Maximize Technology ROI Through Better Operations Management

GlassHouse Whitepaper

By John Richardson,
GlassHouse Technologies,
Managed Services Principal

Technology upgrades are a constant in any IT environment, often driven by the need for more capacity or a platform's end of life. Over the past few years, IT shops have been upgrading to improve their return on investment (ROI) as well. According to a recent Enterprise Strategy Group study, 35% of organizations plan to purchase new technologies with improved ROI as a cost reduction or cost containment strategy. [1]

In many cases, the improved ROI comes from consolidating multiple disparate, under-utilized resources into a centralized, shared pool for optimal utilization. These consolidation strategies reduce data center footprints, reduce maintenance costs, and sometimes have positive impact on licensing costs. But they also bring pressure to bear on the IT teams who will manage them in two key ways.

First, if you are implementing the latest consolidation technologies, you will need to be more engaged with the business users in order to properly rationalize the new capabilities to business requirements. Second, you will also need to upgrade operations management capabilities to address the technologies' higher levels of capability and risk.

Enterprises that don't properly rationalize the new capabilities to business requirements may face political challenges over mismatched levels of service. Those who don't upgrade their operations management approach will be met with stability and performance issues resulting from the greater sensitivity of the highly consolidated environment. Keeping up with both sets of demands will tax many IT shops, who want to focus their internal resources on business alignment and leave the growing technical challenge to external service providers.

How New Platforms Drive These Changes

To understand the inter-relationship between new technologies and operational management, we'll look at specific technology changes. By examining the effects of an established consolidation technology, like the Storage Area Network (SAN), we can see the ongoing impact of that change years after deployment. If we look at server virtualization and data deduplication, we have examples of platforms that are just beginning to mature. To glimpse into the near future, we explore automated storage provisioning.

[1] 2010 Spending Intentions Survey, Enterprise Strategy Group, January 2010

Storage Area Networks

Until SANs came along, open systems storage was attached directly to servers and managed and maintained by the server administration group. Now, by consolidating storage to a central pool, you receive several benefits:

- **Increased utilization of storage** – huge amounts of unused storage are no longer marooned in directly attached storage
- **Improved data protection** – having the storage in a pool makes it much easier and more cost-effective to back it up on a regular basis
- **Tunable levels/tiers of performance** – SANs allow storage to be provided at varying levels of price/performance, so that high end applications enjoy higher levels of performance while other requirements receive “economy” price/performance

But this transformation also brought along new business engagement and management issues:

- With differing tiers of storage, IT now needs a way to define specific storage service levels, “market” them to the business, and devise a logical way to map business requirements to storage tiers.
- Determining how much storage to request and to allocate is as much art as science, and most organizations that go through the effort of identifying over-allocated storage can reclaim significant amounts of costly storage due to the inexact nature of requests and the human tendency to hoard scarce resources.
- Beyond these business-alignment issues, the task of provisioning storage is technically challenging – organizations implementing their first SANs soon recognized the need for new skills to manage this new technology.

In spite of these new complexities and requirements, SANs have been credited with “simplifying” storage management. But is management simplified?

Previously, server administrators took care of storage, and probably didn’t spend a lot of time on it. With a SAN, we need a dedicated server administration function, policy for distribution of storage, process for provisioning, tools for keeping track of how storage is being used, and management practices to manage demand, optimize utilization and charge storage back to the business as a discrete service.

Anyone who is in the midst of storage management would likely agree that it was easier in the “good old days” before SANs, but would then quickly call out how SANs have also improved utilization rates and thereby reduced storage costs and improved data protection and performance.

That’s a look in the rear view mirror at a somewhat mature example. Let’s look forward to some current and future examples.

Data Deduplication

Adoption of data deduplication technology is growing rapidly - according to ESG, growth was over 200% from 2008 to 2009, with enterprises adopting at a rate of 45% and midmarket companies implementing at a rate of 29%.[2]

Among the touted benefits of deduplication are:

- Reduced backup volumes
- Improved backup performance
- Reduced tape consumption and recall
- Reduced failure rates

While manufacturers claim that management time is reduced by as much as 60% and that the technology truly is "set and forget," a number of issues have to be addressed to take advantage of deduplication's promise, including:

- Choosing the right configuration
    - Source vs. target processing
    - Hardware vs. Software vs. combination
- Validating interoperability with the environment
    - Is the particular choice compatible with current hardware/software?
    - Will the deduplication processing require re-architecture or upgrades?
    - Will existing communications lines support the load for remote offices?
- Understanding tape-out options
    - How easily can tape-out be accomplished?
    - Do tapes need to be re-integrated in order to restore?
        - How much time does this take?
- Once the technical issues are addressed, IT needs to engage with the business to determine if any data would be restricted from this process based on:
    - Time to restore
    - Discovery considerations
- In order to keep schedules within backup windows, administrators will need to manage performance keeping several variables in mind:
    - Amount of new data to dedupe
    - Dedupe ratios achieved vs. planned
    - Speed of dedupe
    - Network bandwidth (for remote applications)
    - Appliance capacity (when used) leaving adequate room for system and processing overhead

[2] Oodles of Data, Enterprise Strategy Group, May 2010

Again, we have a combination of technical and business engagement considerations that need to be addressed in order to take advantage of a technology's promise. Does deduplication "simplify" backup management, or have we again raised the management bar – both technical and business – in order to achieve hardware savings? As we continue to look at these technologies, it becomes increasingly evident that getting the desired ROI is dependent on both managing the technology effectively and matching the new capabilities to specific business needs.

Server Virtualization

Just as SANs pooled what was once "just a bunch of disk" into a centralized platform, server virtualization brings many of the same benefits to formerly decentralized, underutilized servers:

- Increased server utilization
- Reduced data center footprint
- Decreased energy consumption
- Reduced server maintenance costs

In a virtual server environment, there is more complexity to manage and fewer margins for error since you're putting more eggs in each basket. You need to be prepared for:

- Balancing the critical function of performance management while maximizing utilization
- Addressing increased complexity in storage and backup management as you abstract the hardware
- Reassuring apprehensive business users , who may insist on dedicated physical servers for their applications
  - Be ready to provide performance SLAs to business units to convince them that virtualized servers are acceptable
  - Measure and report against SLA performance at a business level
- Leveraging tools and processes to prevent failures before they occur due to increased density and a higher level of interdependence
- Managing higher risk and performance issues related to the increased speed with which virtual servers can be spun up can accelerate the rate of change in an environment

These challenges mean that enterprises are quickly identifying a need for:

- **Skill sets experienced in virtualization –** learning as you go with production servers is not the preferred model, yet there still is a wide gap between supply and demand for these skills
- **New monitoring and reporting tools that are optimized to virtualization –** legacy toolsets can often see individual VM guests but don't provide host-level monitoring and reporting critical to maintaining performance and stability.

Organizations are increasingly reluctant to invest in new tools given the historically rocky track record of previous monitoring investments.

- **Proactive processes** – IT shops that relied on "firefighters" to save the day will have to shift gears and focus on process, documentation and problem avoidance as the cost of downtime escalates with consolidation density.
- **Management practices** – with so many variables in play, close monitoring of performance data, SLA achievement and improvement isn't optional – it's required in order to keep customer confidence and system stability where they need to be.

Server virtualization clearly tells the story of how enterprises are taking on increased management complexity in order to gain critical ROI in the data center.

The previous three examples show how management complexity can increase with new technologies, but as we look to the future we're going to see an example of how successfully automating low-end tasks can require even deeper knowledge of higher level business rules.

Automated Storage Provisioning

The advent of solid state storage, which provides very high performance for very high prices, is driving the evolution of Automated Storage Provisioning. The underlying concept is that a few applications will need this very fast and very scarce performance level, and the most efficient way to accomplish this is to have a multi-tiered array decide when to move data in or out of each tier to maximize efficiency.

Benefits of this approach would include:

- Making a relatively small amount of solid state storage be virtually available to a relatively large number of applications
- Resting inactive data on low-cost storage but then automatically moving it to the appropriate higher tier on demand
- Reducing the amount of manual provisioning done by storage administrators

Getting the value from this technology will require:

- Technical skills to understand how the technology works
  - How is the array instructed, what parameters are available?
  - How long does it take to move data from one tier to the next?
  - How sensitive are the triggers – could we have our data playing "bumper cars" and hurting overall performance as the array continually optimizes the data locations?
- Business skills to understand
  - Which applications need high performance
    - On which schedule

- Bringing both pieces together to handle issues such as:
  - Which applications should never reach tier 0
  - Which applications should never reach tier 3
  - How to handle seasonal requirements
  - Where to put new data sets
  - How to handle anomalies (poorly written SQL queries which could trick the array into giving them tier 0 priority)

Looking across all four of our examples, we can see some patterns emerging. Across the technology spectrum, the gap between high performance and low cost continues to widen, and so the demand on IT shops to rationalize these levels of performance to business requirements increases. While performance levels vary, different forms of risk come to bear at all levels of performance.

Ironically, most IT shops spend the bulk of their efforts evaluating the differences in technologies with little attention paid to the realities of management. As our consultants encounter infrastructures in need of repair, we do occasionally see poor product selection, but more often we see adequate products installed well, but not continually "tuned" as the enterprise grows and usage patterns change. The ongoing platform tuning often requires a sophisticated mix of technology and business knowledge, and yet often internal resources are left this responsibility with little or no training.

Why is there a focus on technology selection and initial design instead of ongoing management? It could be a characteristic of the manufacturers' selling process – their priorities may be to position their products favorably, focus on the ROI-driving features, and not emphasizing the management requirements. It may also be that buyers' selection committees favor visible hard ROI over hard-to-measure ongoing management considerations.

Whatever the reason, enterprises of all sizes are feeling more pressure than ever to both deeply understand business needs and deliver to those needs with increasingly complex technology. They are making good choices on technology from an ROI perspective but rarely are making concrete plans on how to manage stay out in front of management of the new platforms.

A War On Two Fronts

The real work of technology operations management is keeping the platform "tuned" to the business through the various changes in infrastructure and user application requirements. Keeping track of both constantly-changing user needs and technology platforms can be like fighting a war on two fronts. As business dependence on IT continues to increase and the speed of change in technology ratchets up again, there is unprecedented intensity to the age-old problems of IT management.

Choosing an approach to this challenge is a critical strategy for IT leadership. Those who elect to try to do both internally will find that the gap between those who know the business well and those who know technology is widening. Every enterprise has some "gold-collar" employees, who know both technology and business - and can speak both fluently - but few have enough to respond to the increasing demands.

Best Practice Example

The best practice vision for this was clearly described by one of our midmarket banking customers several years ago: "I'd like my team to be focused entirely on our internal customers, making sure we know exactly what they need and staying in front of their requirements. As far as delivery is concerned, I don't want to be in that business – I'll 'resell' managed services to meet my customers' demands."

IT shops who choose to focus on their business units and leverage external service providers for the technology management exploit two major advantages:

- They are in the best possible position to understand business requirements at a detailed level
- External providers can bring scale, focus and experience that makes for better performance and lower risk than they could deliver internally

By focusing internal resources to gain high levels of internal customer intimacy, IT shops will get a clear view of current and future business requirements. More time spent with internal customers creates a deeper context for requirements, allowing for more effective provisioning of services and a greater level of collaboration. Offloading technology issues to external players allows IT to look at business alignment in an entirely new way.

External service providers potentially have several advantages over managing changing technology internally, including:

- **Skills and Scale –** External providers are forced by the market to stay in front of emerging technologies. In any given technology, they are most likely to have deep firsthand experience to call on – valuable 'hands on' experience that can go deeper than any support group or manufacturer tech support team. The sheer numbers of skilled consultants and engineers that providers have solving problems in real multi-vendor environments is impossible to duplicate in a single enterprise environment.
- **Tools –** Just as you'd expect your doctor or auto mechanic to have better tools than you do at home, service providers can afford more sophisticated tools for monitoring, automation and reporting because they are shared across multiple customers. In addition, providers have the skills to select and implement powerful tools, and the crucial experience to interpret the data they provide in a way that is meaningful to their customers.
- **Process/Documentation –** While it's often considered optional for thinly-stretched internal staffs, external providers have no choice but to use documentation and standardized processes to deliver results. The process and documentation has benefit beyond just compliance and consistent levels of

service – it also can help support business decisions such as upgrades that depend on detailed knowledge of the environment.

- **Vendor Leverage** – manufacturers know that answering a question for a single service provider is worth answering the same question dozens of times for individual customers. And they also know the trusted advisor status that service providers have with their clients, so most provide their best levels of support to those providers, looking to both reduce their cost of support and build client loyalty for future purchases.
- **Commercial Strength** – service providers solve the same problems that internal staff do, but they do so under the constraints of legal contracts with their customers. They are bound to produce specific outcomes regardless of the degree of difficulty. Service provider and sales staff, looking to ensure future business with each customer, provide additional pressure to deliver high levels of customer satisfaction. Most service providers avoid hiring support staff that don't have external experience due to the very different level of accountability and pressure inherent in the job.

Conclusion

Upgrading technology usually drives needs to also upgrade customer involvement and technology management approaches. Enterprises need to decide if they want to "double down" with new investments in tools, processes and skill sets to keep up with the technology. Meanwhile, customer requirements must be continually rationalized to these new platforms.

External service providers can provide a management layer to IT shops that recognizes the true complexity of new technologies and deals with all of it for them. The freedom to concentrate on internal customer requirements while knowing a professional provider is handling the technology details creates a new paradigm for IT leveraging technology. Visionary IT leaders are exploiting this strategy for the good of their enterprises and their careers.

GLASSHOUSE

GlassHouse Technologies is a global provider of data centre consulting and managed services. In a rapidly changing data centre environment, GlassHouse partners with customers to define a strategy, execute that plan and operate their environment. Our constant focus is on cost efficiency, risk mitigation and service improvement. We provide this through Transom, our unique business model comprised of proprietary software tools, methodologies and domain expertise.

We help deliver on the promise of agility and usage-based spending in the next generation infrastructure paradigm. This journey is significant and requires an experienced and pragmatic guide to help you achieve your goals. Our experience is based on thousands of projects in addition to ongoing daily operations of customer environments.