

WHITE PAPER

Eliminate Shortcomings in Your Cloud Architecture with Smarter Storage

Sponsored by: IBM Corporation

Laura DuBois

January 2013

EXECUTIVE SUMMARY

Enterprises are deploying private and hosted clouds to respond faster to business demands for new applications, to improve service levels for given workloads, and to reduce cost across the datacenter. The enterprise storage architecture deployed is paramount to ensuring agility, performance, and reliability of a private or hosted cloud. Integration of storage infrastructure across hypervisors, cloud orchestration software, and open source architectures is required. Storage architectures for private and hosted clouds must support a range of mixed application workloads and the dynamic placement of data on the right storage tier. These essential storage functions enable faster service response times, improved service-level agreements (SLAs), and greater agility, performance, and reliability for private and hosted clouds.

IBM has a differentiated storage portfolio, including XIV, Scale Out NAS (SONAS), SAN Volume Controller (SVC), Tivoli Storage Manager (TSM), Tivoli Storage Productivity Center (TPC), and SmartCloud solutions. These offerings enable firms to optimize their storage systems in order to achieve their private and hosted cloud objectives. This paper examines the benefits of and barriers to private cloud deployments and identifies the storage challenges enterprises face with these deployment models. It also discusses IBM's robust storage solutions with a focus on the state-of-the-art XIV storage architecture. XIV's ease-of-use, reliability, and performance attributes eliminate storage vulnerabilities that firms face when implementing private or hosted clouds. The offering is being used in some of the largest private and hosted clouds today. IDC spoke with several IBM customers using XIV in private and hosted cloud deployments. Their comments are included in this paper.

SITUATION OVERVIEW

As a result of business pressures and economic realities facing IT organizations today, technology executives need to continuously implement more efficient ways to run IT services for their companies. Increased reliance on IT systems to run business operations has heightened the need for reliability and performance across the datacenter. The consumerization of IT and the speed of global business have brought about demand for self-service, personalization, and instant access to IT. Acceleration in the delivery of products and services dictates faster availability of IT services and infrastructure. Business units and executives expect IT to quickly respond to new application and project requests while meeting service-level requirements. These factors are driving IT organizations to rethink the way infrastructure is provisioned, allocated, and managed for applications and lines of business. This forces IT executives

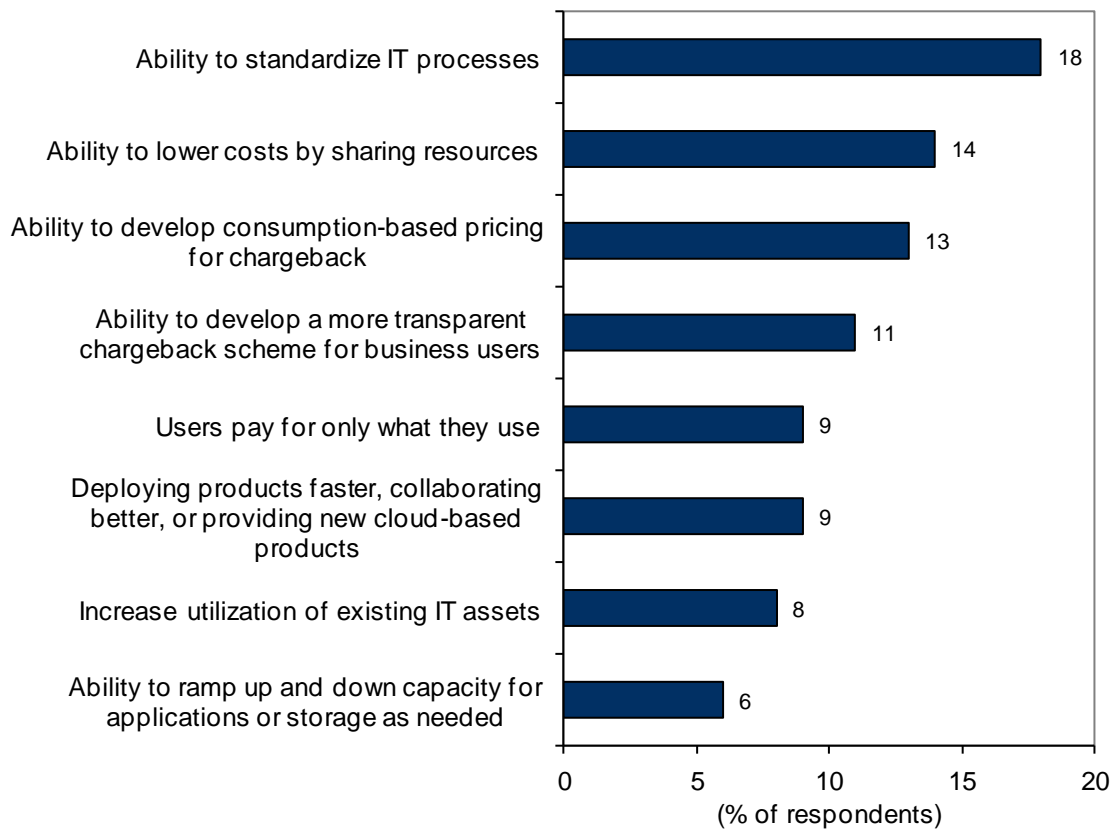
to implement cloud projects that leverage shared infrastructure and the necessary automation and management controls to run IT as a service.

An IDC cloud study of 800 firms conducted in June 2012 cited cost savings and faster time to market as the two most critical benefits of cloud. However, cloud attributes including user self-service, standardization of IT services, elastic scaling, metering/chargeback, and flexible migration between public and private clouds were also important. Among firms considering or deploying cloud services, concerns with public cloud include lack of interoperability standards, data/system portability, and lack of migration choice/flexibility. However, private cloud infrastructure, both on premise and hosted in a virtual private cloud or across a dedicated infrastructure, can address these concerns. As Figure 1 shows, datacenters are considering and/or deploying private or hosted clouds for three leading reasons: to standardize IT processes, to lower costs by sharing resources, and to develop consumption-based pricing for chargeback.

FIGURE 1

Drivers for Use of Private Cloud

Q. Which of the following best describes the primary reasons why you are using — or planning to use — a private cloud?



n = 800

Note: Responses of 4% or less were removed.

Source: IDC's *CloudTrack Survey*, 2012

Private Cloud and Hosted Services Cloud Infrastructure

In principle, a private cloud is implemented by a single company but has the same infrastructure attributes of elasticity, shared infrastructure, secure multitenancy, self-service provisioning, and so on as a public cloud. A private cloud can be managed internally or by a third party and hosted internally (within the company's datacenter) or externally. However, obstacles of security, service reliability, and internal IT staff knowledge required to reorient IT operations to private cloud are concerns with private cloud. Security concerns stem from leveraging shared infrastructure and broader access to critical IT services. Private cloud features, including stringent role-based permissions, integration of service catalogs, secure multitenancy, two-factor authentication, and encryption can address security concerns. Limited internal staff capabilities can be met with training/certifications and/or the deployment of a private hosted cloud that leverages service provider skills. A system that is easy to learn and use gives firms an advantage in terms of staff proficiency. Private cloud service uptime, response time, user downtime, and reliability SLAs are satisfied, in part, with the right underlying storage infrastructure.

As this research highlights, firms starting to implement a private cloud infrastructure often face issues in the infrastructure and storage domains, which can stymie the overall functionality and benefits of a private cloud offering. Storage vulnerabilities in private and hosted clouds are identified in the following section.

Storage Challenges with Private and Hosted Cloud Deployments

Deployment of a private or hosted cloud requires a range of technology components. Compute and storage virtualization are table stakes. Also required are self-service provisioning, automated provisioning, workflow orchestration, SLA management, monitoring, metering, and chargeback. Achieving these capabilities requires the integration of storage management functions with hypervisors and cloud orchestration tools. Lack of such integration is the single biggest storage challenge in private and hosted cloud deployments.

Automation for Standardized Storage Services

With virtualization, the time to provision new systems has been dramatically reduced. However, separate workflows and provisioning processes are still required to carve out and present logical unit numbers (LUNs) to the hypervisor. This creates separate and often disjointed workflows that can introduce delays and errors. For a private cloud deployment, a single provisioning action and workflow process should provision the virtual machine (VM), networking, and physical storage. This requires that a set of storage services (or a storage services catalog) be presented, which allows an IT generalist to easily select the right storage policies based on a given VM or application profile. Firms deploying private clouds at scale require automation via cloud management frameworks, such as OpenStack.

Enterprise Storage Complexity

Configuring storage for a private cloud deployment can be anything but efficient and automated. Administrators must select RAID groups, set up multipathing and host configurations, determine thin or fat volumes, and configure snapshots and/or replication.

Depending on the storage used, painful LUN resizing or expansion processes may be required. If more storage capacity is required, a rebalancing project may be needed to manually balance usage of the storage infrastructure. Once storage is configured, ongoing tasks such as distribution planning, load balancing, performance monitoring, and capacity or performance upgrades must be conducted. This specialized skill in managing storage inhibits the ability of general-purpose IT staff to provision storage for their own assets.

The issues we had were four different types of storage: It became harder and harder to manage the data. To get TBs free to link them to a new server was very difficult. Storage configuration was tricky. We had trouble extending the current storage capacity — being able to leverage storage that was underutilized while expanding storage in other places; doing this without bringing down any systems was hard."
— Executive Director, Educational Institution, Germany

Insight in Storage Resources Among Tenants

Metering, chargeback, and showback are fundamental benefits to leveraging shared storage infrastructure across multiple departments or organizations. IT organizations as well as lines of business seek visibility into how storage resources are being consumed and perhaps ultimately to pay for only the resources they consume. However, it remains very difficult for a cloud team to get an accurate picture of storage resource allocation, usage, and cost according to "logical" organizational or departmental boundaries. Leveraging private cloud infrastructure across different IT groups, business units, and subsidiaries requires a level of granularity in resource tracking and utilization.

Standardizing Services, Satisfying Business-Driven SLAs

IT professionals seek to provide faster time to infrastructure deployment, more rapid response to new business projects, and the ability to scale infrastructure dynamically to maximum capacity periods. The challenge is meeting business-driven time-to-market requirements while having only weeks to get infrastructure services up and running. To meet specific SLAs, IT organizations are defining standard services (or what one firm called "common patterns") within the infrastructure.

The ability to meet documented SLAs on performance, resiliency, availability, and recovery is heavily dependent upon the storage infrastructure. Cloud SLAs mandate nondisruptive maintenance windows, including microcode and hardware upgrades. All components in a private or hosted cloud storage system must be able to be added without bringing down user and application access or introducing performance degradation. Performance requirements vary based on the storage tier, but performance response times are measured in milliseconds, even during overhead processes such as RAID rebuilds. Storage performance needs to keep pace with the downstream demands that cloud I/O and virtualization place on storage processing. Hotspots can be prevented with the right storage architecture. When hotspot conditions are encountered, the system must be able to dynamically respond, ensuring that processing is balanced across applications and users.

Dynamic Scale, Mixed Workloads, and Performance and Capacity Optimization

While improvements in CPU processor speed and computing power have been realized, many workloads with higher IOPS and/or lower latency requirements become spindle bound. Server virtualization and client virtualization exacerbate the performance problems even further by increasing the amount of and randomization of I/O. The use of techniques such as short-stroking drives, adding more spindles, and/or introducing SSDs to deal with performance constraints is common. However, private cloud services seek not only to deliver IT efficiency and faster response times but to do so in a cost-effective manner while maximizing utilization of infrastructure.

Private and hosted clouds seek to deploy storage architectures that can support a range of workloads. With traditional storage deployments, the storage vendor wants to optimize the system for the workload, databases, Exchange, etc. With a cloud offering, the cloud architect or provider likely will not have insight into how the customer will be using the storage. Thus, manual tuning for a specific workload is not a viable option. The customer will ask for storage and for an SLA, and that's it. However, not many storage vendors can deal with mixed workloads.

We can't just buy a general storage box. We need 'specialized hardware,' and what that means is that this is SLA driven. I need hardware that can mix any workload and be able to support a certain SLA for our performance tier. — Storage Architect, Leading eCommerce Company, United States

IBM STORAGE SOLUTIONS OPTIMIZE CLOUD DEPLOYMENTS

IBM has a differentiated portfolio of storage solutions including scale-out, block-and-file storage systems, storage management software, data protection and recovery offerings, and cloud management being deployed in private and hosted cloud architectures today. At the center of IBM's private and hosted cloud strategy is the IBM XIV storage system, which can be deployed standalone or behind an IBM SVC virtualization engine or SONAS gateway. XIV's scale-out storage architecture offers storage automation, ease-of-use, performance, and reliability attributes that are required for private and hosted cloud deployments. XIV also integrates with TSM, TPC, and SmartCloud storage offerings to provide a comprehensive data protection, storage management, and cloud management strategy. Details of the integration between XIV and IBM's other private and hosted cloud storage offerings are included in the next section. First, we provide an overview of the XIV architecture and its distinguishing features.

IBM XIV AND HOW IT ELIMINATES STORAGE VULNERABILITIES WITH PRIVATE AND HOSTED CLOUD

From IDC's perspective, the following features distinguish XIV from other block storage architectures available in the industry:

- ☒ A distributed grid-based storage approach with massive parallelism

- ☒ Unique data distribution architecture/algorithms
- ☒ Performance and reliability attributes
- ☒ Virtualization, cloud, and open source integration
- ☒ Easy-to-use GUI and provisioning simplicity
- ☒ Modular design in which each module has its own multicore CPU, RAM cache, SSD as cache (optional), and commodity HDD capacity, serving as a single elastic data resource
- ☒ No tiers to manage, which eliminates complexities and offers strong performance gains over traditional RAID

The XIV architecture is virtualized from the ground up, with the ability for massive parallelism to gain enterprise performance from commodity components in a single (drive) tier of storage. The XIV management model is differentiated by the ease-of-use and storage management features the architecture affords.

Another factor differentiating the XIV approach from market alternatives is the inclusion of data protection functions such as snapshots and replication services natively with the array. The XIV system architecture is not based on internal storage tiers; instead, the whole system, comprising both HDD and SSD, is presented as a single tier to applications and users. Thus, in effect, more drives are working for each application, which offers massive parallelism and performance advantages.

XIV Architecture

The XIV distributed architecture comprises a series of modules that are interconnected and communicate over an internal, redundant InfiniBand or Ethernet network. Each module consists of its own multicore CPU, RAM cache, SSD as cache (optional), and HDD capacity, which all work in parallel. This parallelism is important because the XIV system distributes storage processing among the nodes or modules. Thus, as the number of storage modules increases, so does the amount of parallelism available in the system. The XIV system leverages these modules to present a single large, elastic data store to the application network. A flexible, distributed cache allows the XIV system to leverage large slots for reads while managing a smaller slot size, resulting in a superior cache hit ratio and, consequently, better performance. Aggressive prefetching is enabled by the large cache-to-disk bandwidth available within each module, together with the extremely large aggregate module interconnectivity bandwidth that is available on the XIV backplane.

As data enters the XIV system, it is randomly and evenly distributed across all modules and disks in the system through XIV's data protection scheme. XIV stores and distributes data by breaking it down into 1MB chunks called "partitions," each mirrored for redundancy to another module, ensuring no single point of failure. The system distributes all the partitions automatically and uniformly across all the disks by means of a pseudo-random distribution algorithm. This approach to resiliency offers performance gains over traditional RAID rebuild times while eliminating the overhead in initial RAID configuration. The XIV system automatically distributes the application load across all modules evenly,

putting the combined power of all modules at the service of all requests concurrently. When the number of functional disks or modules changes, the system is rebalanced to ensure optimal data layout and, consequently, optimal use of system resources at all times — without the need for manual intervention or capacity rebalancing.

With the XIV architecture, the storage system does not need the tuning or changes that other enterprise systems require. When changes to an application or its I/O pattern occur, XIV automatically rebalances resource utilization across the system. Having a single drive type and pool within the storage system negates the need for storage tiering complexities. To ensure application performance in the storage pool, XIV can set quality-of-service (QoS) characteristics through Performance Classes, limiting IOPS and MBps throughput per host. This eliminates the need for storage administrators to monitor and tune storage for performance over time. The addition of disk modules is accompanied by a corresponding increase in processing power, cache, storage capacity, and connectivity. According to IBM internal research, for OLTP-type workloads, SSD caching provides up to 4.5 times the performance without much extra cost. The system monitors its components continuously, reacting to existing and potential issues by activating self-healing as needed and returning to full redundancy rapidly and without human intervention. As part of this proactive approach, the system uses disk-resident diagnostics to predict potential disk failures. It accesses all disk drives and areas to assess health status and increases protection levels by retiring suspect disks before they fail and rebuilds their data while a redundant version is still available. Further, XIV storage is server agnostic, working well with IBM hosts and other vendor systems.

XIV Integration Within the Cloud Ecosystem

Hypervisors form the underlying platform for a private or hosted cloud. Datacenters with private and hosted clouds must support a range of hypervisors. XIV supports all the major virtualization platforms, including VMware vSphere, Microsoft Hyper-V, Xen, and KVM. Many leading private and hosted clouds are using open source technology and standardizing on OpenStack. As a result, integration with the OpenStack Cinder API or other cloud frameworks such as VMware vCloud and IBM cloud management tools is essential. In OpenStack environments, which currently use only the iSCSI protocol, XIV enables high performance by means of high-bandwidth 10GbE connectivity. XIV's support for this API differentiates the offering from other storage arrays, although they will follow in time.

For cloud management, XIV can be integrated with the IBM SmartCloud storage family. Additionally, XIV can serve as the high-performance storage architecture behind IBM SmartCloud Enterprise+ deployment. XIV supports VMware vStorage APIs, including VAAI and VASA, and also offers plug-ins to virtualization management frameworks, including vSphere clients as well as support for vCloud Director and Hyper-V System Center, among others. Microcode within XIV supports Microsoft System Center Virtual Machine Manager (SCVMM) 2012 out of the box. The IBM plug-in for vCenter (IBM Storage Management Console for VMware vCenter) combines the benefits of the XIV architecture with storage management functions that can be performed from the vCenter client. Storage functions such as volume creation, autodiscovery, capacity reporting, and monitoring can be performed. This integration with virtualization is essential for private and hosted cloud deployments because it enables greater visibility into and management of storage by a broader set of IT administrators.

XIV has an adapter for VMware vCenter Site Recovery Manager (SRM) that helps support business continuity solutions by syncing storage failover with SRM failover. IBM SVC also contributes virtualization advantages by supporting VAAI and SRM and offering visibility through the IBM vCenter plug-in. For overarching cloud management, enterprises can leverage the IBM SmartCloud storage family, integrating XIV storage, SONAS, and SVC. XIV and SVC system integration with OpenStack Cinder/Nova-Volume is an IBM strategy for addressing the growing trend to adopt cloud management frameworks.

XIV is ideal for cloud delivery of services. We use all of the tools within the XIV device to support cloud strategies, including virtualization and thin provisioning, so we can add capacity very quickly and scale up very easily. It's very, very elastic — it's very, very scalable — all of the things that cloud defines in terms of reactivity, and it has responsiveness to requirements. XIV is designed to support those sorts of business requirements. — Phil Clark, Channel and Partnership Director, niu Solutions, United Kingdom

IBM SONAS and SVC and XIV Integration

Storage for private and hosted cloud infrastructure from IBM includes IBM XIV, SONAS, and SVC. SONAS is IBM's scale-out NAS offering, ideally suited for applications ranging from high-performance analytics and rich media to computer-aided design, oil and gas exploration, and genomic sequencing that require SONAS' large addressable namespace and high-performance scalable storage. SONAS is being used in private cloud deployments in financial services, research, and educational institutions to provide petabytes of storage under a single management interface. XIV supports integration with the SONAS gateway, which enables administrators to dedicate some or all of the available storage capacity in the XIV storage system for network file serving. This enables XIV to provide a unified storage approach, with Fibre Channel, iSCSI, and NFS/CIFS protocol support, and extends the value of IBM in private and hosted cloud deployments.

IBM SVC is an innovative block-based storage virtualization platform that virtualizes and combines storage capacity from different storage architectures into a single resource pool. An SVC system is configured as a series of LU nodes and deployed in pairs that sit between application servers (hosts) and the storage arrays. The SVC system presents itself as a SCSI target to the hosts and as a SCSI initiator to the storage arrays. The SCSI LUNs that the hosts see are virtual disks or volumes, which are allocated by SVC from its virtualized storage pool. Administrators can manage a range of storage systems with a common approach and also optimize capacity utilization, which makes provisioning easier and enables nondisruptive data migration and copy services between disparate storage products. SVC is being used in private clouds today for data migration and ongoing virtualization of storage. An XIV system or multiple XIV systems can sit behind an SVC cluster to offer a scale-out block storage architecture, as well as coexist with other heterogeneous storage.

IBM TSM, TSM for Unified Recovery, and FlashCopy Manager and XIV Integration

For optimal protection and recovery of data and systems in a cloud, IBM offers TSM, TSM for Unified Recovery, and FlashCopy Manager. TSM is an enterprise-class protection and recovery solution for a broad range of platforms, applications, and virtualized infrastructure in both centralized and distributed locations. As more virtualized environments use array-based snapshots and clones to enable fast protection and recovery of applications and images, FlashCopy Manager offers a centralized, robust application-consistent hardware-based local and remote snapshot protection approach. TSM for Unified Recovery is a suite that includes all IBM's protection and recovery offerings — TSM, FlashCopy Manager, and FastBack — for end-to-end protection and recovery from a centralized management framework. For protection and recovery, XIV comes standard with its own snapshot and replication functions. However, XIV snapshots can also be centrally managed using FlashCopy Manager. For centralized backup, XIV data can be protected in an application-consistent manner using TSM, which supports tape, disk, and deduplication target appliances such as ProtecTier.

IBM TPC and SmartCloud and XIV Integration

XIV systems integrate with the IBM SmartCloud storage family to enable the automation of cloud storage services, including provisioning, monitoring, and access in a private or hosted cloud deployment. IBM TPC is the company's management solution for virtualized, heterogeneous server and storage environments, offering storage analytics, reporting, and control. TPC is the management framework to control all IBM storage systems, including XIV. Management of multiple XIV systems can be centralized and functions automated through TPC. Critical information on XIV health, performance, and utilization can be monitored and reported. The IBM SmartCloud family includes SmartCloud Monitoring, SmartCloud Provisioning, and SmartCloud Storage Access.

SmartCloud Monitoring provides visibility into physical and virtual private cloud infrastructure, monitoring health, capacity, and performance of heterogeneous private cloud infrastructure. What-if modeling and analysis of changes made in the environment allow cloud architects to mitigate risk and ensure that SLAs are not compromised. Policy-driven analytics ensures infrastructure is balanced according to performance and security rules. SmartCloud Provisioning enables consistent, dynamic application deployment, automated administration, and virtual machine image management in private and hosted clouds. It can be used to deploy cloud middleware and applications.

SmartCloud Storage Access is a self-service software portal for NAS storage provisioning and monitoring that includes reporting. In a few steps, users can provision storage and share files over the Internet. Administrators can easily monitor and report storage usage and classify storage for different organizational users. SmartCloud Storage Access also works with the IBM SONAS offering and with XIV behind a SONAS gateway, which enables fast time to value for firms deploying private cloud infrastructure-as-a-service capabilities.

WHY XIV FOR PRIVATE AND HOSTED CLOUD DEPLOYMENTS?

The successful deployment of a private or hosted cloud hinges on the right storage architecture. XIV can help firms overcome storage vulnerabilities associated with cloud deployments by offering the following features and benefits. Comments from customers evaluating and using XIV in private and hosted cloud deployments are included.

Agility and Elasticity

The XIV system can handle changing provisioning patterns across mixed, dynamic workloads and provide resource sharing across multiple tenants while maintaining consistently high performance. This is important as tenants enter or exit the cloud. The XIV architecture not only can scale to support peak workload demands that occur at the close of a quarter or as a result of marketing campaigns and seasonal activity but also can optimize for average workloads, thus supporting changing business patterns.

With XIV, it's easy for me to meet the SLA requirement [for that performance level] because it can absorb load and come back with a relatively fixed latency, and the throughput is relatively consistent.
— Storage Architect, Leading eCommerce Company, United States

Starting out with a minimum configuration of 56TB and scaling to 243TB across 15 XIV modules, the system achieves linear growth in *both* capacity and performance. With the XIV management model and user interface, it's quicker to create a virtual volume than a virtual machine, supporting faster time to market. Together, XIV and SONAS support scale-out configurations required by private and hosted clouds today. SONAS, with its global namespace, provides almost infinite scalability, while XIV performance scales with storage capacity. With these features, storage complexity, dynamic workload SLAs, storage provisioning overhead, and performance and capacity obstacles with private and hosted cloud deployments are overcome.

The challenge was looking for a storage vendor and offering that handled mixed loads, and that had consistent SLAs, for performance, throughput, and latency, as well as tiering them around availability appropriately. We checked, and only three vendors came out that could do what we wanted done. XIV was one of them. — Storage Architect, Leading eCommerce Company, United States

Ease of Management

The XIV architecture provides storage resources that require no tuning. Monitoring and management of the storage array are not required to keep performance consistently high and hotspot free. Storage can be provisioned and used by IT generalists.

It [storage provisioning] used to take one to two days to initiate. The provisioning operation of running the storage can now be done by email. It takes time to process the email and determine the priority of the request, but now it is down to minutes. The technical time to deliver storage is less because negotiation is no longer necessary. It can be configured in just minutes. — Executive Director, Educational Institution, Germany

Performance is paramount in private cloud deployments, and the XIV architecture ensures that application response times, as measured in milliseconds, can be sustained. This zero tuning approach is essential for cloud providers, where usage is not known, and thus a cloud architect cannot safely tune or optimize the storage. Other storage architectures leverage storage tiers, moving data between tiers based on usage/performance patterns. The XIV approach has no tiers to manage, so intensive storage planning is not required and data does not need to be moved between tiers, introducing performance overhead.

XIV's main advantage was it was one big storage [asset] and we could create slices from it and build internal drives. It solved a lot of complexity problems for us. — Executive Director, Educational Institution, Germany

Multiple XIV systems can be managed centrally, via the XIV UI or through IBM TPC. XIV supports integration with cloud management systems, including OpenStack via the Cinder API to provide storage automation, minimizing the effort and overhead associated with storage provisioning and deprovisioning.

We looked at management. Because of the bar we are setting on APIs, a lot of storage vendors fell off. IBM was willing to make an effort to integrate with Cinder and give us the open stack and onboarding capability. APIs are one of the things that we're standardizing on. — Storage Architect, Leading eCommerce Company, United States

So in terms of cloud services, we run predominantly private cloud services for our clients. We don't necessarily know what the application stack is that those clients are going to be accessing. We do know that it's XIV in the back. So when a client requires another four or five or six or ten virtual machines, they simply go into a portal, spin up those machines, and they're done. There was no consideration whatsoever about what the storage impact of that was going to be. It's just there, and it just runs. That's probably the best thing I can say about the product; it just works. So, from our point of view, the management overhead of actual sign-in storage on those cloud-based systems is minimal. We provision space once, and then that's it. The rest is automatic." Stephen Bedford, Technical Director, niu Solutions, United Kingdom

The XIV system captures detailed used-capacity information, which can be integrated with external cloud billing systems for showback or chargeback use cases. Although the XIV UI is advanced, cloud architects and providers want to automate storage management functions via CLI commands (or APIs). XIV management tasks can be automated through scripts using these commands. Further, by leveraging SVC, private cloud applications and data can be nondisruptively migrated from legacy storage architectures to XIV.

Availability, Reliability, and Performance

XIV architecture affords extremely high performance. The self-healing architecture and internal data redundancy prevent and minimize downtime, which is a critical SLA factor in cloud deployments. The XIV system can be upgraded in terms of both hardware components and microcode/software without any downtime, thus minimizing performance degradation. The system rebuilds in the event of disk or module failure, and the rebuild times can be extremely fast compared with RAID alternatives. There is almost no performance degradation during rebuilds.

When there is a drive failure, we rebuild within a short period of time ...
If we can minimize the exposure time, then the risk to the application is minimized. It's about an hour of rebuild, which is very good. — Storage Architect, Leading eCommerce Company, United States

We found we would be able to shut down one half of the entire storage system, and the services would still be up and running. We also had high-performance, high-availability storage. — Executive Director, Educational Institution, Germany

Leading private cloud deployments demand millisecond response time for top-tier applications. The overall XIV architecture and the system's even and distributed use of resources, dispersing data across the entire tier, afford impressive price/performance with no hotspots and without the need for tuning.

Industry performance benchmarks, including those from the Storage Performance Council (SPC), found that IBM XIV systems rank high in price/performance, while the International Technology Group (ITG) found that "XIV performance was, by wide margins, reported to be superior to that of existing disk systems."

XIV added SSD modules into the arrays. One of the SLAs asked for sub 6 milliseconds to get us to this gold standard. XIV proved it can achieve the 6 milliseconds in our large MySQL deployment. IBM XIV became a very good storage offering that allows us to do what most MySQL companies can't do. — Storage Architect, Leading eCommerce Company, United States

We have seen big advantages in the performance; half the time we found that things ran 50% faster. — Executive Director, Educational Institution, Germany

The XIV system performance optimization enables superior utilization of system resources and automatic workload distribution across all system hard drives with no tuning. The XIV system's ability to maximize load distribution across all disks for all workloads, coupled with a powerful distributed cache implementation, facilitates high performance that scales linearly with added storage enclosures. In addition, because this high performance is consistent, cloud environments get the same performance during the typical peaks and troughs.

One really big advantage is distribution of hotspots. We don't always know how heavy the concurrent usage might be. Dynamic allocation and monitoring for hotspots on the hard drives, and redistribution to support our workloads, has been a very unique feature. It makes it much easier to plan the information on an SAP system because we do not have to check before how the workload will be for this special SAP system. — Executive Director, Educational Institution, Germany

Low Total Cost of Ownership

As cited previously, cost savings is a large motivator for firms to consider private and hosted cloud deployments. The XIV system offers low cost per TB enhanced by its use of high-density drives. With its standard snapshot services, the XIV system affords optimal capacity utilization through space-efficient differential snapshots. Because XIV delivers consistently high performance independent of capacity, cloud administrators do not need to hoard capacity reserves for the sake of performance. With XIV, 90% capacity utilization can be achieved without noticing performance degradation, whereas traditional storage architectures typically require a 60% utilization safe zone to accommodate space reshuffling and avoid tiering policy violations.

Other TCO features include XIV's efficient power utilization and effective cooling within a small footprint. In addition, the intuitive XIV management schema minimizes training requirements and allows more administrators to use the XIV system more quickly.

The usability of the solution ... it really has a very good user interface, where you can get monitoring information [and] easily create LUNs that you can mount on your OS, and it's very convenient to handle. I would say this is state of the art in terms of the user interface and how you can manage a storage system. — Executive Director, Educational Institution, Germany

Cloud providers as well as private cloud deployments require secure multitenancy; that is, reliable and performance-sustaining use of shared resources, including storage, among tenants. Tenants can include different organizations or companies. Thus, it is essential that storage architectures enable the use of their resources by different groups and workloads as well as support SLA requirements. The XIV storage architecture ensures that virtual machines and datastores are evenly distributed across every drive in the XIV system, allowing changes in performance or application requirements to be managed without manual tuning. The resource

utilization of the IBM XIV storage system provides ideal load balance — without degrading volume utilization and thus optimizing utilization.

IBM SVC in the cloud eliminates storage islands, improves capacity utilization, and streamlines management, which can limit expenses and conserve time. With SONAS in place, the resulting consolidation, centralization, and automation reduce the administration of petabytes of files for financial and time savings. SONAS floor space conservation helps lower capital and operational expenditures.

CHALLENGES AND FUTURE OUTLOOK

Firms are still analyzing how best to leverage cloud services in a private, public, or hybrid fashion, and they are considering what workloads are best suited to each of these deployment models. Firms that have reached an extensive level of virtualization are looking to take the next step in deploying a private cloud. However, this requires integrating different infrastructure components and streamlining, through integrated workflows, the provisioning of new applications and infrastructure. Today, many of these workflows remain standalone and separate.

However, private clouds offer firms the longer-term opportunity to integrate disparate processes and the ability to manage IT infrastructure as an integrated set of services. IBM can offer workshops and advisory services to help an enterprise formulate a cloud strategy. IBM's range of cloud management and automation offerings can integrate separate workflows around cloud services and provisioning. IBM is taking a leadership position in open source cloud initiatives such as OpenStack.

Driving to cloud is a business decision. The next step firms are taking with their private clouds is adding business logic into the infrastructure and bringing business intelligence into the cloud. This presents a strategic opportunity for IBM to help businesses transition private clouds from an operational advantage to a revenue-generating mechanism. Firms can leverage IBM's core competencies in consulting services, business analytics, and Big Data technologies as they build private and hosted clouds of the future.

CONCLUSION

IT organizations and service providers alike are moving to cloud architectures. Cloud objectives of improved IT efficiency, better business responsiveness, and greater economic advantages can be stymied with the wrong storage architecture and integration. IBM's Smarter Storage strategy eliminates storage vulnerabilities that firms can face in private or hosted cloud architectures. IBM's block and file-based scale-out storage architectures form the underpinning of a highly agile, performance-based, and reliable cloud infrastructure.

IBM XIV storage systems are optimized for cloud deployments because of their ease-of-management, reliability, and performance characteristics. The IBM SmartCloud storage family enables the automation of cloud storage services, including provisioning, monitoring, and access. The TSM family optimizes recovery of data and systems in a cloud deployment. IBM's storage portfolio offers integrated cloud offerings and is well

positioned today and into the future to enable enterprises to deploy cloud architectures. Increasingly, firms are selecting XIV, including solutions with SVC and SONAS, as the underlying storage architecture for private cloud and hosted cloud deployments.

This document was developed with IBM funding. Although the document may utilize publicly available material from various vendors, including IBM, it does not necessarily reflect the positions of such vendors on the issues addressed in this document.

Copyright Notice

External Publication of IDC Information and Data — Any IDC information that is to be used in advertising, press releases, or promotional materials requires prior written approval from the appropriate IDC Vice President or Country Manager. A draft of the proposed document should accompany any such request. IDC reserves the right to deny approval of external usage for any reason.

Copyright 2013 IDC. Reproduction without written permission is completely forbidden.