDMF transfers*

Migration groups

TDMF software supports the definition and concurrent migration of online volumes. Because migration projects may involve several hundred volumes supporting a range of application types, TDMF software provides the ability to logically group volumes for efficient operational control.

Many group migration parameters can be independently configured and controlled to best suit the business applications supported.

Monitoring features

TDMF software provides a Time Sharing Option (TSO) ISPF monitor for managing and viewing a migration process from start to finish. Statistical information includes details such as elapsed time, copy rate, percentage complete and so forth.

Shared DASD

TDMF software works in a shared DASD environment. The DASD can be shared by individual LPARs running on multiple physical CPUs, shared within a sysplex or shared across multiple sysplexes.

Provision of “guaranteed” data integrity

Softek TDMF software was designed to maintain physical data integrity. Source volumes remain untouched by the Softek TDMF session and current data is always available up to the point of the workload transfer when a swap migration is initiated. For a point-in-time copy, the source volume is never touched.

Highlights

The integrity of the data on the source volumes can be maintained even in the rare event of an aborted migration.

Using the DYNAMIC SUSPEND feature of Softek TDMF software, a migration may be temporarily stopped in the event of a hardware or network failure. The Softek TDMF software will continue to monitor the source volume for updates until the migration can be continued from where it was suspended. If, for whatever reason, the migration is aborted or terminated, then the integrity of the data on the source volume(s) is maintained. Unlike other migration tools, Softek TDMF software enables a migration for transferring workloads to be terminated at any time.

TDMF extended distance data migration

TDMF software may be used with existing channel extension technology to implement the product's benefits between data centers separated by thousands of miles. Common uses include remote backup, disaster recovery and data center relocations. Essentially, the use of TDMF software remains unchanged with the exception that some kind of extended distance network connection is set up between the data centers containing the source and target DASD subsystems involved in the migrations. The connection may consist of one or more point-to-point leased lines such as DS-3 (T3) links or may be achieved by connecting to a tariff-based packet switching network such as an asynchronous transfer mode (ATM) "cloud."

In addition to channel extenders or tariff-based packet switching, the latest release of TDMF software can also support the migration of data via TCP/IP. This white paper does not include any topic involving data movement over TCP/IP.

TDMF software can be set up to perform either a push or a pull operation. The terms "push" and "pull" define the direction data is flowing across the network with respect to the TDMF master application.

Highlights

The use of TDMF software remains essentially unchanged by distance.

Softek can be set up to either push or pull data across the network. When pushing, the master resides locally with a source subsystem.

In a push operation, the TDMF master writes data across the network out to the target subsystem(s) residing in the remote data center. In this environment, the master resides on an LPAR in the local data center along with the source subsystem(s).

In a pull operation, the TDMF master resides on an LPAR running at the remote data center and reads data across the network from the source subsystem(s) residing in the local data center.

Diagram 1 illustrates an example of a push operation for a single LPAR environment existing in the local data center. The extended distance link used is a DS-3 (T3) point-to-point leased line. A push operation is useful when target DASD subsystems are running in a remote data center that does not yet have an operational MVS LPAR.

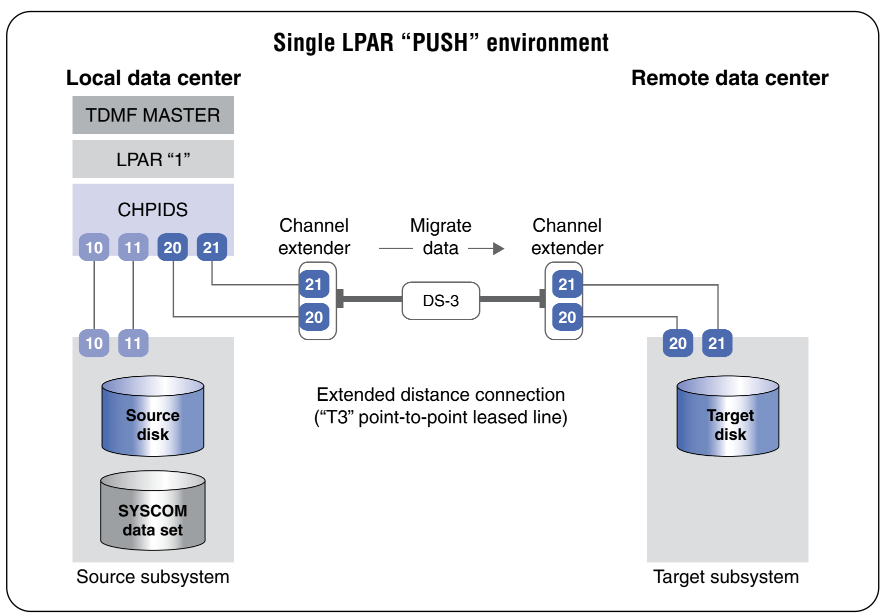

Diagram 1
Single LPAR push environment

Highlights

When pulling, the master resides in the remote data center.

Diagram 2 illustrates an example of a pull operation for a multiple LPAR environment. The extended distance link used is a DS-3 (T3) point-to-point leased line. In this example, both the local and the remote data centers each have an operational MVS LPAR.

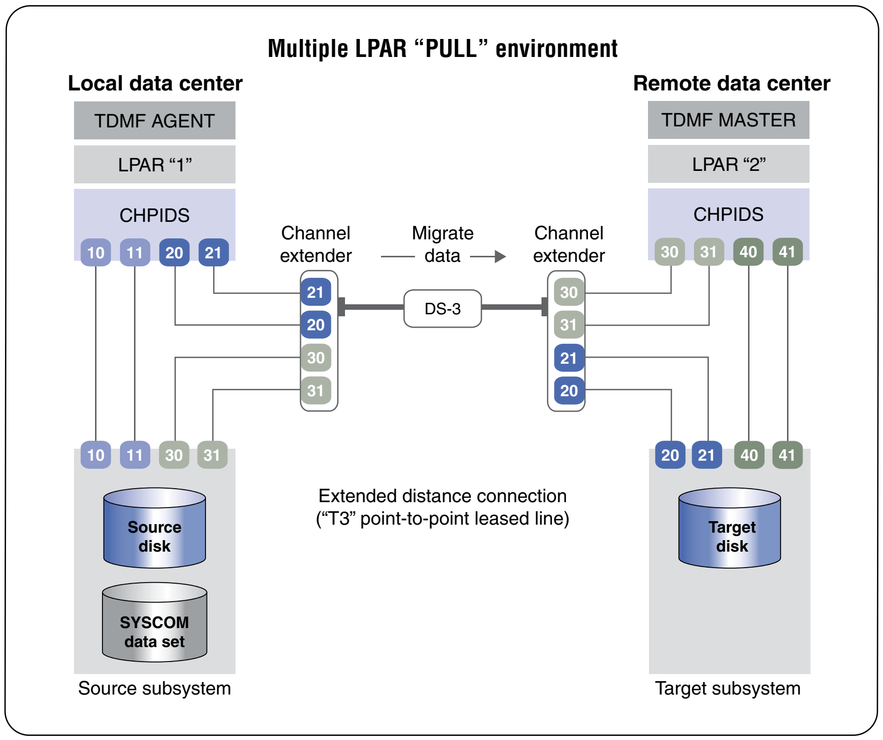

Diagram 2
Multiple LPAR pull environment

Highlights

When relocating a data center, certain operations can reduce both the time and cost of relocation.

When relocating from an existing local data center with an established workload to a new remote data center being brought up for the first time, certain things can be done to reduce both the time and the cost of the relocation. Setting up a pull operation means the TDMF master can take full advantage of the CPU resources installed in the remote data center without impacting the application workload still running in the local data center. This will reduce the amount of time required to perform the relocation.

If the relocation is coordinated by application, it is possible to reduce the network bandwidth needed to move the data. When point-in-time migrations are performed, genning and connecting the local LPAR(s) to the remote target storage (CHIPIDS 20 and 21) is not required. The removal of this requirement can result in less traffic between the data centers, further reducing the network bandwidth needed to migrate the data.

Note: Although not recommended, for swap type migrations, each of the target subsystems must be genned and connected to each of the agent LPARs involved in a session.

Diagram 3 shows a pull operation involving multiple source and target subsystems connected to a pair of DS-3 (T3) links. In this example, the migration will be coordinated by application to reduce the number of DS-3 links required. Volumes related by application will be moved as a group migration requiring a prompt to signal TDMF software of a point-in-time.

Highlights

Softek simplifies relocations, which can result in less traffic between data centers. This can reduce the network bandwidth needed for migrating data.

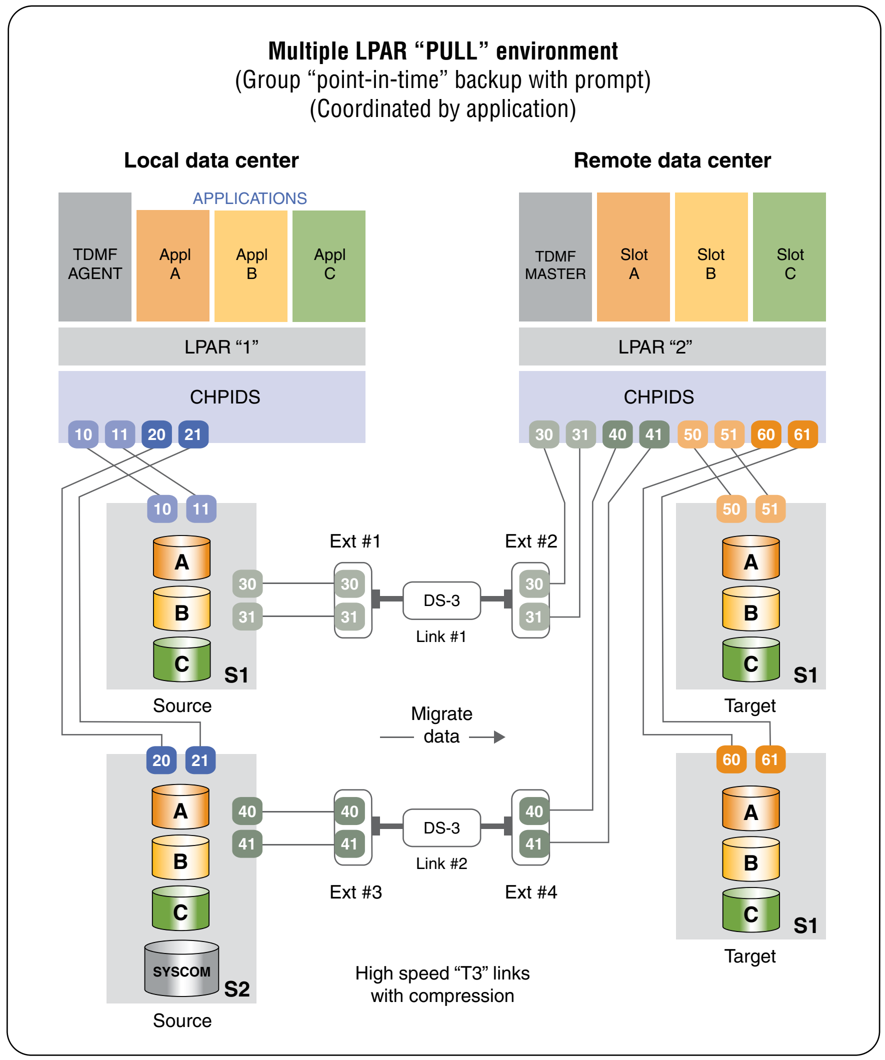

Diagram 3
Multiple LPAR pull environment (Group point-in-time backup with prompt) (Coordinated by application)

Highlights

Group point-in-time migration begins with source volumes used by the first application.

Re-clipping target volumes to the original VOLSERs and varying the source and target volumes are necessary before moving to the next application for migration.

Using TDMF swap migrations can eliminate offline/online and re-clip procedures of point-in-time procedures.

With applications APPL A, APPL B, and APPL C running on LPAR 1 in the local data center, a group point-in-time migration session will be started for the source volumes used by APPL A (depicted as A in source subsystems S1 and S2) out to the target volumes chosen to receive them (shown as A on subsystems T1 and T2). When the master reaches the point where it is ready to end the group A migrations, it will issue a prompt to the TDMF monitor. At this point, the user would bring down APPL A on LPAR 1, then respond to the prompt to complete the migrations for all of the group A volumes.

Before bringing up APPL A in SLOT A of LPAR 2, it will be necessary to vary the source volumes offline to LPAR 2, re-clip the target volumes to the original source VOLSERs, vary the target volumes online to LPAR 2 and catalog data sets on the target volumes as needed. The same procedure would be repeated for migrating APPL B and APPL C. Since Diagram 3 is set up for point-in-time migrations, it is not necessary to gen or connect the target subsystems (T1 and T2) in the remote data center to LPAR 1 running the TDMF agent in the local data center.

If the extended distance network is set up so that source and target volumes are accessible to hosts running in both data centers as in Diagram 2, the relocation of volumes coordinated by application could be carried out using TDMF swap migrations. This would eliminate the vary offline/online and reclip steps described for the point-in-time based procedure since the target volumes would already be online with the correct VOLSERs to the remote host following completion of the swap migrations. Catalog issues for LPAR(s) in the remote data center would still need to be addressed.

Highlights

With TDMF and the right circumstances, it's also possible to migrate database volumes with no further application interruptions.

Using channel extenders can reduce performance degradation. Channel extenders compress data flowing across a link.

If the application's load libraries were copied prior to moving its database volumes, and the appropriate VTAM and/or TCP/IP network(s) were in place, it would be possible to temporarily run the application out of the remote data center accessing database volumes in the local data center and then migrate the database volumes over to the remote data center with no further interruption to the application.

It must be noted, when doing a swap migration over channel extenders, the user could be exposed to server performance degradation after the swap of a volume. It is for this reason Softek always recommends doing a point-in-time migrations involving channel extenders.

A DS-3 (T3) link has a raw (actual) bandwidth of 44.736 Mbps (Mega bits per second). The DS-3 rate divided by eight yields approximately 5.5 MBps (Mega bytes per second) as compared to the effective bandwidth of 17.5 MBps for ESCON channels. The channel extender boxes attached to either end of a link are able to compress the data flowing across the link. The compression ratio achieved is dependent upon the data patterns being moved and will determine the effective bandwidth of the link. Effective bandwidths yielded through compression are typically two to three times the raw bandwidth of a link. For DASD applications, it is usually safer to use a 2-to-1 compression ratio for estimating throughput.

The key to efficiently migrating data across high speed DS-3 links is setting up an environment that allows the capacity of the link(s) to be saturated as nearly as possible. The percent utilization of a link's capacity can be monitored via the channel extender or derived by dividing the actual data rate across the link by the link's effective bandwidth (taking compression into

Highlights

To efficiently migrate data across high-speed DS-3 links, you need to use 90 to 95 percent of a link's capacity.

account). In a favorable environment, between 90 and 95 percent of a link's capacity can be utilized. If the links depicted in Diagram 3 could be made to sustain a 91 percent average utilization factor, the actual data rate across both links could be estimated as follows:

- *Link utilization factor = 0.91*
- *Raw bandwidth/DS-3 link = 5.5 MBps/link*
- *Compression ratio = 2/1*
- *Number of DS-3 links = 2*
- *0.91 x (5.5 MBps/link) x (2/1) x (2 links) = 20 MBps*

Many factors can cause the percent utilization of a link's capacity to change over time. By monitoring the link(s), it is possible to tune migration performed to help sustain a high percentage of link utilization over the course of one or more migration sessions. Leased point-to-point lines such as DS-3 circuits are most cost-effective when they are being utilized at near capacity. The user pays the same amount to lease the line whether it is actually in use or not. These circuits are ideal in situations such as data center relocations where large amounts of data are being moved within a fixed period of time.

Multiplexed T1 circuits may be used for smaller data migrations. Additional hardware may be needed if channel extenders are used.

In situations where smaller amounts of data are being migrated, multiplexed T1 circuits may be used. If data compression is to occur in the channel extenders for data being sent across T1 links, additional hardware in the form of inverse multiplexors may be required. The raw bandwidth of a T1 link is 1.544 Mbps. Approximately 28 T1 links equals the raw bandwidth of a single DS-3 (T3) circuit (44.736 Mbps).

Highlights

Diagram 4 shows compressed data packets that are being transported via a multiplexed high-speed serial interface.

Diagram 4 shows the use of n number of T1 circuits being multiplexed into a high-speed serial interface (HSSI) for the purpose of transporting compressed data packets. Compression ratio estimates for T1 and T3 circuits carrying DASD data are comparable (typically 2:1). If compression is not required, the inverse multiplexors are not needed, and the T1 circuits may be connected to the channel extenders via modem using a V.35 interface.

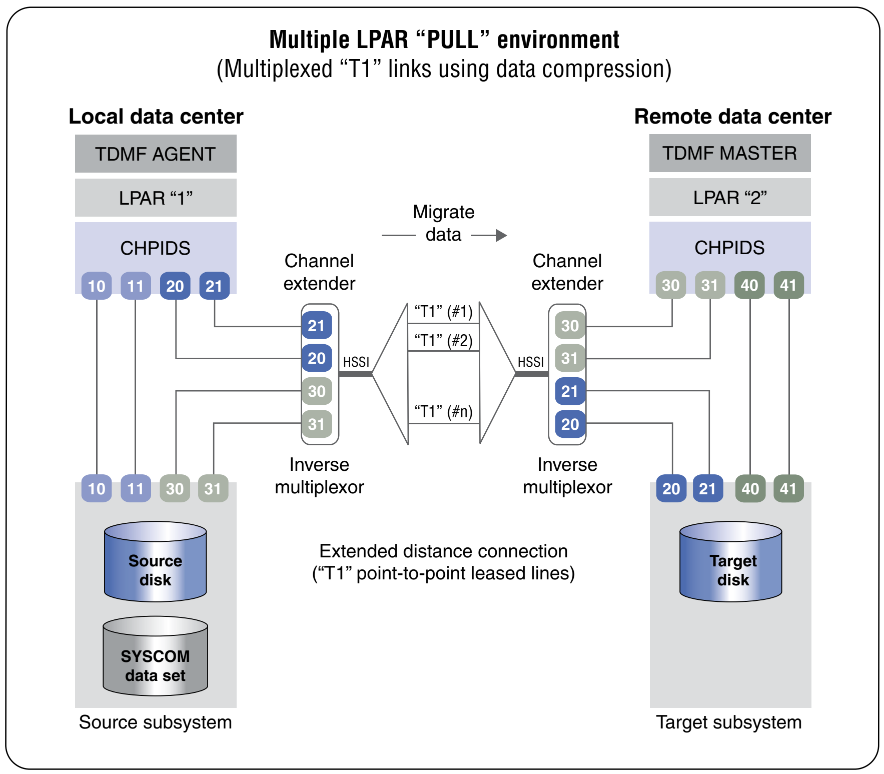

Diagram 4
Multiple LPAR pull environment (Multiplexed "T1" links using data compression)

Highlights

Diagram 5 illustrates a solution for users who have the need to migrate varying amounts of data over extended distances on an ongoing basis. This may be achieved by connecting the data centers involved to a tariff-based packet switching network referred to as an ATM "cloud." The cost of migrating data depends on the number and size of data packets sent. During periods of little or no migration activity, the user does not incur the high cost of leased line(s). Another advantage of the ATM cloud network is that it can be used to connect more than two data centers together.

Diagram 5 depicts an ATM cloud network that can be useful for users who need to migrate varying amounts of data over distances on an ongoing basis.

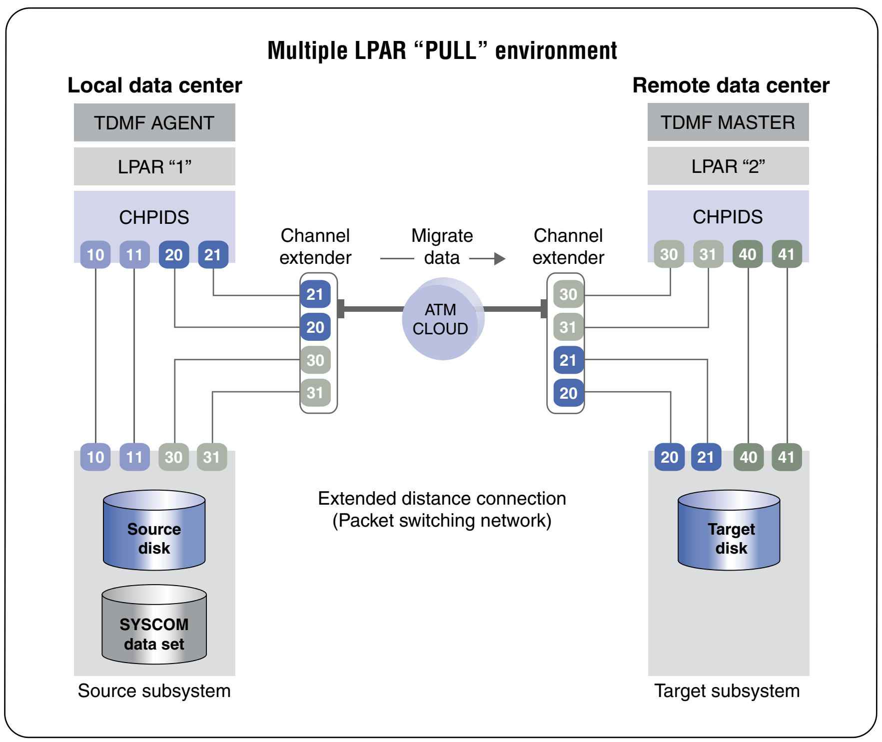

Diagram 5
Multiple LPAR pull environment (ATM "cloud" packet switching network)

Highlights

Achieving a 2:1 compression ratio with 85 percent efficiency doubles the Gigabytes transferred.

ATM clouds may be implemented on SONET (Synchronous Optical NETwork) fiber networks capable of running at 155 Mbps. The throughput yielded from connecting to an ATM network will depend on the backplane capacity of the channel extenders being used. User data throughput across an ATM cloud implemented on a SONET fiber network typically ranges from 80 Mbps to 135 Mbps (10 MBps to 17 MBps). Consult the channel extender vendor for data rate(s) supported by the channel extender equipment to be used.

Sample estimate timings

Table 1 shows the number of gigabytes that may be moved with no compression (1:1 ratio) and an efficiency rate of 85 percent.

Bandwidth	Number of links	Gigabytes/1 hr	Gigabytes/12 hrs	Gigabytes/24 hrs
E3	1	13	156	312
	2	22	265	530
	3	33	397	795
	4	44	530	1061
DS3	1	16	210	403
	2	28	343	686
	3	42	515	1030
	4	57	686	1373
OC3	1	59	711	1422
	2	118	1422	2845
	3	177	2134	4268
	4	237	2845	5691

Please note that DS3 is the formatted digital signal used on T3 lines.

Highlights

If it is possible to achieve a 2:1 compression ratio with 85 percent efficiency, then the Gigabytes transferred in the table above can be doubled.

For specific estimates on the transfer rate that can be achieved in an environment using channel extenders, please contact the hardware vendor.

A typical data center relocation

This extended distance migration example allows relocation with minimal downtime.

The following is an extended distance data migration scenario using TDMF software. It would allow the user to relocate a data center and minimize downtime to the amount of time required to bring down LPAR(s) operating in the local (originating) data center, switch over the network and IPL LPAR(s) in the remote (destination) data center.

The actual migration of data would be performed while the user's production workload is running. The key to this relocation technique is in providing sufficient extended channel bandwidth that would enable the user to run production applications in the remote data center that would temporarily access data residing in the local data center.

The TDMF master would be run in the remote data center. The stage 1 environment can be set up without disrupting data processing at the local center.

Essentially, the user's data would be migrated in two stages:

- *Stage 1 would be to migrate static data (system packs and application loadlibs and parameters) using TDMF software's point-in-time option.*
- *Stage 2 would be to migrate dynamic data (application data) using TDMF software's swap option.*

The TDMF master would be run in the remote data center. The reason for this is that point-in-time migrations do not require target DASD subsystems to be genned to LPARs running agent copies of TDMF software. This allows the

Highlights

The first four steps for performing a data center relocation include setting up the remote data center environment, performing a freeze to static data in the local center, performing Stage 1 of the relocation, and preparing the production environment in the remote data center.

environment for Stage 1 of the relocation to be set up without disrupting data processing in the local data center, because a new gen would not be required for the production LPAR(s).

Outlined below are the steps required to perform the data center relocation:

1. *Setup initial remote data center environment. This includes:*
   - *Creating a skeleton MVS in remote data center capable of running the TDMF master*
   - *Installing target DASD subsystems in remote data center intended to receive the data.*
   - *Genning source DASD in local data center and target DASD in remote data center to the skeleton MVS.*
   - *Establishing extended channel connectivity between the skeleton MVS in the remote data center and the source DASD subsystems in the local data center.*
2. *Perform a freeze on update activity to static data in the local data center which will be needed to setup a production environment in the remote data center. This data includes system packs as well as volumes containing application loadlibs and parameters.*
3. *Perform Stage 1 of the relocation (use TDMF software's point-in-time option to migrate static data from the local (original) data center to the remote data center).*
4. *Prepare a production environment in the remote (destination) data center using the static data migrated from step 3. This includes:*
   - *Setting up system definitions for production LPAR(s) to run in the remote data center. Cataloging source VOLSERs that will be required by applications that are to run in remote data center LPAR(s).*

Highlights

The last two steps include bringing down production LPAR(s) in the local data center and performing Stage 2 of the relocation using TDMF software's swap option.

Watch the network requirements between the two locations and keep the communication links to a minimum of T3/E3 links.

- *Testing the functionality of applications brought over from the local data center.*
- *Adding the target DASD subsystems to the I/O gen(s) of production LPAR(s) intended to run in the remote (destination) data center.*
- *Setting up network access to applications to be brought up in the remote data center.*

5. *Bring down production LPAR(s) in the local (original) data center, switch over the network, and IPL production LPAR(s) in the remote (destination) data center. Note: At this point, the user would be running production applications in the remote data center accessing "dynamic" data across extended channel connections in the local (original) data center. It is important that the user has set up sufficient extended channel bandwidth to prevent a bottleneck between data centers.*

6. *Perform Stage 2 of the relocation [use TDMF software's swap option to migrate dynamic application data from source DASD subsystems in the local (original) data center to target DASD subsystems in the remote (destination) data center]. Because these migrations are transparent to applications using the data, they may be broken up into smaller groups of volumes to reduce the use of extended channel bandwidth being shared with other applications. When all of the data has migrated, extended channel connections to the local (original) data center may be removed thus completing the data center relocation.*

Data center relocation considerations

Below are some considerations to keep in mind when planning data center relocation using TDMF software:

1. *Network requirements between current and new location.*

- *There are really only a few players in the channel extension business today. The average lead time for building the extenders is approximately two months*

Highlights

based on requirements. The same channel extenders do not support parallel channels. Usually the minimum contract is two months if the user is planning on leasing.

- *Communication links should be at the very minimum T3/E3 links. T1 links are too slow. Using existing networks within the user's environment is not recommended as the traffic could impact production applications. There could be a lead time associated with the lease of these links. Additionally, the term of the lease could come into play.*

Ensure that you have the right CPU memory for Softek TDMF point-in-time copies.

2. *CPU memory requirements to complete the Softek TDMF point-in-time copies to the new location.*

- *The total amount of memory necessary for a migration is dependent upon the number of LPARs participating in the sessions and the number of volumes to be moved. The intent is to move the largest amount of volumes using the least amount of memory. To this end, it is important to consider whether or not any LPARs (such as test LPARs) might be shut down during the migration. Also, it is suggested that a review of all volumes be done to identify any volumes that need not be moved, such as page packs, work packs, etc.*
- *It may be that there is not enough memory available in the production environment. Having the test LPAR(s) shut down would make that storage potentially available either by DRM if all storage is dynamically re-configurable or a power-on-reset (POR) may be necessary in order to access that storage.*
- *If the total memory requirement is more than any one LPAR could handle, what can be done in this situation is to balance the load across the participating LPARs so that no one partition is saturated. This would result in more than one LPAR running master sessions in a push scenario as laid out in a worksheet.*

Highlights

Make certain to use the point-in-time function rather than the swap function.

When a long-term session occurs, keep in mind that there is no function within the product to pick up where it left off.

Determine whether you want a push or a pull process.

3. Swap migration versus point-in-time copy.

- *It is recommended that the point-in-time copy function be used rather than the swap function. The point-in-time function provides a fall back position in the event of a problem. Additionally, this would provide the ability to ensure that system and application feature/function work at the new site without impacting production at the current location (assumes net-new or asset swap equipment).*

4. Long-term sessions.

- *Determining the duration of the migration is dependent on the bandwidth available on the links, the channel configuration over the channel-extenders and the amount of transferred data being moved. Using Softek TDMF software will result in only allocated tracks being transferred; that is, if only 70 percent of a volume is allocated, then only 70 percent of the volume needs to be migrated.*
- *Keep in mind that during the data migration, Softek TDMF software will not survive an IPL of one of the participating systems; all sessions affected would have to be started from the beginning. There is no function within the product to pick up where it was left off.*

5. Push versus pull process.

- *A push operation is where all the sessions are at the current location and the master sessions are pushing the data to the new location. No other LPARs participate outside of the current location.*
  - *The benefit of a push operation is that there are no extra systems to be included in a session.*
  - *The target devices would be defined via channel-extension to the LPAR(s) executing the master sessions.*
  - *Agent sessions would be run on all LPARs participating in the migration.*
  - *If the communications link(s) fail for whatever reason, the sessions may be suspended and then continued when the links are reestablished.*

Highlights

A multi-hop migration with loaner storage equipment may be more costly but can allow users to use a vendor's remote copy facility. Benefits of multi-hop migration include exercising the new environment prior to cutting over to it.

- *A pull operation is where the master session(s) are executed in the new location, and the LPARs at the current location participate as agent systems.*
  - *This assumes that there are new mainframe server(s) at the new location with nothing executing on them outside of the LPARs executing the Softek TDMF sessions and, therefore, there would be sufficient memory available for the master sessions.*
  - *The benefit would be that the current production LPARs would only be acting as Agents systems and would significantly reduce the memory requirements on these LPARs.*
  - *The target devices would be defined only to the new LPARs. The source devices would be defined via channel-extension to this environment. The agent systems are not required to have the target volume defined in a point-in-time situation.*
  - *If the communication link(s) fail, a 15-minute limitation would come into play as the master system(s) would not be able to communicate to the communications data set (COMMDS). If the link(s) have not been reestablished within the 15-minute limit, the sessions will fail.*

Alternate method—a combination of local and remote migrations

Although this may be a more costly solution, organizations have used a combination of TDMF local swap migrations along with hardware controller-based remote replication to accomplish a data center relocation. Some have referred to this as a multi-hop migration using loaner equipment (storage) at the current site.

Highlights

The basic reason for using TDMF software in this scenario is to place the data on a storage subsystem that is compatible with a remote storage subsystem. Users now have the ability to use a vendor's remote copy facility such as IBM's Peer-to-Peer Remote Copy (PPRC).

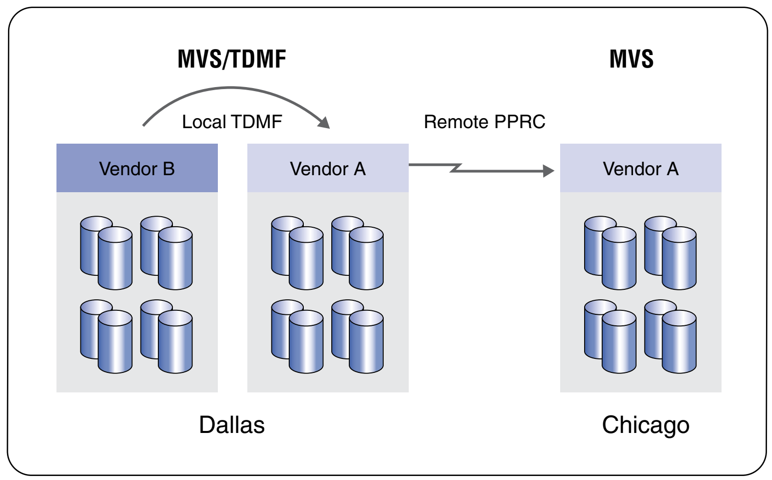

There are some advantages with this type of migration. Assuming that the new storage is identical at both the current and new locations, there is an opportunity to exercise the new storage with the production environment prior to cutting over to the new remote site. That is, do a TDMF swap migration to the new local disk and allow PPRC to do the remote copy to the new location. After the local migration is complete, the user can continue to run production while monitoring the new subsystem at the current site for both functionality and performance. When satisfied that everything is working fine, the user can schedule a shutdown at the local site and bring up production at the new remote site.

Another variation might use TDMF software's perpetual point-in-time option.

A variation of the above could be to use TDMF software's perpetual point-in-time option. In this scenario, it is possible to create a local replication of the production DASD using perpetual point-in-time and allow PPRC to copy the data to the remote site. At some point, a controlled point-in-time mirror

Highlights

image of the production environment can be created at the new location. Stop the PPRC session and test the relocation procedure at the remote site. When satisfied that everything will work, start a new fresh PPRC session and create a new point-in-time mirror image for the final migration. Using perpetual point-in-time, it is not necessary to recopy all the DASD from the start but just copy the delta changes since the last perpetual point-in-time copy was created. It is possible to repeat this testing for an unlimited number of times.

Alternate method—use of TDMF TCP/IP option

Using the TDMF TCP/IP option may be a cost-effective method of moving data.

In some environments, the use of the TDMF TCP/IP option may be a cost-effective method of doing a data center relocation. The cost savings is gained by not having to acquire expensive channel extenders or establishing costly hardware replication features such as EMC's SRDF.

It should be noted that, generally speaking, the transmission bandwidth is much lower and more difficult to achieve parallel data transmission, thereby extending the length of time a migration will take. Also, if the TCP/IP lines are shared with other functions such as online communications then the overall time it takes for a TDMF migration will be impacted.

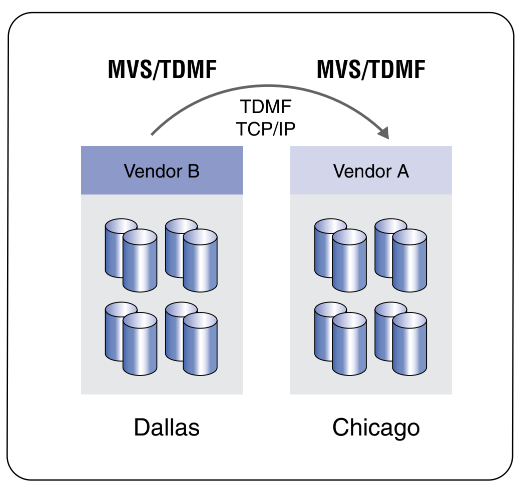

Highlights

With this type of replication, no hardware is available to allow the local or remote operating system access to the DASD.

Shutdowns would be necessary to create the final TDMF copies of the source DASD.

With a TCP/IP replication, it not possible to do a swap migration because as a volume completes the copy/synchronization phase it is not possible (nor desirable) to physically swap that volume to the new remote location while production continues to run at the current local site. Using TCP/IP does require a point-in-time replication.

Note that with this type of replication there is no hardware in place for either the local or remote operating system to have access to the DASD at the other location. The way TDMF software accomplishes this replication is by having a TDMF master read the data off the local source volume and then pass the data to a TDMF master running at the remote location. The remote TDMF master will, in turn, write the data to the target volume.

When all required volumes are replicated to the remote site, it is necessary to shut down all applications, create the final TDMF point-in-time copy of the source DASD onto the target DASD and then shut down MVS at both locations. The shutdown is required because it is necessary to clip or re-label the volumes at the new remote location to the VOLID of the original source volumes. TDMF software can create the ICKDSF control statements to re-label the volumes.

One of the drawbacks to this type of relocation is that in order to test the relocation effort, it is necessary to terminate the TDMF session, clip the volumes at the new location, and test and then restart the TDMF relocation effort again from the beginning. Production at the current location does continue to run at all times other than a short outage to synchronize all local and remote storage just prior to terminating the TDMF session.

Highlights

It's possible to use a hardware mirror facility to create a second copy of the DASD, allowing continued production.

There is an option that could help facilitate the testing effort without restarting the entire migration again from the beginning. It is possible to use a hardware mirror facility such as IBM's PPRC to create a second copy of the DASD at the remote location. This would allow the user to continue production in the original local data center, continue to keep replicating current source volumes to the remote target volumes, and test the relocation effort using the second PPRC copy. Note that it would still be necessary to quiesce the applications in the production environment to create a consistent copy of the test volumes. To further assist in this effort, one method is to use the optional perpetual point-in-time feature. This would allow the creation of as many point-in-time copies of the source DASD as required, all of which can be accomplished without restarting the TDMF replication from the beginning but rather just sending the updates that occurred since the last point-in-time.

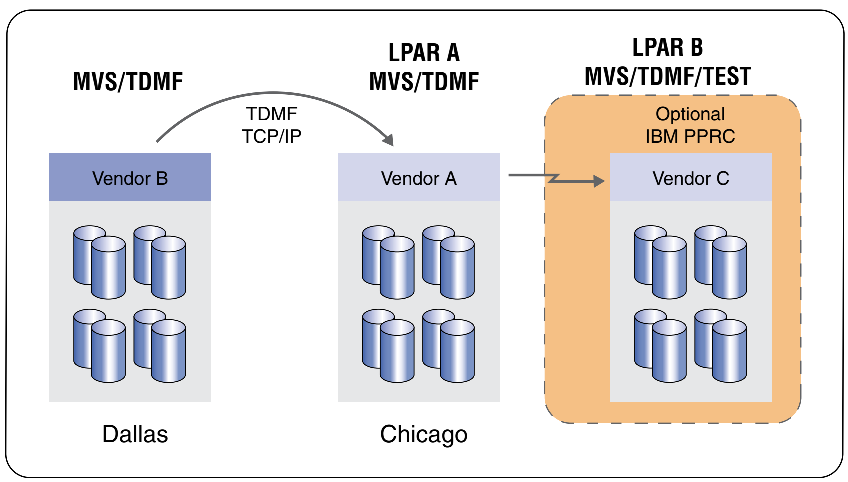

When satisfied that the testing went well, the user can execute the procedures as outlined above and complete the final TDMF TCP/IP relocation copy. Again, this could all be accomplished by starting the TDMF replication once and then just sending the new source updates to the target when required between testing period and the final relocation cutover.

Highlights

When only an application needs to be relocated, difficulties arise because the application resides intermixed with applications that don't need to be relocated.

Using Softek LDMF and TDMF software can help relocate an application with little effort and virtually no outage impact.

Alternate method—use LDMF software to isolate an application for relocation

In some cases, relocating a data center is not granular enough when all that is really required is an application or two to be relocated. The issue most of the time is that the application(s) is not nicely located on its own volumes but, rather, intermixed with many other applications that are not being relocated. Some examples of this could be a business unit is being relocated or sold or an application that needs to be moved from a test/development environment to a production environment located in a different data center.

Using a combination of Softek Logical Data Migration Facility (LDMF™) and TDMF software can accomplish the relocation with little effort and no outage impact to the application while it continues to run production until the final TDMF cutover. At cutover, the outage will be only for a very short period.

Once the files of the application have been identified, LDMF software can be used to migrate those datasets onto dedicated volumes. (See the white paper *Implementing LDMF z/OS for simple, effective, and nondisruptive data migrations*.) Now that the application(s) is isolated on its own volumes, there are multiple options for relocating the application. One way would be use the traditional tape dump/restore, but that would require an extended outage. Another choice is to use TDMF software as discussed previously to relocate the application volumes to a new remote location with only a minimal application outage required.

Other Softek TDMF solutions

Softek TDMF z/OS technology is designed for volume-level migrations involving movement over local or remote distances (also known as "global migrations"). Data can be moved across a TCP/IP network (LAN or WAN) or channel extenders.

Highlights

A benefit of using Softek TDMF z/OS technology is its ability to reduce out-of-service database backup time.

TDMF software can perform backups while the database is active and available to users.

TDMF software has multiple other uses including migrations for packs that need maintenance and for rebalancing workloads for better performance.

Reduction of routine database backup time

One of the benefits offered by TDMF software is its ability to significantly reduce routine out-of-service database backup time. TDMF software's ability to perform a point-in-time backup with a prompt, coupled with its ability to handle a group of related volumes as a single entity, enables it to perform the majority of backups while the database region is active and available to users. The region would have to be down for only a small percentage of time at the end of the backup, because database transactions are buffered in CPU memory before being written out to disk. When TDMF software signals with a prompt that it is ready to complete the migrations for a group of related volumes, the database region may be brought down for a short period, which will allow any remaining buffered updates to be written out to disk. This ensures that the logical and physical images of each of the source volumes match.

Once the region is down, TDMF software's prompt may be responded to, allowing TDMF software to finish backing up the relatively few remaining updates contained on the source volumes. When the backup is complete, the region may be brought up and put back into service. The point-in-time copy of the database created on the set of target volumes may then be written to tape using the user's conventional disk-to-tape package.

Maintenance

Users may use TDMF software to migrate off of packs that have not yet failed but are giving indications that maintenance is required.

Performance tuning

As the dynamics of various DASD subsystems within a shop change, TDMF software can be used to nondisruptively rebalance workloads for better overall performance.

Highlights

The software can also be used to help reduce lease overlap and to create copies of production data for various reasons.

Softek LDMF software can assist with data migration as well. And Softek TDMF for open systems works with several platforms including IBM AIX, Windows NT and other environments.

Lease considerations

For users who lease DASD, TDMF software can be used to help manage and significantly reduce lease overlap between old and new equipment.

Application testing

TDMF software may be used to create copies of production data that application programmers may use to develop modifications required for processing data.

Other Softek migration solutions

Softek LDMF software

Softek LDMF z/OS has the ability to nondisruptively switch files from a source volume to a target volume. LDMF software moves the applications dynamically onto new storage. This switchover feature, which occurs under user control, results in redirection of an application's I/O function (e.g., from the original source to the target). This occurs without any disruption to the application. Softek LDMF software is designed and optimized specifically for local data migrations at the file extent level. Softek LDMF and TDMF software can both be used in the same environment to address different migration project requirements. LDMF and TDMF software provide a very fast, easy, optimized migration solution for data migrations

Softek TDMF software for open systems

For open systems, Softek also offers TDMF software for the open systems platforms. The platforms include IBM AIX®, HP-UX, Sun Solaris, Linux® and Microsoft® Windows® NT®, 2000, 2003.

Softek Replicator for open systems

Softek Replicator is versatile multiplatform data replication software that enables local and offsite disaster recovery and eliminates backup windows.

Highlights

Softek TDMF software helps with your data migration needs while maintaining application availability.

Summary

Today's business-critical applications must be available 24x7, with no down-time window for data migration or relocation. Softek TDMF software, the standard in data migration, gives users the freedom and the power to move data from any storage to any storage, on any platform, over any distance, at anytime—with no interruption to active applications.

Softek TDMF z/OS is easy to use, flexible and transparent to applications. Whether the need is to upgrade to new storage and/or new server hardware (from any vendor, to any vendor), consolidate or relocate a data center, or find a more effective way to migrate data based on cost/performance, TDMF software provides the solution, all while maintaining application availability.

TDMF definitions

The following terms are commonly used when discussing a migration strategy for transferring workloads:

TDMF session: A migration group(s) and/or migration pairs to be processed in a single Softek TDMF execution.

Master: The Softek TDMF system running as an MVS batch job that is responsible for the data copy function. There can be only one master system in a Softek TDMF session.

Agent: An associated Softek TDMF MVS image running in a shared storage environment with the master. To ensure data integrity, any MVS LPAR that has access to the source or target volumes must be running the TDMF master or one of the agent systems. The master and associated agent systems communicate via a shared system communications data set (COMMDS).

Highlights

These are some of the terms used when discussing a data migration strategy.

Source: The originating DASD volume(s) containing the data to be migrated.

Target: The destination DASD volume(s) receiving the migrated data.

Migration pair: Source and target volumes for a single migration.

Migration group: A group of volume pairs with the same group name.

Synchronization: The collection of the final set of updates from all source volumes in a group or session applied to the target volumes. For a point-in-time copy, synchronization involves quiescing any source application systems to ensure that all buffers are flushed.

Swap migration: A migration session in which the source VOLSER is switched to the target volume at the end of the session.

Point-in-time migration: A migration session in which the VOLSER on the source volume is not switched, and remains on-line to the application on the original volume.

Local: Equivalent to source in an extended distance migration.

Remote: Equivalent to target in an extended distance migration.

Push migration: An extended distance migration in which the Softek TDMF master runs on the local system and makes point-in-time copies to remote volumes using channel extenders and communication links. The remote volumes do not have to be attached to any processor.

Pull migration: An extended distance migration in which the Softek TDMF master runs on the remote system and makes point-in-time copies from the local volumes to the remote volumes using channel extenders and communication links. The local system(s) run only a Softek TDMF agent.

Gen/Genning: To generate or produce something according to an algorithm or program or set of rules (the opposite of parse).

Mbps: megabits per second

MBps: megabytes per second

Loadlibs: load libraries

For more information

For more information about TDMF z/OS, visit:

ibm.com/services/storage/migration

IBM Global Services
Route 100
Somers, NY 10589
U.S.A.

Produced in the United States of America
06-07

GTW01276-USEN-00