HITACHI
Inspire the Next

Reducing Costs and Risks for Data Migrations

Data Migration Best Practices and Recommended Storage System–based Data Migration Techniques

White Paper

By Patrick Allaire, Justin Augat, Joe Jose and David Merrill

February 2010

Executive Summary

Forrester 2007 research has shown that in excess of 70 percent of enterprise IT budgets is devoted to maintaining existing infrastructure. Organizations need to accelerate innovation and reduce operating expenses to increase their competitiveness or maintain their current market position, CIOs are under pressure to identify and adopt best practices to lower their IT operating expenses and redirect the savings in support of new investments.

Major drivers of growing data center operating expenses include ongoing power, cooling and labor costs. IT organizations are in a constant flux, handling change management due to storage growth, consolidations, mergers and acquisitions, vendor product life cycles, and interoperability and supportability requirements. For example, most enterprise IT shops have to plan for their storage platforms' end of life and migration before the end of the four to five years of the storage product's useful life. This white paper is intended to provide IT decision makers with information on costs, risks and considerations regarding the migration of data from old storage platforms to new storage to reduce this significant area of operating expenses.

To understand the strategies and best practices organizations were using to reduce cost and risk with storage migration, Hitachi Data Systems sponsored a survey in partnership with TechValidate, an independent market research firm. Key findings from the survey that illustrated the expense of migration and best practices to reduce the cost and risk include:

- Migration project expenditures are on average 200 percent of the acquisition cost of enterprise storage. With an average of four years useful life, the annual operating expenses associated to migration represent ~50 percent of acquisition cost.
- Enterprise storage migration costs can exceed US$15,000 per terabyte migrated.
- Storage migration projects required four to six hours per host, from internal organization resources. Of these hours, four to five hours were used to plan the migration and one to two hours (~30 percent) were used to execute the migration.
- Duration of the migration is mainly due to limited maintenance windows. Common migration techniques require application outages due to either SAN rezoning and/or host reboot activities.
- The two biggest concerns that organizations face during a data migration are the risk of downtime or extended downtime/impact to the business and the budget overrun of the migration project. Seventy percent of customers reported schedule overruns of about 30 percent while 64 percent reported average budget overruns of 16 percent.
- The leading indicator of schedule and budget overruns was the team member experience.

Organizations are spending considerable resources and assuming more risk than necessary to conduct data migrations. There are many new approaches using virtualization that greatly reduce cost and risk. To successfully migrate data with these new approaches, enterprise IT executives should leverage a vendor who understands the all of the approaches and has experience migrating various platforms with multiple technologies. Hitachi Data Systems has a broad portfolio of hardware, software and services and has experience with all industry leading software and hardware to help organizations reduce risk and costs for their data migrations now and for the future. Also, with emerging technologies such as Hitachi High Availability Manager software, which enables storage controller–based migration, organizations will be able to completely eliminate the outage window due to data migrations.

By reducing risk and costs for data migrations, organizations will reduce operating costs due to data migrations and more easily prevent technological obsolecence. Enterprise IT executives should use the ideas in this paper to research this topic further to understand how much data migration costs their organization. These costs will only increase. Hitachi Data Systems expects the industry average cost of enterprise storage migration to continue to rise due to increasing 24/7 application availability requirements and increases in cost of labor. On the basis of this research, it is highly recommended that IT executives evaluate new storage platforms in conjunction with migration solutions to lower data center operating expenses.

For best results use Acrobat Reader 8.0

Contents

Reducing Costs and Risks for Data Migrations

Data Migration Best Practices and Recommended Storage System–based Data Migration Techniques

White Paper

By Patrick Allaire, Justin Augat, Joe Jose and David Merrill

The Importance of Data Migration

Data Migration is an important event that consumes significant budget and labor and occurs very regularly. The combination of the frequency of and resources consumed in a data migration results in data migration taking a significant amount of the IT budget. As storage infrastructures become larger and more complex, data migrations are becoming more complex, risky and labor intensive. Organizations must begin managing this growing portion of their IT budgets more effectively.

According to a 2005 survey conducted by ESG, data migration projects are constantly being conducted by IT managers: "Attesting to the fact that data migration is a fact of life for the IT manager, 39 percent of those surveyed said that they perform migration on a weekly or monthly basis." While organizations are continually migrating data, they are not considering data migrations a core competency as data migration is usually the result of another event in the data center, such as an application upgrade, data center consolidation project or technology refresh. According to Gartner:

> *Interest in data migration and conversion risks and best practices is on the increase as a result of contemporary business drivers, including a mandate for IT modernization (replacing legacy, risky and non-strategic technology) and cost optimization (reducing the cost of IT through consolidation and efficiency improvement). Gartner client inquiry trends clearly reflect this, with the volume of inquiries about data migration issues and practices up more than 50 percent in 2008, compared to 2007.*
>
> *— Gartner, "Risks and Challenges in Data Migrations and Conversions," February 2009, ID Number: G00165710*

Data migration projects can be very complex, large scale projects requiring many in-house and contractor personnel. As a result, the labor, consulting, software and hardware for data migration have become a very large market. The overall market for data migrations can be calculated by identifying amount of data migration activity that results in large data migrations. Many data migrations are a result of technology refreshes. For example, an average FORTUNE 1000® company has an average of 800TB of network attached storage (NAS) and nearly 3PB of storage (InfoPro Wave 12–Q2, 2009) with, on average, 300TB per storage system. As the useful life of most storage systems is three to five years, this size of organization will often, at any given time, have multiple storage systems at the end of their useful life spans and requiring a refresh. Therefore, these large enterprises could be in a position where they would be always conducting a data migration of multiple storage devices at any given time.

We can assess the overall market for data migrations as a result of these storage technology refreshes by using the overall storage market revenue forecast at US*$7.5 billion in 2009 (for storage systems > $150,000) and data from our survey. Data from our survey indicated that when all the labor, resources and equipment required for conducting a data migration are included, the cost of the data migration was twice that of the acquisition price. Therefore, the overall data migration market would be $15 billion in operating expense to support

**All monetary figures in this document are provided in US dollars.*

storage technology refresh in 2009. This type of market size indicates a very large operating expense; organizations should focus more attention on defining best practices and technology to reduce this large expense.

Description of Enterprise Storage Migration Survey

In all cases, the statistics and numbers referenced in this document should not be use in a rigid ways; the reader should look at these findings as an industry average, which fluctuates based on sampling frame. Changes in size of organization respondent, industry, geographic location, and labor cost and migration technique selected are just examples of factors which influence migration project statistics. Due to the fact that Hitachi Data Systems did not force an equal sample of respondents for each migration technique, comparisons between migration technique (as shown in Figure 1) are limited to host and virtualization migration techniques only; the number of observations was too low for the other techniques to make a valid comparison.

Figure 1. Migration Technique Used

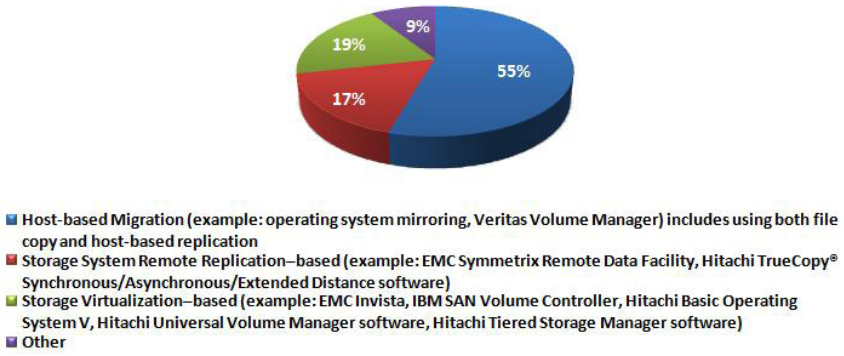

In the sampling frame for this research were respondents who have storage responsibility (from CIO to storage administrator) and were part of an enterprise organization (>1,000 employee US companies, >500 employee non-US-based organization); 50 percent were Hitachi Data Systems customers and 50 percent were non-Hitachi Data Systems customers. The respondent organizations needed to be located within countries that were English (reading) proficient (for example, United States, Canada, UK, Germany, Switzerland, Belgium, Australia, New Zealand, Singapore, etc).

Respondents' industries and sizes of migration projects (in terms of size of data set migrated, number of storage frames and number of hosts impacted) were widespread. For a detailed breakdown of respondents, go to *Appendix A — Survey Demographics.*

Considering that respondents reported in excess of 70 percent of the storage migration cost is driven by labor cost, it is important to factor in the full time equivalent (FTE) costs reported in this research of $126,000 in contrast with IT average labor cost for an industry or organization. FTE cost does not equal employee salary; frequently, an organization's FTE cost may be closer to a ratio of 1.25 to 1.5 times the employee salary. This salary-to-FTE ratio varies based on organization overhead and geographic disparity in social benefits. For a more detailed breakdown of the wide range of reported FTE cost please go to *Appendix A — Survey Demographics*.

Key Risk Factors for Data Migrations

Data migration projects are critical to the success of the initiatives that the migrations support; they impact business critical data, applications and systems, and result in significant cost. The data migration project itself possesses significant risks and requires proper planning and attention to ensure success of the data migration and the initiatives that it supports, which could be an enterprise-wide application upgrade, a data center consolidation or an infrastructure upgrade.

However, as Gartner indicates, many organizations do not place the proper importance on and take proper consideration of the data migration:

> *Analysis of data migration projects over the years has shown that they meet with mixed results. While mission-critical to the success of the business initiatives they are meant to facilitate and support, lack of planning structure and attention to risks causes many data migration efforts fail.*
>
> *— Gartner, "Risks and Challenges in Data Migrations and Conversions," February 2009, ID Number: G00165710*

In our survey research, survey participants identified many risks for data migration projects, as indicated in Figure 2.

Figure 2. Perceived Risks in the Migration Planning

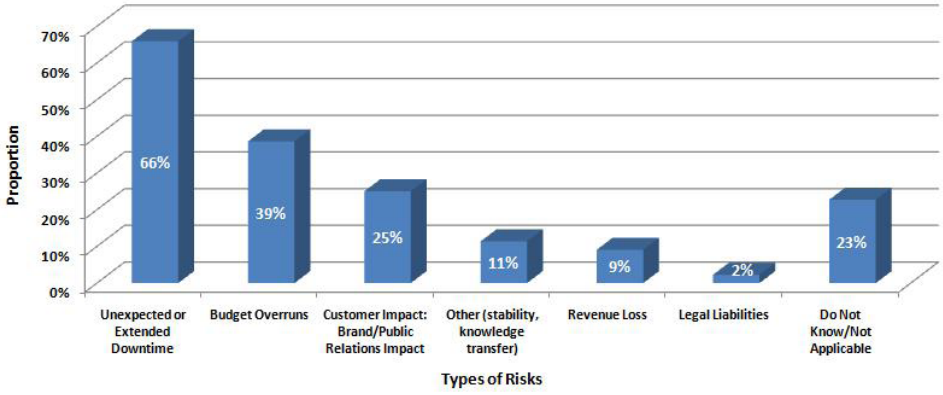

We will focus on a few key areas of risk for data migrations: downtime or loss of data, schedule overrun, budget overrun and customer/brand impact.

Unexpected or Extended Downtime: Organizations Always Need to Have Their Information

Very detailed and careful planning needs to take place to clearly identify windows in which downtime is acceptable and ensure that no data is lost. For data migration projects that include mission critical business data, the risk of impacting sales operations is high; the loss of availability or access to the data could directly impact the profit and loss of the business.

Trends in compute architectures and the adoption of virtualization have increased the number of applications running on a single host. This has resulted in decreasing maintenance windows where the downtime to a host can result in multiple significant application outages. Figure 3 illustrates the increasing density of applications due to virtualization and adoption of multicore computing architectures.

Figure 3. Increasing Application Density

Schedule Overrun: IT Project Management

According to the participants in our survey, only 64 percent of data migration projects were completed within 10 percent of the planned date. When data migration projects last longer than planned, resources are used for longer than anticipated. The schedule overrun can also impact or delay other planned IT projects and impact the business. Often data migration is part of a larger overall project, and the delay in the data migration project affects the success of the larger overall project.

Budget Overrun

As with all IT projects, a significant risk is the cost associated with any budget overrun. This can directly affect budgeting for other IT projects and affect the profit-loss statement for the business. As identified previously, data migrations are part of a larger project. The budget overrun of the data migration project reduces the overall cost/benefit for the larger overall project.

Customer/Brand Impact

Data migration projects can involve customer business data. If any of that data is lost or a customer's access to the data is interrupted, there could be very severe negative public relations impact. The trend for organizations to have more interconnected applications between customers, suppliers and partners, where downtime in one application affects multiple applications, increases the magnitude of the negative impact of a loss of data or availability of data. Also, there could be legal implications and revenue impacts for the loss of data or availability of data.

Government Organization Migrating to a New Virtualized, Multitier Environment

Situation

- Migrating 60 servers with 47TB from an HP XP 1024 system to a Hitachi Universal Storage Platform® V

Solution

- Virtualized existing storage with Universal Storage Platform V technology and remotely replicate data
- Initiated replication in both directions and migrated single tier to three tiers

Key Migration Challenges Overcome

- Required combination of host-based and virtualization methods
- Involved two separate migration windows and two sets of procedures
- Comprehensive plan addressed multiple host-based migration methods; a difference in migration speed was seen, with the virtualized storage approach at 1TB/hr per server and the host-based method at 4hr/TB

Key Cost Factors for Data Migration

Storage Economics from Hitachi Data Systems is a methodology to identify, isolate, characterize and measure costs (of storage) so that actions can be taken to reduce total costs. Within these methods, Hitachi Data Systems has defined 33 different types of costs that apply to storage total cost of ownership (TCO).

Some costs are hard costs, which means they directly impact budgets and expenditures. An example of a hard cost would be power and cooling costs for storage systems. This is a real cost that appears in the IT organization to provide power. Soft costs, on the other hand, are those that are quantifiable but may not produce tangible results in a budget. Reducing a storage management action may save an administrator two hours of work, and although the savings of two hours has many benefits, the administrator would not be paid any less due to this change.

Storage costs are also highly dependent on the organization that pays for the cost. Not all costs are simply rolled into a single management budget. Rather, they are often spread between several organizations. The distribution of costs can sometimes add to the softness of savings to be measured. Best practices in cost reductions tend to occur when an “economic hero” emerges to take ownership and provide common cost metrics for the entire organization. These econometrics can provide the stimulus for continuous improvements as costs are consolidated, measured and used for future actions or improvements.

Since 2002, Hitachi Data Systems’ Storage Economics methods have provided a characterization framework of various hard and soft costs associated with storage system migration and data remastering. The determination of hard or soft, or the direct or indirect nature of migration costs, is a decision left to the organization. The following sections outline the nature of the cost of migrations without necessarily attempting to quantify each element; that is left to the reader to calculate based on local parameters. These sections identify cost estimates based on analysis of the survey results where participants were able to provide valid data.

Change Control/Remediation

Change control, version control and configuration management tasks are an integral function with all IT organizations. When a storage system reaches end of life or is fully depreciated (or its lease ends), there are certain fixed tasks associated with planning the move and removal of the old asset. Further, there are tasks associated with absorbing the new assets, including porting of existing software utilities and scripting to the new environment. Testing, certification and verifying the new assets aside, there are several operational tasks essential to replacing any hardware. Microcode, interoperability, security and operating system alignments are just some of the necessary tasks involved with change control and configuration management.

To the extent that the tools and methods used to complete the migration are limited (assets besides those being replaced or introduced), most migration methods consume roughly the same change control effort, time and cost. If the migration effort requires new appliances, tape resources or specialized software to enable the migration, then the added effort with these tools needs to be factored into the total migration cost.

In the Enterprise Storage Migration survey, survey participants identified how much effort was required to perform remediation activities and migrate existing scripting to the new storage environment. To enable organizations to use the results of the survey to develop their own rough cost models, the analysis of the costs from the survey are documented in costs per TB. This allows organization to tailor this data to their environment. Using the cost of labor that was identified by the survey participants and the cost of external consulting services, Hitachi Data Systems calculated the total cost for Remediation and Scripting to be $6,733 per TB. The breakdown of this total amount was: for Remediation — $4965 per TB and the amount for Scripting — $1767 per TB.

Server/SAN Outage Cost

In most migration methods or processes, the servers need to be rebooted to end the connection with the old storage system and to be re-established with the new resources. As a target to the server, the storage may present LUNs to many applications on the server, so a single server outage usually involves secondary impact to application outages (see below).

Each server reboot has to be carefully scheduled, usually in low usage times on the weekend, and often has to be planned well in advance of the actual outage. High availability server clusters or virtual machines require special time and attention to this detail. Some operations require a server reboot only during scheduled maintenance windows, which may or may not align with the storage migration timetable. New business demands are constantly shrinking or eliminating server outage and maintenance windows.

The cost or impact to the business of the server outage may be hard to calculate and is a function of the business revenue or operation costs of the servers' applications being offline. If scheduled during a routine maintenance window, these outage costs are minimized (planned and expected). But if they are outside of this window, the business impact can be measured in terms of lost revenue, opportunity loss or business disruption costs.

Ideally, new migration methods that reduce or eliminate server and SAN outages will enable continuous business operations. Eliminating server, SAN and application outages is a primary economic benefit of migration techniques that are done with continuous operations kept intact. Figure 4 provides a rough measure of the massive impact downtime can affect by industry and specific applications within those industries.

Figure 4. Downtime By Industry and Specific Application ($/hour)

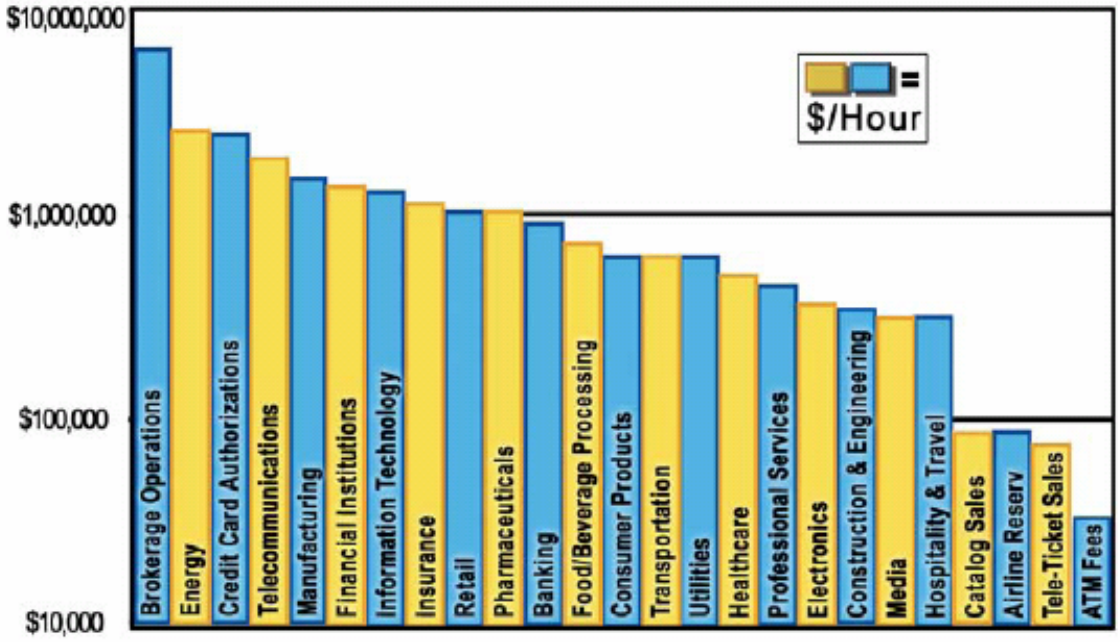

Source: Understanding Downtime, Business Continuity Solution Series, p. 4. Vision Solutions, 2006

Application Outage Cost

Just as in server outages, some applications may be impacted with a change in storage resources. Similar impact to applications can be drawn from the above description of server outages, in terms of:

- Scheduling within an approved maintenance window
- Impact to the business when the application is offline
  - Opportunity loss cost and revenue lost

Migration methods that can be achieved without disrupting servers, SANs, applications and databases are far superior to other methods because of the lightweight impact to business and IT operations.

Labor for Data Movement/Migration

Data migration requires human intervention. From planning and change control through the data movement and server/application restart process, there are several levels of staff labor and effort. This section highlights the often tedious methods of planning, copy, move and verification tasks that are required to be completed by the IT staff and/or external consulting resources.

Sometimes due to time constraints and other priorities, some migration tasks can be outsourced to third party agencies that will perform much of the heavy lifting involved with data migrations. Therefore, the costs of migrations needs to be viewed and tracked for internal and external consulting costs.

Costs of labor vary widely from region to region around the world. When calculating internal costs of migration labor, the administrator's salary can be used with addition of a 50 percent overhead (office space, insurance, tax) or burden factor to determine the true costs associated with the tasks. Nights, weekends and overtime rates are common for nonexempt personnel and need to be factored into the costs of migration.

Advanced methods of data migration (utilizing virtualization techniques) can reduce the labor time and effort drastically. The arduous tasks of copy and move are often automated to the extent that the administrators can assign the target devices and allow the system to move and migrate in the background. Overtime, weekends and late nights are eliminated.

In the survey, the participants identified the following costs for the labor effort necessary to conduct the migration. The participants identified both internal staff and external consulting costs. The internal IT staffing cost to conduct the migration was $2095 per TB and the external consulting cost was $3552 per TB.

Specialized Hardware and Software for Migration

Another factor of traditional migration methods is the introduction of specialized tools or appliances to complete the migration. These can be purchased, leased or used for one-off migrations. Larger organizations that are in a state of continuous migrations make long term commitments to these tools and appliances to meet the constant demand. Examples of these specialized tools include:

- Swing hardware (storage, virtualization appliances) that serve only to move data from system to system
- Backup systems (hardware and software) including tape media
- Copy and replication software
- Additional SAN or network devices (ports) to support the data movement process
- Dedicated network circuits

Survey participants identified that the cost for hardware and software to assist with the migration was $5,099 per TB.

Added Environmental Costs

A cost (hard cost) that is often overlooked is the additional floor space and power/cooling costs associated with storage system–based migration. When migration efforts take three to nine months to complete, there is a doubling of environmental costs associated with having both the old and new storage systems side by side for that time. Calculating the total environmental costs during a long migration period requires adding the power, cooling and floor space costs for the new storage system for this period. Data centers that are limited with power or space will experience secondary costs since long migrations will consume power and space resources that could otherwise be used for operational growth or strategic IT investment.

Useful Life Impact on Storage System

One of the soft costs associated with long migration times involves reducing the useful life of the storage asset. When looking at the time involved to get on the storage system and later get off the storage system, the actual useful live of the storage system is greatly diminished. It is not uncommon to see the asset's useful life negatively impacted by 20 percent to 33 percent due to the time and effort involved with migration. Reduction in the useful life has a direct impact on the return on asset (ROA) for the storage assets and the entire IT organization. Long migration times result in overlapping assets and costs to the organization that are not necessary, especially when techniques and methods exist now to nearly eliminate migration impact.

Figure 5. Assessing the Storage Asset's Useful Life

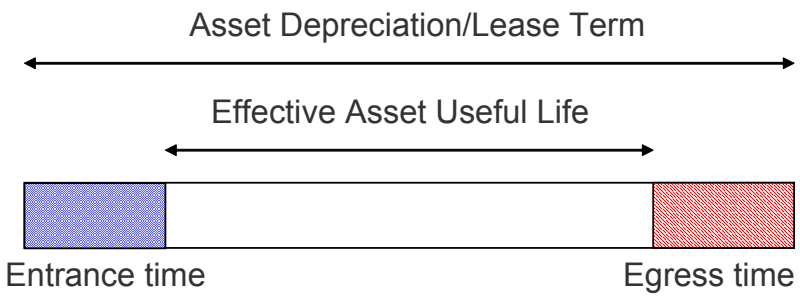

Using the survey data, we were able to calculate that the useful life impact costs for the storage system in the average data migration for the survey participants was $236 per TB. This calculation is introduced in the *Cost of Migration Comparison* section below and explained in *Appendix D — Cost Analysis*.

Storage System Maintenance

We often see that the migration process does not get started until the end (or near end) of the asset warranty period. Extending the terms of the hardware and software maintenance to cover the period of migration is a hard cost.

Combining People, Processes and Technology to Reduce Data Migration Cost and Risk

Data migration projects are complex projects that possess significant cost and risk. To successfully complete a data migration project, organizations must develop a comprehensive plan encompassing people, processes and technology. There are many methods to migrate data, each method having different levels of cost and risk

with varying advantages and disadvantages. It is important to identify which methods are optimal for your environment. To assist in the selection of the best methods, Table 1 describes and identifies the advantages and disadvantages of some of the basic data migration methods.

Table 1. Data Migration Options

Migration Method	Description	Advantages	Disadvantages
Tape Migration	Use a backup copy of data from the old systems and restore the backup to the new systems. The associated storage systems and/or servers must be offline to ensure that no new data is added.	• Can restore to a different host; very compatible to all operating systems. • Generally lower cost solution; all organizations have backup software.	• Major outage is required for the restore process and can be disruptive to or impact host performance at two separate times. • Process is very slow and does not scale very well when large amounts of tapes are necessary.
IP-based Host-based Migration Using Replication Technology (not array-based replication)	Involves replication of storage volumes. Volume management is used primarily to control disk resources by mapping the logical view of storage space with the actual physical disks.	• No impact on server performance. • Method optimal if leveraging existing investment and for small amounts of data.	• Significant investment is required. • Method does not scale to large amounts of data. • Software license fees could be required.
Host-based Migration — Block ***(There are a number of tools: VMware, Veritas Logical Volume Management)***	Host-based mirroring or replication solutions generally focus on file-by-file data movement to create a secondary data copy for disaster recovery purposes.	• Initial setup and cost is low; for a few migrations involving a few hosts this technique works well.	• Method affects server performance. • It can become difficult to manage in heterogeneous operating system environments. • Method becomes difficult as the amount and size of environment increases.
Host-based Migration Using Server Virtualization ***[There are a number of tools: VMware, IBM® System Storage® SAN Volume Controller (SVC)]***	VMware has created a tool set with VMotion to assist with server migrations that can assist with migrating the data as well. It is used for a storage area network (SAN) environment that is supported by VMWare and when migrating a VMWare virtualized environment.	• VMotion solution is very good for small source and target VMWare environments.	• Server and SAN must support specific software. • Performance impact on host; does not support remote migration. • Not fit for changing or consolidating operating systems. Focused only on x86 servers; currently only 33 percent of x86 servers run VMware and are compatible.
Host-based Migration — File Copy and NAS ***(There are a number of tools: Hitachi NAS Platform, powered by BlueArc®, using ScriptLogic Secure Copy, Hitachi Dynamic Replicator software, Arkivio AutoStore and Microsoft® Robocopy)***	There are a number of unique methods for specific host-based migrations that leverage NAS technology for file-based migrations either through NFS or through CIFS. To migrate CIFS or NFS file systems to or from an existing production file server NAS system to another, files are moved based on the network or file system protocols.	• Useful in moving from one operating system to another, or from one NAS platform to another platform; for example: NetApp to Hitachi NAS Platform. • Leveraging existing namespace enables transparent data migration and simplifies administration of target environment.	• Method is dependent on network, application and security infrastructure (authentication and permissions).

(Continued on next page)

Table 1. Data Migration Options *(continued)*

Migration Method	Description	Advantages	Disadvantages
File System-based Data Replication ***(There are a number of platform-specific tools: Inmage, NAS specific – Hitachi NAS Platform incremental data replication, NetApp SnapMirror and EMC Celerra Replicator)***	Data replication allows copying or relocating of both file data and file system metadata, depending on the type of server operating system, replication software, or NAS solution. Replication can include not only the data but also the policy, rules and schedules for the file system environment.	• Does not require third party tool. • Automation and scheduling are built into product. • Method is easy to use. • If you can leverage existing namespace infrastructure, you can simplify configuration of target environment.	• Vendor lock-in • Homogeneous support for platform, not heterogeneous platform support
Appliance-based Migration	Method uses hardware or software technology focused on conducting data migrations. There are a few different types of appliance-based approaches which vary based on the technology involved.		
	IBM TotalStorage SAN Volume Controller virtualizes storage.	• Method minimizes downtime and can scale very well.	• There is vendor lock-in.
	Storage area network (SAN) method uses a SAN switch – Cisco and Brocade. The SAN technology splits or mirrors the writes to both target storage systems until both are identical and stops writes to the source; and then the target storage environment is ready.	• Method supports heterogeneous storage.	• A highly scalable switch and large SAN are required, as well as a comprehensive suite of Cisco storage software. • Both storage systems must be online throughout the migration. • There is vendor lock-in.

Storage Virtualization Approach

The use of heterogeneous storage virtualization technology to conduct a data migration minimizes risk and cost. Using the heterogeneous storage virtualization approach, the new storage system is connected to the SAN along with the old storage system. Then, the new storage system discovers the connected hosts and is properly configured and tested. The data is copied to the new environment and data is transparently redirected to the new storage system. No added outage is required. Once all the data is on the new storage system, the old storage system can be decommissioned. This approach minimizes downtime and enables high data throughput. This approach is very flexible to allow for various outage windows and can scale to large amounts of data.

Table 2 provides a high level comparison of three of the major data migration techniques. For a customer environment with 100MB/sec Ethernet and average server data size of 650GB, we have found the following comparisons.

New Emerging Technologies

With the release of the Hitachi High Availability Manager software, organizations will able to completely eliminate the outage window. Host connectivity is configured at the switch level to two storage systems, so no rezoning is required. Since all volumes from the primary system can be replicated using Hitachi TrueCopy® Synchronous to a secondary system, host migration effort is greatly reduced; the migration execution is handled at the storage frame level, not the host level. The host discovery, review and planning phases of the migration are still required; in large enterprises it is possible to have a host connected to several storage frames. The Hitachi High Availability Manager enables a simplified and nondisruptive migration, which is 90 percent less effort and cost compared to the industry average.

For data migrations using host-based migration, the two outages for the host-based migration are for the discovery of the new storage system and the release of the old storage system after the migration has completed. These outages can vary depending on the operating system and the type of application. For example, old versions of Sun Solaris (Sun 5.8) would require a reboot after the data migration, resulting in an application outage. For new versions of Sun (Sun 5.10) the hosts can dynamically release the old storage system without an outage.

Table 2 illustrates that the Storage Virtualization approach has the least amount of impact on the environment; however, it requires a significant upfront investment in the storage virtualization technologies. If a customer is migrating to a virtualized storage environment, then this method has many significant pros with very limited cons.

Table 2. Comparison of Three Data Migration Approaches

	Storage Virtualization	***Replication Technologies***	***Host-based Migration***
Cross Platform	Y (any to any)	N	Y
Application Downtime/ Server	<30 min	2 to 3 hours	2 hours (2 outages)
Number of Outages	1	2 (second outage is managed/scheduled)	2
Impact on Server Performance	None	None	High
Technology Investment	High	High	Low

Cost of Migration Comparison

Based on the data that we obtained through the survey, we developed a cost model to represent the overall costs of data migration. This section of the white paper focuses on providing our findings to empower you to adapt this research to your own environment. As described in the cost section, there are additional cost elements to data migration, which were not captured as part of this research due to limits of Web survey capabilities and sample size.

At the highest level, this research indicates that the industry average migration cost for Fibre Channel SAN storage is more expensive than the $5,000 to $7,000 per TB commonly quoted by industry consultants, press or storage bloggers. Most of the posted or printed articles on this topic focused purely on the storage platform itself without considering interoperability requirements that storage vendors impose on customers at the SAN switch, host operating system or host bus adapter level. Additionally, the technology and tool costs required

to perform the migration itself are often excluded as is the effort required to rebuild the automation on the new storage platform with the use of scripts.

Figure 6 illustrates the different cost elements captured in this research.

Figure 6. Enterprise Storage Migration Effort/Costs per Terabytes Migrated

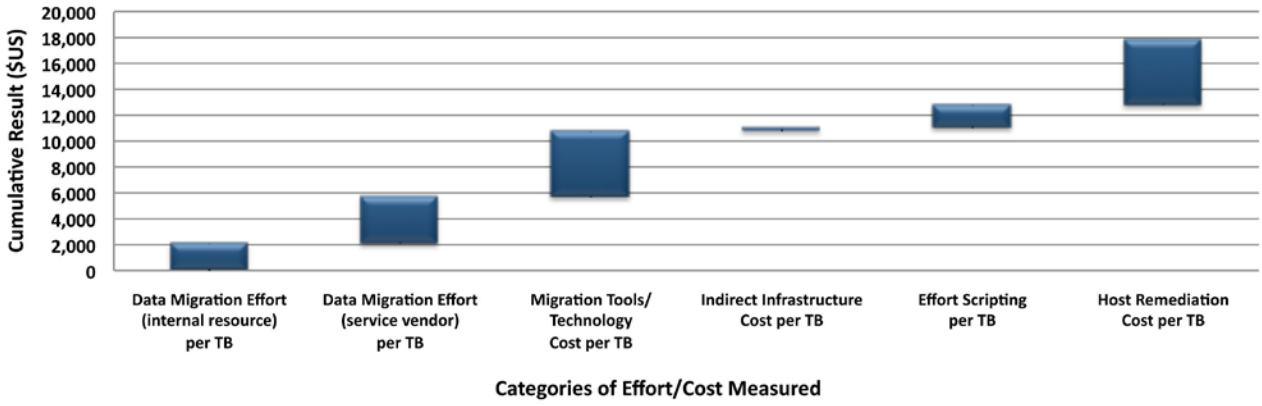

Source: Hitachi Data Systems/Tech Validate surveys, September 2009

The overall average cost of enterprise storage migration easily exceeds $15,000 per TB. The following paragraphs provide detailed explanations for each category of migration cost.

The first element that must be clarified is the “per TB” metric used to express the migration cost. This metric is not based on storage capacity but the size of the dataset migrated. This size is often estimated by using the used allocated storage capacity. For the purpose of this paper, the migration cost model is assuming the data size or the used allocated storage capacity is the same.

For the purpose of adapting this cost model, organizations can translate this cost to a per host basis. The cost of an enterprise storage migration is directly proportionate to the number of hosts impacted. (Note that other factors also directly impact the migration cost, such as multisite implementations, and time and day of maintenance windows.) With the average capacity per host of 1.25TB to 1.5TB, the reported migration cost would range from $10,000 to $12,000 per host.

For the purpose of this research, the survey focused on both internal resources and external contractor spending used to perform the migration. Average cost of internal resources required was constructed from the respondents’ answers to five questions. Table 3 provides the summary of the results from the survey participants on the key criteria.

Table 3. Key Survey Criteria

Full Time Equivalent (FTE) Cost	US$126,000
Size of Migration Team	Average reported internal resource team size was composed of six individuals.
Average Team Loading Assigned on Migration Project	Average reported internal resource team loading was 45 percent.
Project length	Average reported migration project length was 4 months.
Size of Data Migrated	Average of 1.25TB to 1.5TB of data per physical host (not virtual machine) was reported.

Using this data, we calculated the cost for internal resources was $2,095 per TB. For detailed explanation of calculations, see *Appendix D — Cost Analysis.*

Migration technique is a key driver of the overall project cost and length. For internal resources, respondents reported host migration requires four to six times more effort compared to virtualization, as shown in Table 4.

Table 4. Effort Required for Migration Techniques

Data Migration Effort (US$) per TB – Internal Resource Only	***Migration Technique***	
	Host-based Migration	**Storage Virtualization–based Migration**
Low	$1767	$297
High	$2787	$635

Enterprise organizations typically do not delegate host administration and maintenance to contractors for security reasons; therefore, the burden of the effort required for host migration falls mainly on the internal IT staff.

Similar findings were identified for host and virtualization migration technique impacts on project duration. Host migration projects were reported to be almost twice as lengthy as storage virtualization projects, as shown in Table 5.

Table 5. Migration Project Duration per Migration Technique

Migration Length	***Migration Technique***	
	Host-based Migration	**Storage Virtualization–based Migration**
Low	13 weeks	7 weeks
High	19 weeks	11 weeks

Please see the *Migration Approaches* section of this document for more information on advantages and disadvantages of migration techniques.

The respondents reported an average of US$3,552 per TB in service vendor cost. This value is computed from the service vendor spending divided by the average capacity migrated. The use of external contractors and consultants is a common practice in the industry to support off business hours migration. External consultants also help to reduce migration risk through experience and knowledge of migration technologies, tools and processes. The application operational requirements, migration techniques and tools, internal IT staff expertise and availability, maintenance windows and project size will greatly influence the number of contractors and consultants required. Average reported contractor and consultant personnel spending was $251,000 per migration project. For the detailed survey results see *Appendix C — Survey Results: Direct Costs*.

The respondents reported an average of $5,099 per TB in migration tools and technologies spending. This value is computed from the reported migration tools and technologies spending divided by the average capacity migrated. The cost category is highly variable based on the migration technique, vendor tools selected and licensing structure. See *Appendix C — Survey Results: Direct Costs* for the survey results regarding migration tools and technologies costs.

The fourth category of cost labeled “Indirect Infrastructure Cost” was modeled at $263 per TB. This indirect cost is defined as the excess infrastructure required while performing the migration and its associated costs. While performing the migration, target storage platform capacity must be equal to or greater than the capacity

of the source storage system. The same logic extends to port requirements, power, cooling, etc. To calculate the $236/TB, we used industry standard disk price per GB and estimated cost of infrastructure over required time frame. Replace this cost with the appropriate estimate that better reflects your financial liabilities associated with keeping the source storage frame(s) for an extended period while performing the migration. For detailed explanation of calculations, see *Appendix D — Cost Analysis.*

Respondents reported the fifth category of cost labeled “Effort Scripting” as an average of $1,767 per TB. This value is computed from the reported scripting effort in relation to the migration effort divided by the average capacity migrated. The migration effort baseline for this value was the sum of the internal resource effort and service vendor spending. There was high variability in the responses to this question, which could be due to the number of different platforms used by survey participants. Another interesting analysis from the survey respondents was that an average of 31 percent of the migration effort was reported to create new scripts on the target storage system(s).

Respondents reported an average of $4,965 per TB in host remediation effort with an average remediation effort of four days to support the most recent storage migration. This value is computed from effort reported multiplied by the FTE cost divided by the average capacity migrated. Four main drivers were identified to explain the wide range of answers reported:

- Length of maintenance windows due to application outage required
- The age of the respondents’ hosts and version of the operating systems running on these hosts
- Storage vendors’ interoperability requirements exclude older operating systems, host bus adapters and Fibre Channel SAN switches; an organization that did not keep their IT environment up to date may have to upgrade several hosts to support the target storage platform(s)
- Storage vendor interoperability

Lower Risk with Experience and Methodology Based on New Technologies

In our survey of customers’ experiences with data migrations, we have identified a number of critical success factors that are essential ensure successful data migrations. These critical success factors fall into the categories of experience and a number of best practices to ensure a successful migration.

Organizations that reported having greater than 10 percent budget or schedule overrun, unanimously identified lack of experience as their biggest reason for not meeting budget and schedule (see *Appendix E — Survey Results: Best Practices from Successful Data Migrations*). When conducting migrations, leverage a team that has extensive experience with different types of migrations and methods. Organizations should look at companies like Hitachi Data Systems, which has conducted thousands of data migrations; typical consultants have an average of 15 years of industry experience and more than 50 migrations completed. For more information on Hitachi Data Systems migration capabilities, see *Appendix F — Hitachi Data Systems Data Migration Methodology*.

Many of the best practices Hitachi Data Systems consultants have learned through years of experience are practices similar to organizations from the survey. Those who were successful identified the following key data migration best practices:

- Logical grouping of applications
- Splitting migration into multiple phases
- Staging/testing

Best Practices and Lessons Learned from Successful Data Migrations

Hitachi Data Systems has leveraged this methodology to the benefit of many of our customers. A few of these successes are described below. For each customer success, we have described the situation, solution and key migration challenges that Hitachi Data Systems addressed. As shown in Table 6, Hitachi Data Systems has experience with many kinds of environments and very complex requirements.

Table 6. Hitachi Data Systems Data Migration Success Stories

Situation	Solution	Key Migration Challenges
Large Telco Provider Migrating Large Data Warehouse in a Mission Critical Data Center		
• Migrate the virtualized storage to virtualized storage from one data center to another data center within very tight business timelines.	• Replicate virtualized storage to virtualized storage over 50km without impact on production. • Take a best practices approach to replicate and migrate data using block-based remote data migration project.	• Identify appropriate recovery time and recovery point objectives. • Avoid purchasing software not used in future environments. • Handle lack of detailed knowledge of SAN configuration.
Large Bank Migrates Open Systems and Mainframe Data from Old Storage Equipment		
• Migrate multiple EMC Symmetrix storage systems with 50 mainframe and open system servers and ~70TB of data to Hitachi Universal Storage Platform® V.	• Virtualize existing storage with Hitachi Universal Storage Platform V technology. • Migrate data to new storage system in the background, reducing application downtime and customer effort/ involvement.	• Running scripts to change zoning dramatically reduces impact per server. • Migrate/virtualize a large number of servers quickly.
Government Agency Migrates Virtual Server Environment to New NAS Systems		
• Allow host-based migration across the customer's IP network. • Migrate Microsoft® Windows® Storage Server cluster that hosts 28 virtual servers on an HP EVA SAN to Hitachi Universal Storage Platform V.	• Use of Domain Name System (DNS) will enable remapping clients' drives to servers matching those created on Hitachi NAS Platform, powered by BlueArc®, as Microsoft Exchange Virtual Servers. • Microsoft's Robocopy was used to move the source data files and access control lists between the source servers and the target Hitachi NAS Platform. Shares also were copied in the same fashion using RichCopy.	• The use of DNS provided a means for the organization to remap clients drives to a core set of servers matching those created on Hitachi NAS Platform.
Technology Manufacturer Migrates NetApp NAS System to Hitachi NAS Platform		
• Allow host-based migration across the customer's IP network. • Migrate NetApp NAS system with file systems ranging from 45GB to 600GB in size, with a total of just over 60TB.	• Static copy of data was created for initial population on the Hitachi NAS Platform. • Various tools were used to establish a baseline data set on the Hitachi NAS Platform, including: — Hitachi NAS Platform–to–Hitachi NAS Platform Migration Tools: Data Migrator — Windows Server–to–Hitachi NAS Platform Migration Tools: Robocopy, Secure Copy — UNIX Server–to–Hitachi NAS Platform Migration Tools: rsync, Veritas Volume Manager	• Minimize investment in tools. • Required use of multiple tools to address migration requirements.

Conclusion: Develop Comprehensive Strategy to Reduce Both Cost and Risk

Data migration is an ongoing activity at all enterprise IT data centers. Data migration projects consume considerable resources and involve considerable risk. There are key best practices and technologies that organizations can use to lower cost and risk of data migrations. Information contained in this paper regarding the costs of enterprise data migration provides cost modeling information for IT management to evaluate different data migration options. Based on this research and industry data, IT management should evaluate new storage platform vendors in conjunction with migration solutions to lower data center operating expenses.

The cost modeling data and experience Hitachi Data Systems has with data migrations point to data migration solutions that can help organizations reduce risk and costs for their data migrations now and for the future. Choosing Hitachi innovative storage architectures with market leading storage virtualization capabilities and heterogeneous vendor support coupled with Hitachi Data Systems experience and best practices has helped organizations achieve lower costs and less risk for data migrations. Hitachi Data Systems Global Solutions Services (GSS) organization has highly trained, knowledgeable and experienced data migration consultants who have migrated thousands of enterprise systems for our customers. Over the years, GSS has developed heterogeneous storage migration consulting expertise, as well as methodology and best practices that reduce migration risk.

When planning your next data migration, consider leveraging Hitachi Data Systems experience, knowledge and best practices to help you conduct your data migration cost effectively with low risk. By reducing risk and costs for data migrations, you will reduce operating costs due to data migrations and more easily prevent technological obsolescence. This allows you to be more adaptive to change, leveraging existing investments by repurposing existing assets or facilitating technology refreshes to reduce costs.

Appendix A — Survey Demographics

The Hitachi Data Systems research goal was to raise market awareness on the topic of enterprise Fibre Channel SAN storage migration and provide ways by which IT executives could tailor the findings of this research for their own environments. Using TechValidate Web surveys platform capabilities as the primary data collection outreach, the scope of this research was reduced to accommodate respondents' knowledge gaps and the limited time storage professionals can spend answering questions. With this in mind, we excluded several related and soft costs of storage migration from this project to ensure the survey would stay within these limits. The following six figures provide graphic views of survey demographics.

Appendix A — Figure 1 compiles responses regarding the number of source frames involved in the migration project. Enterprise frames are defined as those similar to EMC Symmetrix, Hitachi Universal Storage Platform®, HP Storageworks XP, IBM® DS8000® series, Sun StorageTek 99XX — systems usually available in a greater than two controllers configuration. Midrange/entry level frames include other storage systems that are usually available in a one or two controller configuration.

Appendix A — Figure 1. Category of Storage Frame Migrated

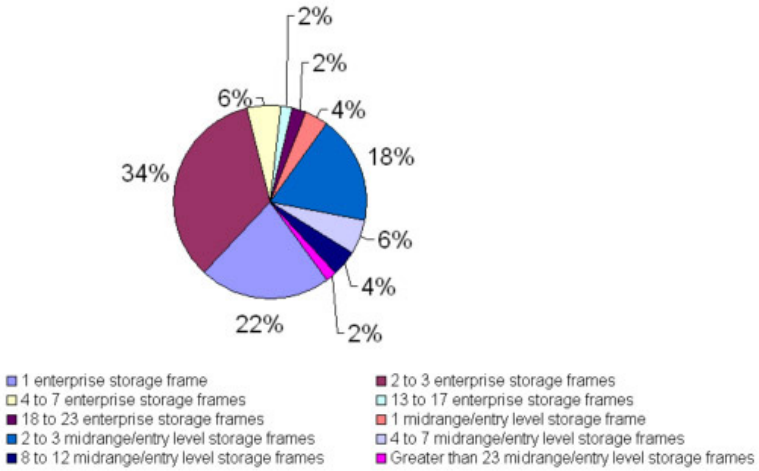

Appendix A — Figure 2 compiles responses indicating the amount of data that was migrated from the source storage system(s) to the target storage system(s) in respondents' migration projects, indicating *utilized* storage capacity.

Appendix A — Figure 2. Respondents' Most Recently Completed Fibre Channel SAN Storage Migration Project

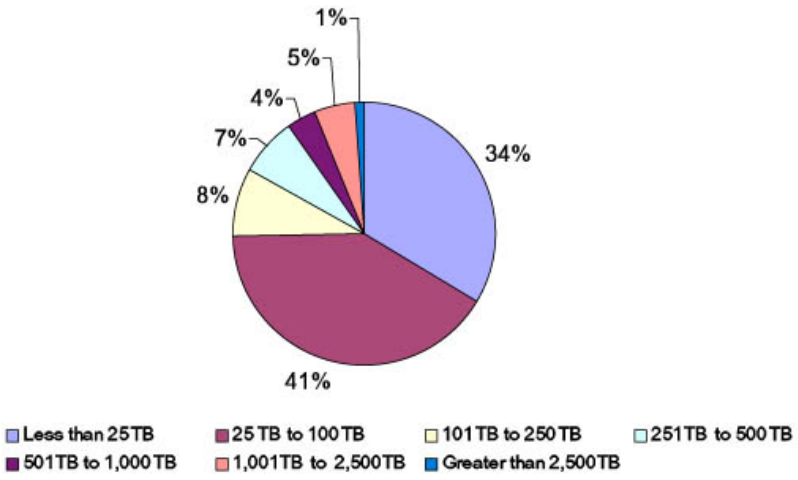

Appendix A — Figure 3 surveys the number of respondents returning surveys from various industries.

Appendix A — Figure 3. Respondents by Industry Organization

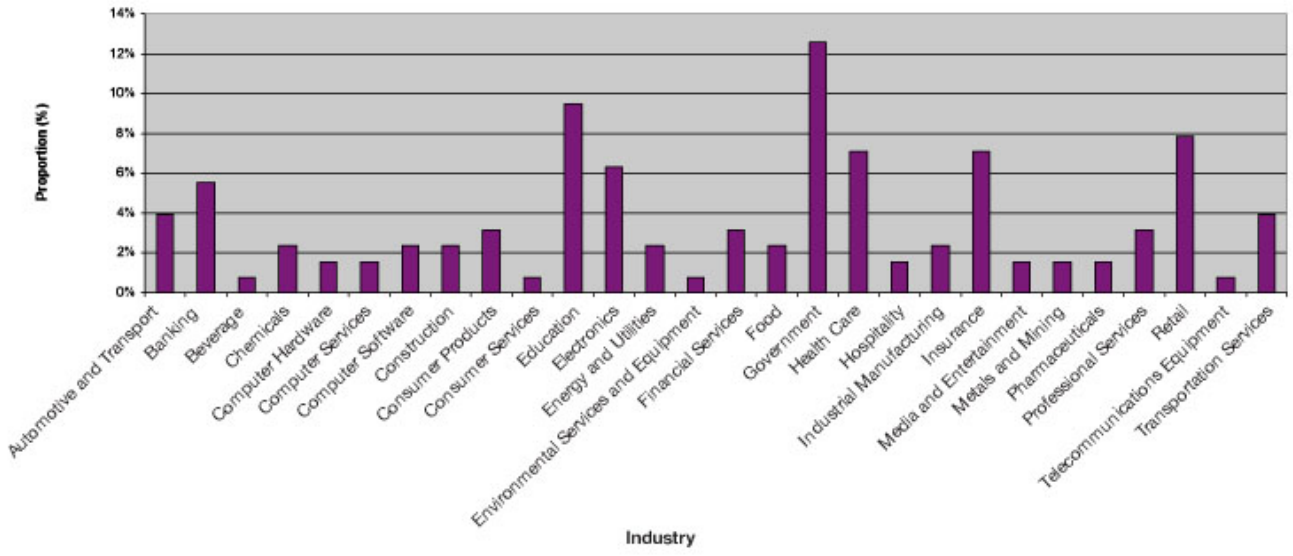

Appendix A — Figure 4 presents the number of hosts the respondents used in their migration projects.

Appendix A — Figure 4. Number of Hosts Impacted by Migration

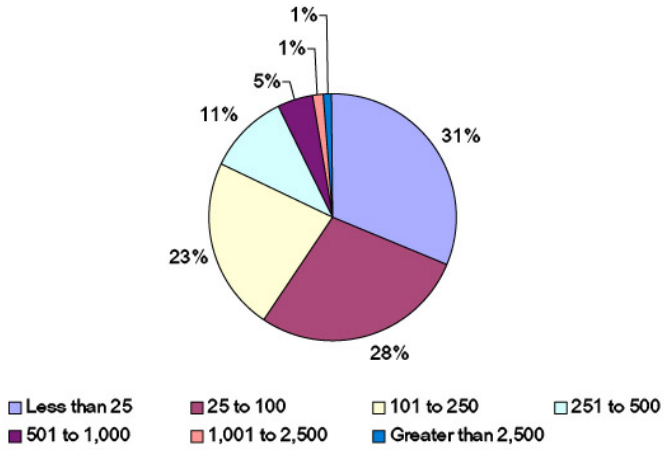

Appendix B — Calculation of Data Migration

Appendix B — Figures 1 through 4 summarize the data collected on internal resource used to support the most recent enterprise Fibre Channel SAN storage migration of the respondent.

The respondents' fully loaded cost or the annual cost for a full time equivalent IT storage employee in US dollars is represented in Appendix B — Figure 1.

Appendix B — Figure 1. FTE Cost

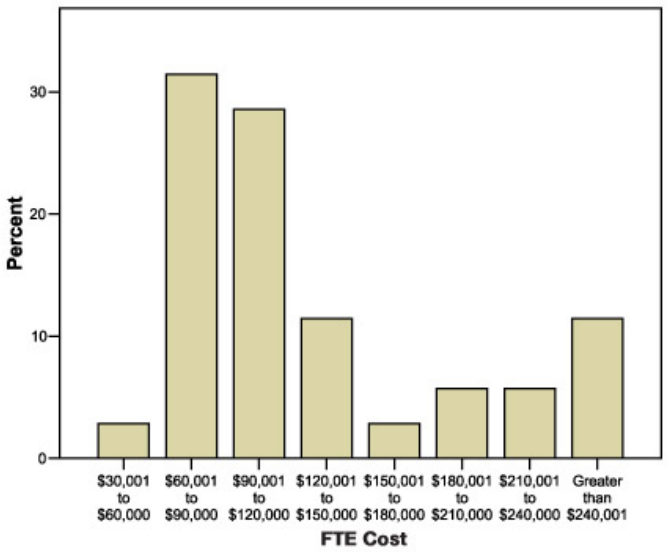

Appendix B — Figure 2 depicts the number of internal staff allocated to respondents' migrations.

Appendix B — Figure 2. Migration Team Size

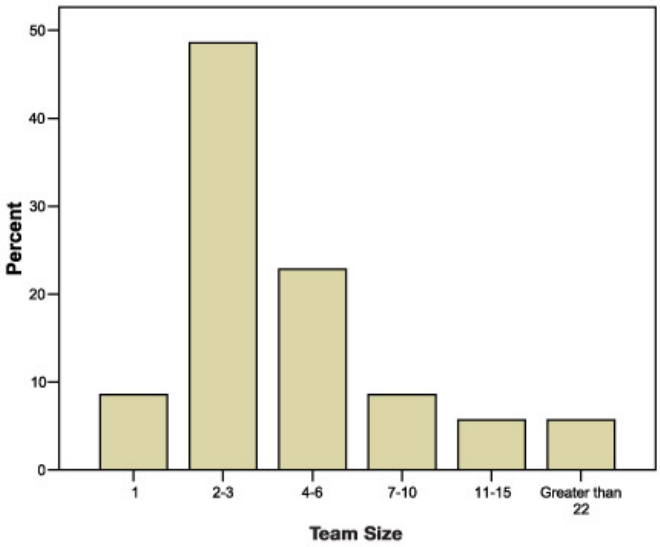

Appendix B — Figure 3 presents the average amount of time each of the respondents' assigned resources spent on this project.

Appendix B — Figure 3. Team Loading

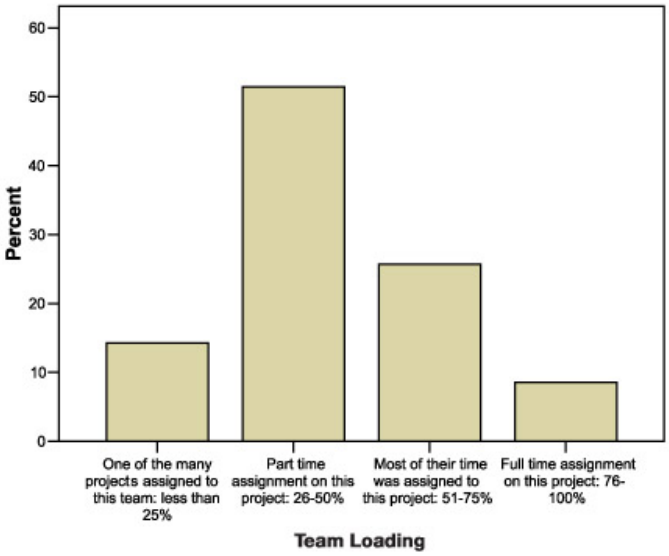

Appendix B — Figure 4 illustrates the time frames of respondents' projects.

Appendix B — Figure 4. Migration Project Length

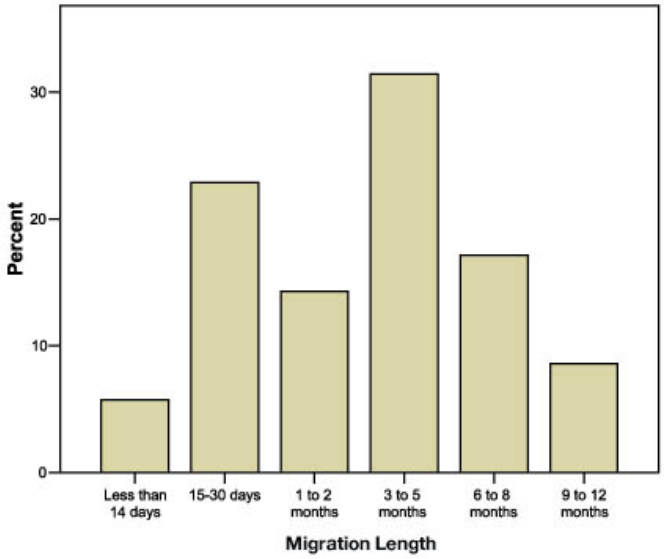

Appendix C — Survey Results: Direct Costs

Appendix C — Figures 1 through 4 illustrate the direct costs respondents incurred as a result of their data migration projects.

Appendix C — Figure 1 depicts the amounts respondents spent on service vendor personnel in US dollars during their migration projects.

Appendix C — Figure 1. Service Vendor Personnel Costs

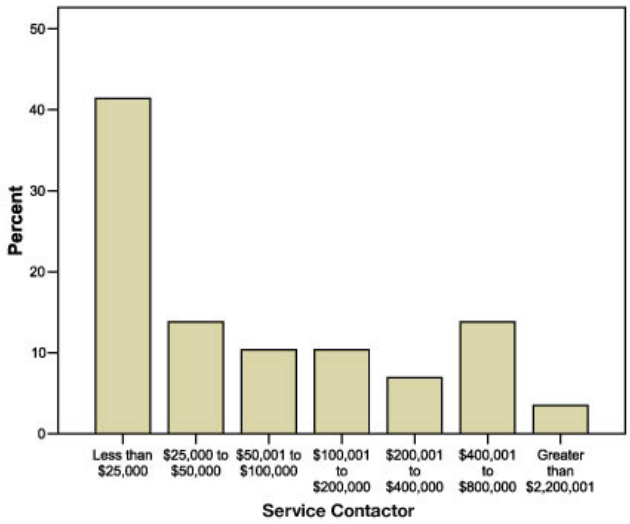

Appendix C — Figure 2 reveals the amount of time respondents spent creating new scripts to duplicate the source of storage system(s) administrative tasks (for example, backup, in-system replication, remote replication, etc.).

Appendix C — Figure 2. Effort Scripting

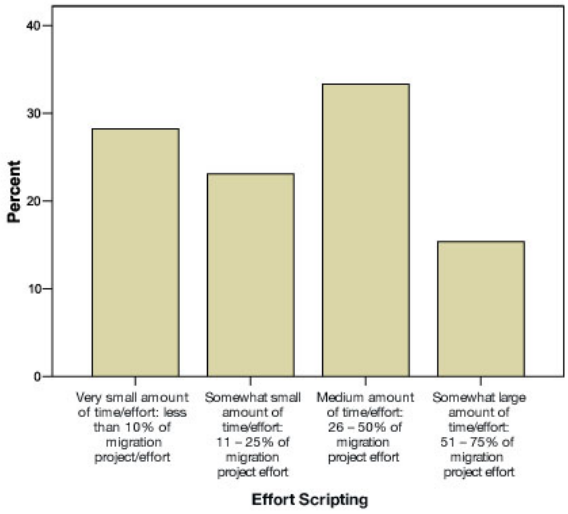

Appendix C — Figure 3 depicts the average amount of time and effort spent on each respondent's affected hosts on remediation tasks, such as host replacement, operating system, patches, driver installation, host bus adapter replacement, etc. This measurement excludes discovery phase, review and design, implementation and testing phase efforts.

Appendix C — Figure 3. Server Remediation Effort

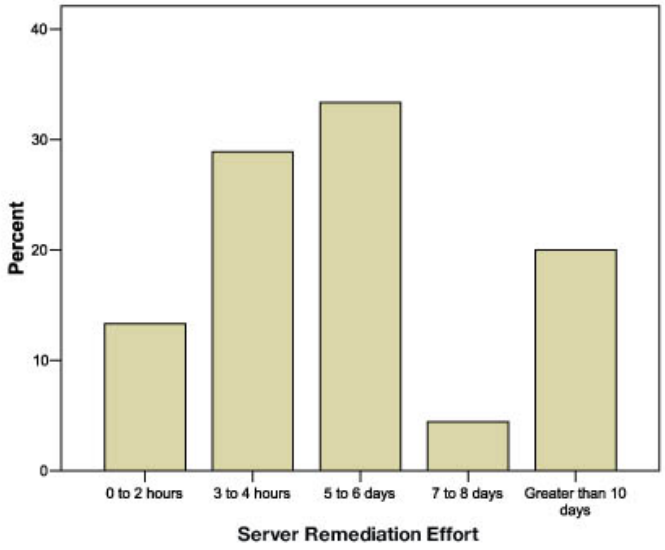

Appendix C — Figure 4 illustrates the respondents' total costs for tools to support their migration projects (host migration software, discovery software, replication appliances, etc.).

Appendix C — Figure 4. Migration Tools and Technologies

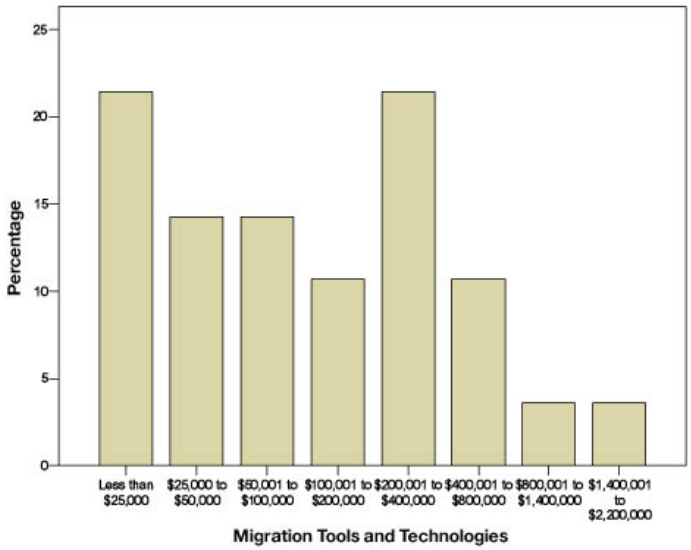

Appendix D — Cost Analysis

Internal Labor Cost — $2095/TB

"((FTE value * Team size * Team loading/100 * Project length * 13/12)/48 weeks of work per year)/Average size of data migrated"

The 13/12 multiplier is a normalization factor to compensate for the project length dimension, which was coded as one month equals four weeks in the survey database. The use of 48 working weeks per year default value is using two weeks of vacation time and 10 days of official holidays. This is representative of United States labor practices. No overtime pay factor was applied to this cost due to lack of knowledge on respondent organizations' labor practices, even if most storage migrations are executed outside normal business hours due to application downtime requirement. The direct internal resources cost will vary greatly based on actual location of the organization, labor pay practices and size of project.

Indirect Costs Calculation — $236/TB

The limited number of survey questions prevented the research team from capturing the specific case of hard and soft costs. Hitachi Data Systems field delivery consultants provided several examples on this topic where the customer had to purchase additional Fibre Channel switches to extend the leasing agreement or pay additional maintenance fee on source storage system(s) to accommodate migration project length. Understanding that the source system(s) may have been purchased three to five years ago, it is challenging to approximate street price at the original time of acquisition to derive additional leasing or maintenance fees. The research team chose instead to use 2009 street price estimates for modular and enterprise storage platforms to compute the excess monthly capacity costs, excluding extra Fibre Channel switches, ports and soft costs like power and cooling. The excess capacity acquisition value was estimated using a modular bit price of $2 per GB and an enterprise storage bit price of $7 per GB; these figures include hardware, software and maintenance. The reported size of data migrated was prorated over the reported number of modular and enterprise source frame(s) to estimate its acquisition value. The estimated acquisition value was then multiplied by a monthly cost factor of $3 per $100 of capacity value and the project length (expressed in months). It is similar to computing a monthly lease payment for the migrated capacity using 2009 street pricing. Due to aggressive storage bit price erosion in excess of 30 percent per year, it is the belief of the research team that we are underestimating this cost by a wide margin.

Appendix E — Survey Results: Best Practices from Successful Data Migrations

Respondents who have successfully completed data migration within 10 percent of budget identified these results.

Appendix E — Figure 1. Successes with Minimal Budget Overruns

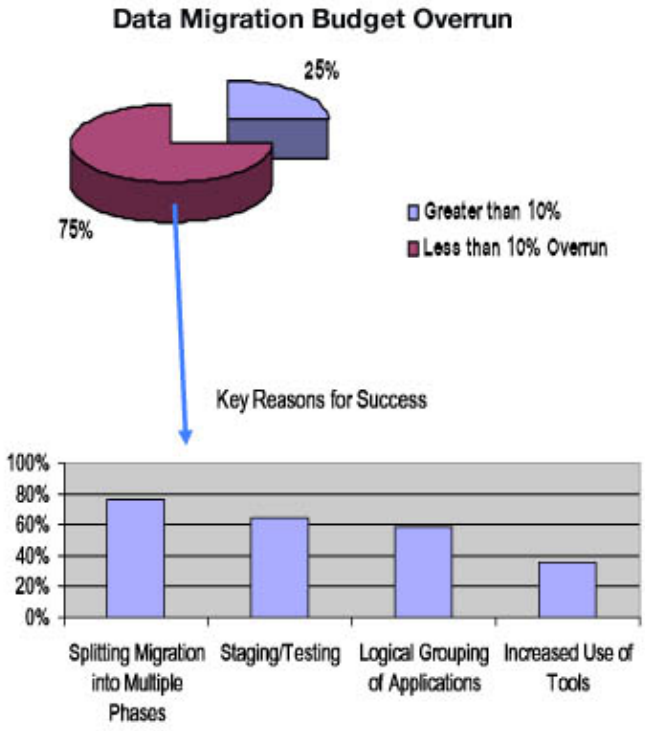

Appendix F — Hitachi Data Systems Data Migration Methodology

These migrations on average involve 40 servers and 10TB, with the largest data migration projects handling over 5000 servers and more than 2PB. We can perform data migrations using all forms of technologies and methods described earlier in this paper: host-based, replication, virtualized and using appliances. Our methodology, shown below, has been proven and tested, can be used for file and block-based migrations and is flexible to accommodate multiple migration technologies and methods. Also, the Hitachi Data Systems migration consulting practice is not focused only on Hitachi storage systems; we leverage many third party products and know how to migrate to and from other storage vendors' products.

Our methodology starts with developing a comprehensive project plan with the organization before the actual data migration. The Hitachi Data Systems migration methodology leverages the best technique or approach based on the customer requirements and existing or planned new infrastructure. Hardware-based solutions from third party providers (Acopia, Arkivio, etc.) are sometimes utilized, depending on previous licenses or upcoming license purchases. Hitachi Data Systems methodology includes the key best practices that were identified in the survey in addition to the following:

- Implementation: through Hitachi virtualization technology since 2004
- Discovery: leverage Hitachi tools
- Minimized disruption: no changes made to applications or data, LUN configuration, etc., on external storage (other than LUN masking)
- Improved commitment: mandatory stakeholder meeting
- Improved change control: Change Advisory Board (CAB) migration events reviewed through our internal CAB for approval
- Audit trail: migration workbook provides extensive documentation

The methodology follows the framework identified in Appendix F — Figure 1.

Appendix F — Figure 1: Phases, Purposes and Results of Hitachi Data Systems Data Migration Methodology

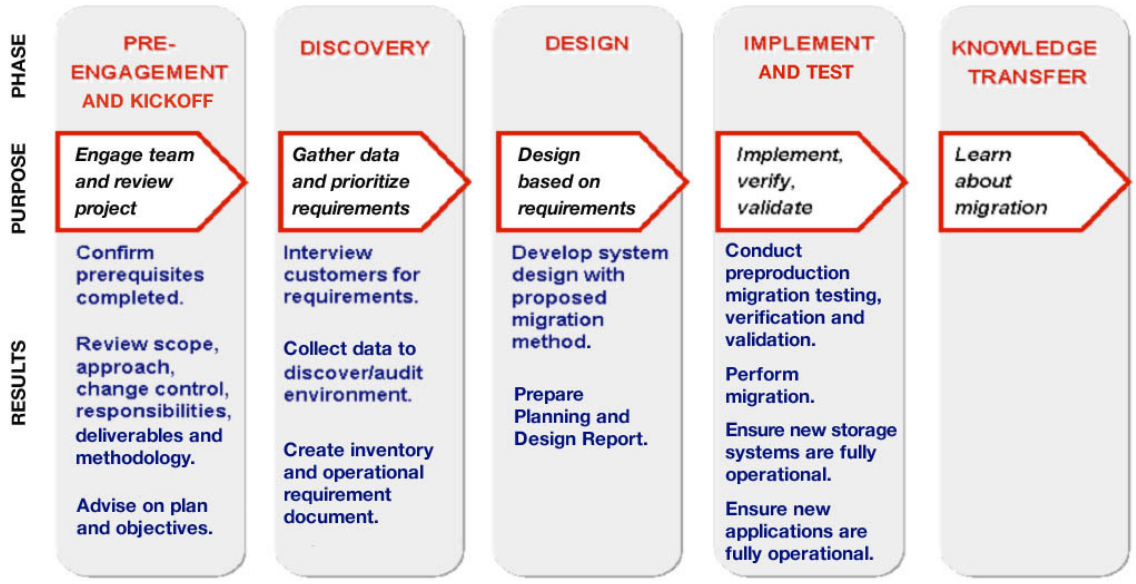

Hitachi Data Systems Corporation

Corporate Headquarters 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Asia Pacific and Americas 750 Central Expressway, Santa Clara, California 95050-2627 USA
Contact Information: + 1 408 970 1000 www.hds.com / info@hds.com

Europe Headquarters Sefton Park, Stoke Poges, Buckinghamshire SL2 4HD United Kingdom
Contact Information: + 44 (0) 1753 618000 www.hds.com / info.emea@hds.com