ITG | International Technology Group

Executive Brief

May 2014

Cost/Benefit Analysis of IBM Virtual Storage Center
Compared to EMC Storage Virtualization Solutions

State of Virtualization

Storage growth has created challenges in most large organizations. As data volumes grow, users face escalating costs. Quality of service is threatened, and processes such as backup, replication and recovery are challenged. Management overhead rises, and storage administration teams become overworked.

What is to be done? This report focuses on one technology – storage virtualization – whose potential has been inadequately exploited. While 40 to 60 percent of servers in large organizations are now virtualized, the comparable statistic for storage capacity appears to be under 20 percent.

One reason why this situation has developed is that disk array vendors have tended to focus on platform-specific virtualization, meaning that virtualization benefits are realized within single frames or small homogenous clusters. While this approach may provide value, it does not allow economies to be realized across entire storage infrastructures. Multivendor support, when available, is generally weak.

There have been some exceptions, including offerings by independent software and appliance vendors. The only solution from a major vendor to address large-scale multivendor virtualization, however, is IBM Virtual Storage Center (VSC). User experiences with this solution have shown that benefits are significantly greater than for conventional virtualization techniques.

This may be demonstrated by comparisons based on input from five multiple-petabyte (PB) VSC users, along with 15 users of EMC storage virtualization solutions. EMC users employed VMAX and/or VNX-based software, including Federated Tiered Storage (VMAX only), and/or VPLEX solutions.

The results are striking. In multiple-petabyte, multivendor installations, overall costs of ownership for use of VSC averaged 72 percent less than for EMC solutions. ***Average five-year savings of more than $11 million per petabyte are realized in costs of storage capacity, personnel and facilities*.** There is no obvious reason why other large organizations should not realize similar gains.

Capacities and Costs

Comparisons are for installations in telecommunications, financial services and manufacturing companies with initial capacities of 1.5 to 4.3 PB of raw disk storage on all tiers.

Hardware and software configurations, numbers of full time equivalent (FTE) storage administrators and costs were calculated for use of EMC solutions and VSC for each installation. Calculations allow for annual growth rates of between 15 and 45 percent, varying by workloads and tiers, over a five-year period. All calculations were based on user input.

A five-year measurement period was employed because the benefits of large-scale storage virtualization are typically cumulative – three-year calculations tend to understate long-term savings. This is particularly the case in that vendors of Tier 1 and (to a lesser extent) lower-tier systems tend to raise maintenance and support costs substantially after three years. EMC increases for high-end VMAX arrays are particularly steep.

After five years, total raw capacity in VSC environments averaged 46 percent less than for use of EMC solutions. This was due to higher utilization levels and more extensive integration of multivendor arrays.

In VSC environments, there was also a more pronounced shift of capacity from comparatively expensive Tier 1 to Tier 2 and Tier 3 arrays. These results are illustrated in figure 1.

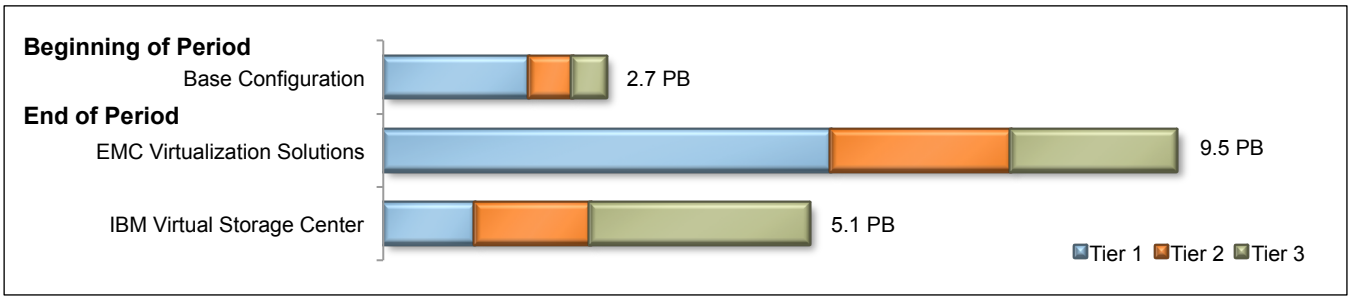

Figure 1: *Capacities by Tier – Averages for All Installations*

Initial capacities, which are the same for use of EMC solutions, averaged 66 percent on Tier 1 arrays, and 34 percent on Tier 2 and 3 arrays. After five years, use of EMC solutions results in 56 percent of capacity on Tier 1 arrays, and 44 percent on Tiers 2 and 3.

By the end of the same period, use of IBM VSC resulted in only 22 percent of raw capacity on Tier 1 arrays, and 78 percent on Tiers 2 and Tier 3. VSC is a great deal more effective than EMC solutions in enabling organizations to transfer data that does not require full Tier 1 capability to lower-cost platforms.

Higher capacity utilization, greater use of Tier 2 and 3 arrays, and lower FTE storage administrator staffing result in significantly lower five-year costs of ownership. Figure 2 illustrates disparities.

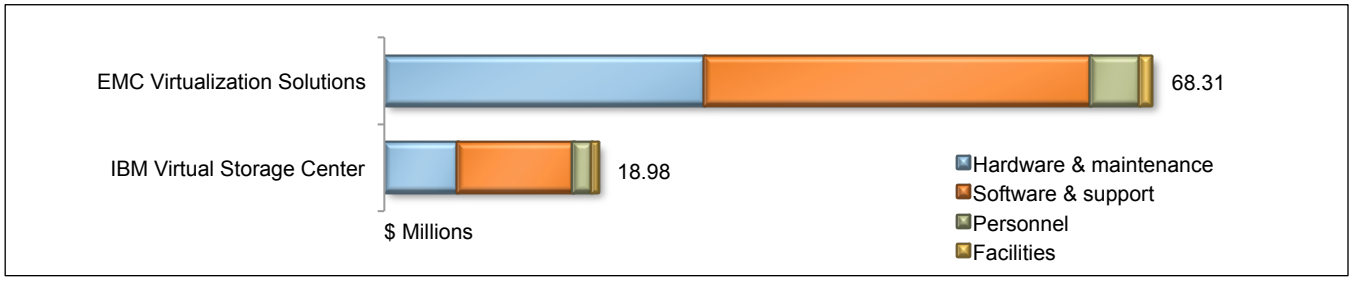

Figure 2: *Five-year Costs of Ownership – Averages for All Installations*

Costs of ownership include hardware acquisition and software licenses to handle growth over five years (costs of beginning-of-period installed bases are not included), along with hardware maintenance and software support, storage administration personnel and facilities costs including data center occupancy and energy.

Software stacks employed for calculations include operating systems and tools for virtualization, tiering, snapshot copying, storage management and, where appropriate, replication for disaster recovery. Costs of hardware, maintenance, and software licenses and support were based on discounted prices as reported by users.

Tier 1 arrays included EMC VMAX, IBM DS8000, and Hitachi VSP and USP-V systems. Tier 2 arrays included Dell Compellent, EMC VNX, HP 3PAR and Enterprise Virtual Array (EVA), IBM XIV and Storwize V7000 and others. Tier 3 arrays included a broader mix of Dell, HP, IBM and other low-end platforms. All Tier 1 and 2, and some Tier 3 systems were configured with tiered mixes of solid state and conventional disk drives.

Calculations do not include all-flash arrays, which were at a relatively early stage of deployment among users of both EMC solutions and VSC.

Further information on profiles, methodology and assumptions employed for calculations, along with cost breakdowns may be found in the Detailed Data section of this report.

Conclusions

Lower VSC costs reflect multiple factors. These include more efficient use of disk resources, data movement and management mechanisms that enable greater use of Tier 2 and Tier 3 arrays, and the ability to virtualize IBM as well as non-IBM arrays in a comparatively simple manner. FTE storage administration staffing is lower, reflecting larger virtualized capacities, and more effective automation and analytics.

Software costs are also reduced for advanced functions such as thin provisioning, tiering, copy services, disaster recovery, replication and storage resource management (SRM). Tier 1 solutions for these functions tend to be expensive – this is particularly the case for EMC. VSC provides comparable cross-platform capabilities as part of a single, lower-cost package.

VSC, moreover, requires only a single storage driver type. Device interfaces are handled through the VSC virtualization layer. No changes are required when data is moved or arrays are replaced or upgraded.

The ability of VSC to reduce dependence on Tier 1 arrays should be highlighted. Most large organizations place a majority of data – the industry norm is around 70 percent – on Tier 1 systems. Large VSC users have reduced this ratio to as little as 30 percent within a few years, while maintaining performance and availability targets.

The cost implications are important. Tier 1 arrays tend to be significantly more expensive than Tier 2 and 3 equivalents. In the calculations undertaken for this report, for example, five-year hardware and software costs for all Tier 1 platforms averaged $15,039 per terabyte (TB), while Tier 2 and Tier 3 systems averaged $2,559/TB.

Tier 1 costs reflect higher-priced hardware acquisition and maintenance (an average of $6,101/TB) as well as licenses and support for advanced function software (an average of $8,938/TB). Comparable costs for Tier 2 and 3 systems were $1,973 and $586/TB respectively.

EMC has consistently sought to differentiate its high-end arrays. The company's Storage Tiering for Virtual Pools (FAST VP) has been employed primarily within VMAX or VNX environments, while VPLEX has been targeted at midsize organizations using VNX arrays. The company's track record in multivendor support has left much to be desired.

In 2013, EMC announced a new scheme, *ViPR*, which the company characterizes as a software-defined storage framework enabling large-scale virtualization of multivendor storage installations.

ViPR was initially presented as an "open" virtualization framework. It is, however, now moving closer to EMC proprietary hardware and software. High-end VMAX arrays remain differentiated, and integration of non-EMC platforms requires that storage vendors, software companies and/or users develop – and update – interfaces conforming to EMC application programming interfaces (APIs).

ViPR is still at an early stage of development, and it will probably be years before it is anywhere near complete. In contrast, VSC components have been deployed together for more than a decade.

A broader difference should also be noted. ViPR is an architecture. VSC is a solution. Although VSC will clearly evolve in the future, it is essentially complete now. It may be deployed "as is" and can typically yield significant bottom-line gains within a few months to a year.

Asked to describe key characteristics of VSC, one user observed simply, "it works." In the world of storage virtualization, this is not a small compliment.

Additional Information

This ITG Executive Brief is based upon the results and methodology contained in a Management Report released by the International Technology Group. For copies of this Management Report, please email requests to Contact@ITGforInfo.com.