WHITE PAPER

Data Migration with SGI

Business Value *as well as* an IT Solution

By Mark Peters, Analyst

March, 2008

Table of Contents

Introduction

To migrate is to move. In the animal world, or with mankind, it is—hopefully—about getting to a better place. Sometimes, the motivation to leave the original situation can be due to onerous or unpleasant factors; other times, it is simply the possibility to improve a situation that leads to relocation. Whatever the motivation, the degree of ease or difficulty of the journey can itself be a considerable factor.

Storage migration in IT is essentially the same—it's about getting data to a better place. And the relative difficulty of the relocation is a significant contributory factor when IT managers decide whether or not to make a switch. The need to move and the attraction of the new location may be clear, but the historically high degree of difficulty in making the transition has been a major thorn in IT's side for decades.

Whenever such movement—migration—is hard to do, the motivation, or pay-off, must be even higher in order to make it attractive. Animals only move because food or water supplies run short and make it imperative; we humans don't typically move houses or regions more than is necessary because of the sheer hassle of doing so. And so it is with data—migration can be a hassle… but what if it weren't? Perhaps a lease has run out or maintenance is being discontinued. In those cases, the issue is that 'something *must* be done. Historically, migration is simply an unpleasant operational issue to be put off if at all possible, then eventually dealt with and fixed. However, if the act of migration becomes easier, then it is possible to view the act as something that *can* be done more frequently and for incremental gain. Migration can become an activity that is associated with opportunity rather than pure necessity. In short, simplified migration can deliver business value.

Improvements in migration tools and methods can both remove obstacles to what we'll call 'required' migrations *and also* expose opportunities to take advantage of 'choice' migrations. SGI recently entered this market with a compelling offering for such an enhanced migration facility. Although SGI is a company that is mainly known in the High Performance Computing (HPC) space, it has (quietly!) been building decent capability in other areas—including storage. This new migration offering is its latest value-add play in the storage market.

Data Migration

A quick glance at the state of the data migration market shows a lot of opportunity, simply because most users seem to be continuing with the traditional ways of doing things—in which migration is a little more than a necessary IT evil.

Storage Lifecycle Management Strategy – the Why and the Value

Why would a user want a strategy to manage the lifecycle of their storage? While it's easy to think of the underlying *equipment* reasons, the real answer lies in terms of an organization's *data*. The value of a Storage Lifecycle Management Strategy is to get more life from that organization's information and to squeeze the best information bang from each storage buck.

To do this, data must be moved off, and between, devices. Potential motivators include:

- Technology refresh (whether host or system upgrades, and whether driven by operational or financial issues, such as the equipment simply being out of date, or because a lease is ending)
- New applications, data center or relocation
- Consolidation (which could also be driven by many factors—finances, space limitations, server virtualization, etc.)
- Processing needs, such as a need for a test environment or to add tiering capabilities
- A need to test a new product and/or supplier
- Straightforward data growth or performance balancing

- Asset management—perhaps a better deal is to be had, or perhaps an installed product is reaching its 'end-of-service' (when migration is often 'encouraged' since many vendors make crossover licensing hard or even impossible)

Conversely, the lack of a migration strategy can slow or stop progress. This could preclude a user taking advantage of advances in technology, which in turn could negatively impact its ability, and costs, to do and expand its business. For sure, avoiding migrations can be attractive—after all, it is associated with disruption, uncertainty, point products and proprietary tie-ins—but the benefits of such avoidance are transitory and likely to come only at a higher future cost in terms of operational pain, money and potential business impact. Users understand the 'better place' to which they aspire—the *real* issue for many is not about having a strategy, but actually being able to do anything about it.

It is imperative that organizations start to embrace and plan for more effective data migration using a tool that can both mitigate or preclude the pain of 'required' migrations (fixing the IT operational problem) and also make available the possibilities of 'choice' migrations (providing added business value).

The State of Migration

The general frustration with traditional approaches to migration has already been mentioned—but what are the specifics? First and foremost, most migrations take applications down. In addition, most migrations are: 1) Uncertain until the full job is finished, 2) Homogenous and inflexible and 3) Consume valuable IT resources (both systems and staff). A review of traditional data migration shows that none of the choices are especially appealing:

1. **Backup –** in which data is backed up from one system and then restored on another. The execution is akin to Disaster Recovery, but of course, the motivation is different. Although a cynic might say the reference to the word 'disaster' is apropos, the fact remains that this migration method is still widely used. It takes applications down, is typically slow, consumes significant server resources and a user doesn't really know whether it succeeded or failed until the job is complete. With frequent delays, it is not at all unusual for a user to get only some percentage of the way through (and whether 1% or 99% makes no odds), run out of operational window and then have to start all over again. Just about the only good news is that back up operations often have their own network and so network clogging can be avoided.

2. **Copy to Disk –** similar to regular back-up in most respects, but faster. Applications are still down and data is still in a state of flux.

3. **Network Copy –** whereby, using a practical example that we've probably all experienced, data on a Home drive is copied to another drive somewhere on the overall network before being restored to a newer/bigger/faster Home drive. Once again, many of the issues are the same, *plus* this can clog the network, which may be an issue depending on your business profile. Users are also often nervous about their data being 'off on a cloud drive' somewhere.

4. **Snapshots –** this more recent option is possible with virtualization (whether within the storage system itself or via host/gateway based options). It is definitely a superior method, especially as it can significantly reduce the application downtime. Since two servers have a point-in-time view of the data, one snap can be used to effect the migration over time before switching across to the new storage system. There will invariably still be some final changes to do using some other method, but clearly, uncertainty is much reduced.

Challenges and Inhibitors/ Needs and Opportunities

There is a range of practical challenges and inhibitors around migration, which is why users fight against doing them if they can. And yet—ironically—such reticence can itself contribute to the need for migrations! The issues can be grouped into two separate sources:

Business Issues

- Operational concerns – these include both the risk from exposure to data loss or corruption plus strain on personnel. Most migrations are done at nights and on weekends; and this is not a small thing—large data centers with, for example, multiple lease terminations can easily end up faced with migration needs multiple times per quarter.
- Financial pressure to get the migration done is invariably at odds with the business pressure to preclude application downtime, with the consequent risk of business and/or customer loss.
- In general, it has not been in the interests of vendors to 'play as nice' as they might. Almost all 'talk the talk' but the walk tends to be slow as vendors seek to protect their proprietary interests. The resulting lack of supported heterogeneity adds a concern to users that do not want increased risk exposure.
- Personnel 'migration hangover' effect – tired staff are less productive on the Monday after a weekend migration effort….*and* will rack up lots of costs in terms of overtime payment and accrued time-off that has ongoing productivity implications

Data Issues

- The risks inherent in migrating data from platform to platform (loss, corruption, lack of availability) are extended by the risks inherent in keeping data on aging storage (is support available? device speed and reliability may be insufficient; increasing costs are likely).
- Data in flux - uncertainty as to what is where and whether a migration worked or not, and/or no in-process knowledge of whether or not the migration is proceeding successfully.
- Managing multiple migration tools from multiple vendors is likely—and potentially confusing—in a typical mixed vendor environment.

With such challenges, it is no wonder that migration is seen as annoying, risky and—often—to be avoided or delayed if at all possible. Because none of the traditional choices are very good, and because of the significant challenges and inhibitors it is clear that there are implicit needs and opportunities for a better approach. The benefits of moving migration to be *online and in the storage network* have been realized for a long time. First and foremost, of course, would be to get migration done online with *applications still running.*

There have been various attempts at new migration methods—pure software plays are somewhat popular, but they can bog down servers; and, although the applications do stay up, performance often gets adversely impacted and so these migrations still typically revert to nights and weekends when the business impact is likely less. Recently, with the logic of network based, heterogeneous migration becoming inescapable, mixed hardware/software solutions are coming to market. All such tools employ some level of virtualization, which permits the copying to be transparent, fast and non-disruptive. This is where the new offering from SGI fits.

Why is Application Uptime so Crucial?

ESG surveyed 128 enterprise customers and the majority could not tolerate more than 1-4 hours of downtime before it would have a significant impact on their businesses (23%). Nearly as many said that they could not tolerate more than an hour (22%) and many couldn't tolerate any downtime at all (14%)

Research showed that large enterprises can tolerate very little, if any, downtime. However, ESG also found that 46% of mainframe administrators perform data migrations offline. Additionally, on average, 54% of companies surveyed take between 1-6 hours to perform data migrations. Simple deductive math says that there are a number of companies performing data migrations offline, which is negatively impacting business and revenue.

The Business Value of Effective Data Migration

A Means, Not an End

It's worth remembering that no IT manager wakes up deciding they want a 'fun day of data migration.' Even if migration were fun (which it hasn't been), it is not done for its own sake. It is a means to a number of IT and business ends—such as asset management and tiering. Although often viewed merely as a 'fix' to an IT 'problem,' best practice data migration implementations can deliver genuine business value. Thus, when evaluating migration solutions, users should not look on a point-product basis, but across a scope of requirements covering such aspects as the breadth of a solution, the depth of knowledge from the vendor and the availability of support.

The 'ends' that migration supports are themselves constantly evolving—what a good migration tool provides is the storage flexibility needed to support new business approaches and endeavors. ESG research[1] shows the impact of some key initiatives in many data centers. ESG's global survey of IT management software and services buyers shows that the convergence of technologies such as virtualization, SOA (service-oriented architecture) and Web 2.0—combined with rising energy prices and increasingly tight IT budgets—is forcing IT decision makers to rethink their data center operations strategies. Specifically, data from 602 global decision makers shows that:

- 76% expect virtualization to extensively or moderately impact their IT management requirements in the next 24 months
- 68% expect Web 2.0 to extensively or moderately impact their IT management requirements in the next 24 months
- 63% expect SOA to extensively or moderately impact their IT management requirements in the next 24 months

The connection to migration is simple since all of the aspects above have significant storage implications and demands. Such changes could be to platforms, technologies or data placement as well as—almost inevitably—an increasing focus on lower (or no) application downtime. And, of course, all these things apply both as these new technologies are implemented and as they are then delivered.

The bottom line? A sophisticated data migration ability—preferably offering heterogeneous support—is a crucial and strategic part of an organization's ability to deliver new applications and to grow revenue.

Business and Cost Advantages

In most organizations, the physical life and applicability of the storage asset is *less* than the useful life of the data. Maximizing the life and value of data assets economically is the driving business motivator behind migration. And the ease of the migration itself is crucial. Think, for example, of a personal music (a.k.a. data) collection—the willingness to adopt iPods and other MP3 devices has been much greater than the LP to CD move a couple of decades back, largely thanks to the ease of data migration (a.k.a. 'ripping' and 'synching').

The Virtualized Storage Migration Solution from SGI is designed to virtualize data off the control/data path (see the 'How it Works' section) and to support storage capacity and performance pooling across an entire SAN storage domain. This delivers a range of pertinent business benefits for migration needs:

- Technology that can span (virtualize) storage devices from multiple vendors and that can provision multiple host systems; such heterogeneity precludes supplier lock- in and allows storage system choices to be based on merit, thus ensuring users are always able to take advantage of the latest improvements in cost and function.

[1] ESG Research Report, *2008 Service & Infrastructure Management Survey*, March 2008

- Centralized management for distributed enterprise storage, with a single common toolset across heterogeneous storage platforms that spans applications across all verticals - finance, manufacturing, telecommunications, entertainment, utilities and transport.
- Additional data management capabilities based on the virtualization functionality include -
  - Mirroring and replication (for business continuity, disaster recovery…even of unlike devices)
  - Snapshot and point-in-time copies (for fast disk-to-disk and disk-to-disk-to-tape backup and restore solutions)
  - Multi-directional migrations (for tiered storage and ILM solutions)

The amounts of money can be very significant: One large end-user in the telecommunications industry tells of spending over $1 million on unbudgeted maintenance costs simply because: a) it did not have a non-disruptive migration tool and b) it could not find a sufficient window to take down a specific application to allow the equipment refresh to take place. Another user that ESG spoke to last year spoke of saving a similar amount of money in capital equipment over 18 months by implementing just a two tiered storage migration approach.

Storage Resources

Storage resources are valuable and should be treated like currency. The Virtualized Storage Migration Solution can be used to 'spend' storage resources wisely. Rich, revenue producing mission-critical applications can be assigned the most expensive, highest performing storage devices, RAID levels and copy services. Frugality can be exercised on lower priority applications and persistent data by using slower drives with dense capacity. The solution makes it possible to move freely between price/performance optimized cost structures while applications remain on-line and accessible.

'Green' Credentials

The 'green' aspect in IT generally, and storage specifically, is certainly high on many data center managers' agendas right now. It certainly does no harm to the overall Virtualized Storage Migration Solution story to find that it utilizes energy efficient Dual-Core Intel Xeon Processors. Aside from its own economical power footprint, the environmental impact of the solution is all about what it can achieve in users' data storage facilities and the power, space and cooling demands therein. All the hype aside, it is a simple fact that optimizing both the use of the storage in any data center and also the placement of different types and priorities of data onto different storage systems is one of the quickest ways to save on the usage of power, cooling and space. However, the best will in the world won't achieve this—an effective migration policy and facility can.

Of course, 'green' is as much an economic issue as it is an environmental one. ESG research shows many examples of why a more considered storage approach is crucial across organizations. Looking at Figure 1, the massive growth in digital archive capacity is shown. It is simply unnecessary, uneconomic and environmentally irresponsible to store all this data—which, after all, is by definition, persistent—on the fastest, most expensive, least environmentally conscious, disk drives. Users might not want to take down their newest applications, it definitely makes all kinds of sense to automatically migrate to more optimized storage tiers.

Indeed, further ESG research (*2007 Data Protection Survey*, published in *Data Protection Market Trends*, January 2008) finds categorical evidence that organizations are looking for secondary storage systems that are both space- and energy- efficient more than ever. For example, about a third of respondents with storage capacities above 25 TB said that the physical footprint (34%) and the energy efficiency (32%) had become more important considerations in their purchasing process for secondary storage systems. And, of course, somehow the data has to be moved to the 'greener' platforms. The succinct summary here is that migration is a key element in any workable 'green' storage strategy.

FIGURE 1. WORLDWIDE ARCHIVE CAPACITY GROWTH

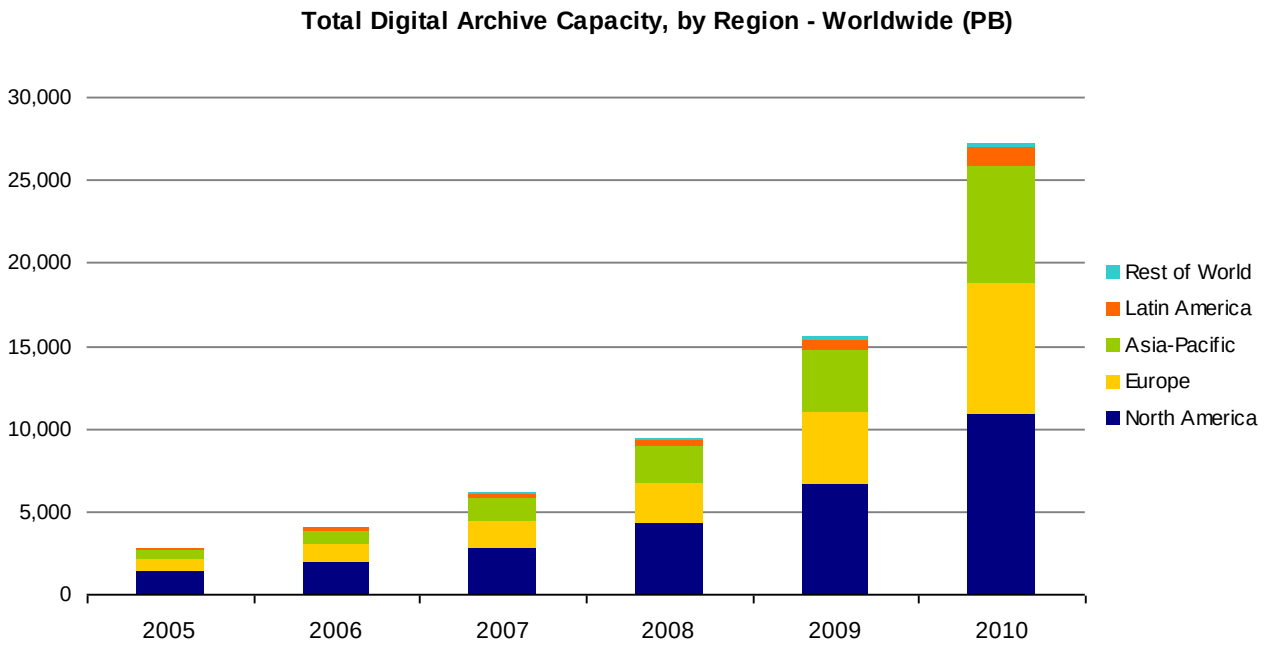

Source: Enterprise Strategy Group, 2007

The SGI Offering

The SGI Value Proposition

To address the obvious needs of a coherent and relevant migration offering (enabling both 'required' and 'choice' migration), as well as to offer additional data management possibilities, SGI is now delivering its Virtualized Storage Migration Solution. Although 'SGI' and 'storage management software' are not phrases that are automatically linked in many peoples' minds, SGI has nonetheless been gradually developing its offerings and a presence in storage. This has been a part of its drive to expand into a more general 'enterprise' space in addition to its well-known HPC capabilities. The Virtualized Storage Migration Solution is clearly an attempt to gain not only direct revenue for the product and services themselves, but also to encourage more users to consider SGI for storage. The overall offering is built on the foundation of the Storage Virtualization Manager (SVM) product from LSI, and as such comes with pedigree (some 6 years of shipping) and proof-points (more than 500 SANs in over 25 countries). Underlying the hardware, processing power comes from the aforementioned Dual Core Intel Xeon Processors. In addition to the hardware and software, SGI is being very clear that its approach embraces a full range of professional services to enable successful implementations. With an aim of central and simple management, Virtualized Storage Migration Solution—as its name tells you—uses virtualization to deliver migration; however, as the word 'virtualization' additionally tells you, migration is not the only thing that the product can do. It is, however, the immediate focus for SGI as it enters the market.

Overview of the SGI Virtualized Storage Migration Solution

The high level attributes of Virtualized Storage Migration Solution are easy to summarize as they are the migration prerequisites defined by all the challenges and requirements listed in this document so far. It offers:

- Non disruptive, online migration—applications keep running

- Heterogeneous storage support—the 'back end' can be a meritocracy
- Test and fall back— a user can track progress and always has a recovery route
- Networked storage virtualization—the product can be extended beyond migration alone
- Tried and tested solution—user risk is much reduced compared to an in-house development
- SAN deployment—offers maximum flexibility, connectivity and supports heterogeneity

How it Works

The key to the attributes that Virtualized Storage Migration Solution can offer is the 'split path' architecture that the solution employs. The term 'split-path' is given because the responsibilities of data transfer and control have been 'split' (see Figure 2), compared to an 'in-band' approach which handles both jobs. Because the data path is implemented in fast hardware within storage switches, a split-path approach should in theory both perform and scale better than an in-band approach. In-band approaches available on the market today are implemented using general purpose commodity servers. Data flows through host bus adapters and device drivers under the direction of a general purpose processor and common memory subsystem. A split-path approach instead uses purpose-built hardware implemented with a high-speed chip to move data. It also allows much more scalability as many data paths move data independently between them. In simple terms, this also means that control commands are not in the data path slowing it down… while the separate control path gives more freedom in terms of the commands that can be used. And, since just about any storage device that can present a LUN will work with Virtualized Storage Migration Solution, the product is both heterogeneous and very flexible.

Key Operational Features and Functions

From a basic migration feature perspective, Virtualized Storage Migration Solution has been tested widely for interoperability, is flexibly scalable (in terms of the control path, data path, metadata and data movement) and easy to port. Although production applications can remain online during migrations, no significant server or application level changes are required (there is just a small client driver code insert to connected servers). When a migration has been completed, the original volume remains intact for back-up purposes. All storage can be managed as one entity, rather than as 'vendor-islands' with different management suites.

FIGURE 2. SPLIT-PATH INTELLIGENCE

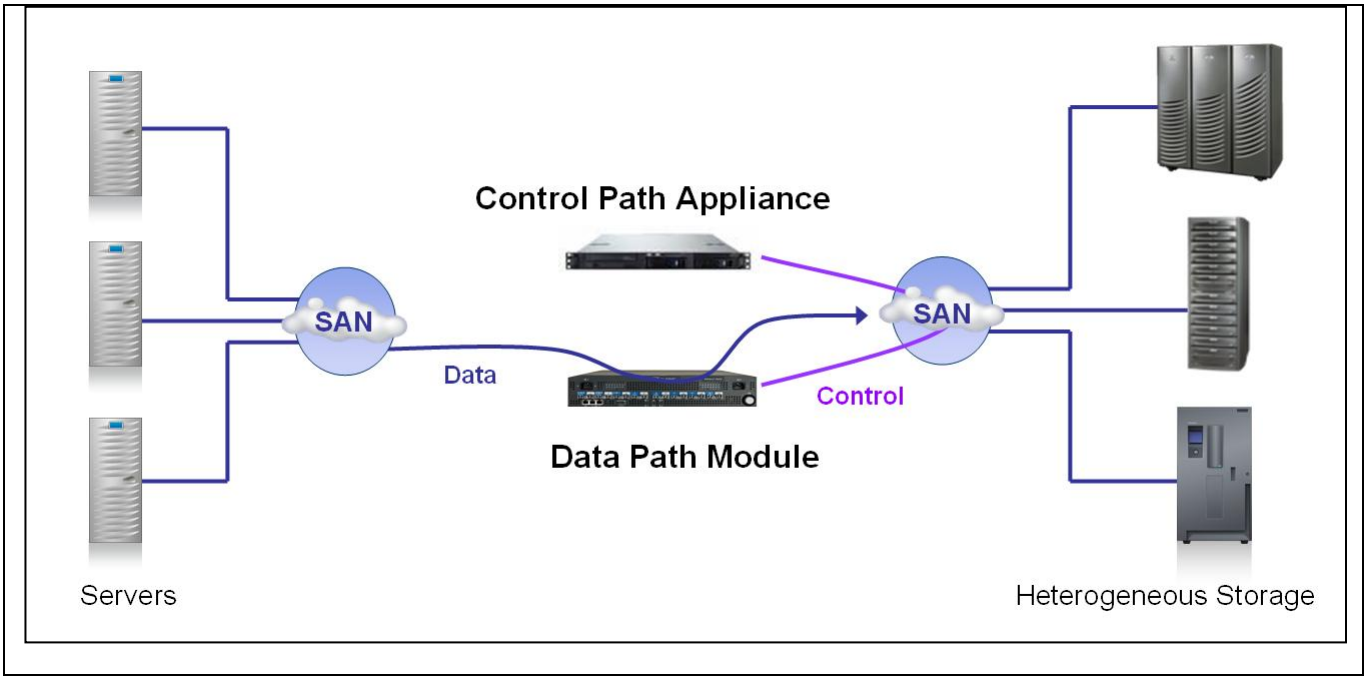

Functionality can be broken down as follows:

'Standard:' Once the Virtualized Storage Migration Solution has given a user a virtualized storage network layer, the functions become extensive. Of course, migration is the current SGI focus, and in this respect, heterogeneity means that underlying technologies and products can change and there's no impact to the data. Moreover, because of the split-path implementation, network elements can also change, to provide additional capabilities, without requiring user data to be moved or altered.

'Advanced:' After 'required migration' has been dealt with, the greater opportunities and efficiencies of 'choice migration' can be enjoyed since the traditional pains and hassles associated with migrations have gone. In this respect, it is easy to test a new storage vendor or to instigate tiered storage policies. This starts to show how this solution can become an 'enabler' as much as a 'fix;' something that is extended by the last functional category…

'New:' Virtualization enables other functions. For example, a user can take snapshots and have multiple copies of the same data—this can be very useful, for example, to support the iterative testing of new versions of applications against a common set of data.

ESG's View

Migration is a market that's ripe for improvement. On the demand side, user adoption of improved techniques has been slow—in part that's attributable to simple user intransigence, which itself can be somewhat explained by a desire to minimize risk (despite the operational frustrations and potential business impacts). However, beyond this, quite simply, on the supply side, the technical challenges of network heterogeneity have been greater than expected—especially in terms of the myriad physical combinations, the intricate and expansive error cases to be understood and integrated, and the difficulties to get speed up and latency down in the associated intelligent switches. Not only are areas of such interoperability and integration hugely challenging, but vendors have marketing reasons to procrastinate as well. Of course, everyone will publicly agree with how important it is to all work together, but why hurry a heterogeneous solution when a proprietary one generates better, and more guaranteed, returns?

In offering its Virtualized Storage Migration Solution product, SGI has a lot of upside and opportunity. It has no significant storage market to protect and it has taken on a tried-and-tested product that is running in hundreds of sites and that meets all the needs of an excellent migration solution. ESG does not see any significant limitations or drawbacks—quite the contrary. If SGI is serious about being in the enterprise space, about being in storage and about using a migration tool as an entrée to help achieve these objectives, then it has taken an excellent product step. The split-path architecture promises better performance, scalability and availability than the more common in-band appliances used for virtualization and migration and of course, this standards-based SAN implementation makes rare-but-enticing heterogeneous support practical. SGI wants users to consider its storage too, but it has adopted a genuine 'no strings' agnostic approach—combined with a pay-as-you-grow model—to enhance its differentiation from the large incumbent vendors.

The challenges for SGI, as so often in this business, are about sales and marketing—not product. Today, SGI's migration offering is not well known and, further, it's not in an area of the business where the company is well known. That said, with IT efficiency (as opposed to just effectiveness) being all the rage right now, it has chosen a good time for its market entry. As demands to reduce expenditure and downtime increase, a network-based, online, heterogeneous and non-disruptive migration tool such as Virtualized Storage Migration Solution makes eminent sense. The virtualization upon which this offering is built actually constitutes a double-edged marketing sword for SGI. The underlying capabilities of the product can handle much more than just migration; however, SGI is only mumbling rather than shouting about this, as it has chosen to concentrate on just one IT pain point rather than try to set itself up too broadly. With its good, but limited PS bandwidth and its lack of presence in storage management, this is, on balance, probably sensible. The problems with migration are easily and immediately understood, it gets focus in its distribution channels and it keeps its marketing claims limited and

therefore credible. If SGI can get some traction by addressing migration issues, it can open up the products' capabilities and take the offering from 'walking' to 'running' very quickly.

Effective, widely used data migration abilities are much overdue—and necessary. Data migration 'done right 'can provide genuine business value in addition to simply smoothing operations. Users seeking to address this issue, and regain control of what storage they use and at what overall cost, should certainly put the Virtualized Storage Migration Solution from SGI on their consideration list.

20 Asylum Street
Milford, MA 01757
Tel: 508-482-0188
Fax: 508-482-0218

www.enterprisestrategygroup.com