

Increase data availability with the IBM Storage Utility offering

Highlights

- Eliminates waste with pay-per-use billing within your HA environment
 - Provides on-demand system provisioning scaling across the HA environment
 - Requires no data center disruption or buffer capacity expansion
 - Enables smooth capacity expansion within the contract period
 - Moves high up-front costs to steady, business-aligned quarterly charges
-

Lack of architectural resiliency in the data center is exposing enterprises to higher levels of risk and application downtime

Business downtime is costly. Some IT industry analysts estimate that the average per-minute cost to businesses of data center downtime can be nearly eight thousand dollars.¹

To mitigate risk and protect enterprises from costly outages, IT departments must understand and address two key issues that expose applications to downtime.

The first issue is the common confusion of the resiliency capabilities of a five-nines or six-nines *array* with those of a high-availability (HA) *architecture*, and the ability of each to meet expected application service level agreements. Regardless of how resilient any particular array is, it still represents a single point of failure. To ensure HA, redundancy must be incorporated at the compute, network and storage layers.

The second issue is the increased amount of data typically stored in a single array—an increase driven by dramatic improvements in both storage array performance and density. These advances have fueled massive consolidation of storage infrastructure, leading to a reduction of the physical separation of data, and the creation of large failure domains. The result is a broader impact to the business if a failure does occur at the array level.

To achieve the level of service availability application owners demand, IT administrators must recognize the risks to availability and adhere to a set of *core tenets for HA* designed to address resiliency at the architectural level:

- Fully redundant architecture in the compute, network and storage layers
- Multi-path access to storage
- Multiple copies of synchronous data
- Separation of copies on different storage arrays
- Elimination of single points of failure

For applications that require HA, an architecture that follows these core tenets for HA must be employed.

How do I build a redundant storage architecture?

IT departments normally take great care to ensure their hosts and networks are deployed redundantly. Storage, however, is typically purchased and deployed as a single array. Administrators rely on the redundancies within the array itself for protection, but what if the array should go offline for any reason, such as human error or mechanical failure? At a minimum, data access is lost. Worse, the data itself could be compromised.

Achieving high availability in a new environment

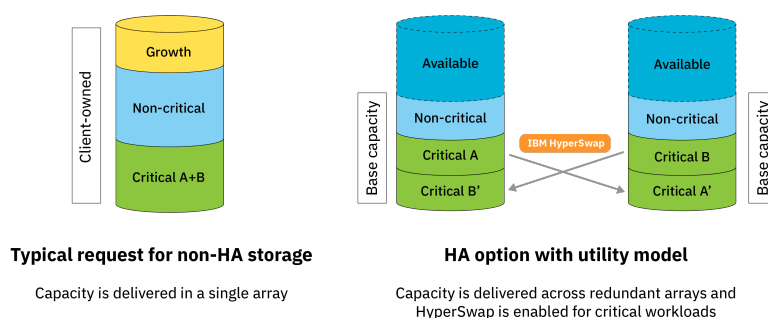


Figure 1: Achieving high availability in a new environment.

Deploying a redundant storage infrastructure and enabling HyperSwap for critical data is a necessary component of a resilient architecture. This configuration prevents the failure of any one array from preventing access to critical data, because a synchronous copy is kept on each array.

Maintaining application availability and containing costs are critical

- For high-availability configurations using HyperSwap and deployed by IBM Lab Services, IBM FlashWatch guarantees zero interruption to data availability during the guarantee period
- Acquire an HA configuration (two systems) with the IBM Storage Utility Offering with a starting rate that is *only 20 percent more* monthly than leasing a single system

IBM Storage Utility advantages:

- Drive capital expenses (CapEx) costs to operating expenses (OpEx); free cash for other purposes
- Pay for what's used, when used
- Always-ready capacity and performance when business dictates
- No disruptions during growth spurts
- Total capacity deployed on day one, no incremental upgrades needed

I already have an array. Is it too late to build a redundant architecture?

To achieve redundancy from existing deployments, the infrastructure for HA, including the appropriate firmware levels, must be configured, and a second array deployed. Next, the critical data will need to be replicated to the second array with HyperSwap enabled.

In this example, a new array is added with the utility model. Adopting a utility model allows for the delivery of capacity on demand. Spreading the data over two arrays provides additional growth capability in the existing array as well as the redundancy created by using a second array.

Achieving high availability in an existing environment

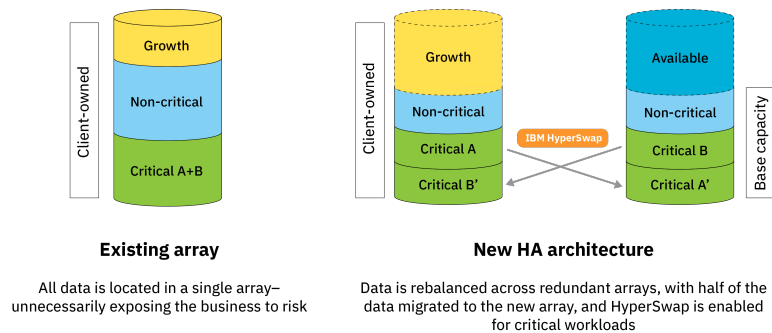


Figure 2: Achieving high availability in an existing environment.

Note that in the case of using an existing array, the client is allowed the High-Availability Guarantee; however, the guarantee period is limited to that of the remaining warranty on the existing system.

I already have multiple arrays. Can I build a redundant architecture without buying a new one?

Existing clients owning multiple arrays need not purchase additional storage to create a replication partner relationship. HA can be achieved by:

- Following the core tenets of HA
- Maintaining the code-level requirements that enable HyperSwap
- Leveraging spare capacity for the synchronous copy of critical data

Clients cannot, however, qualify for the guarantee unless new storage is purchased and HyperSwap is configured by IBM Lab Services.

I would like a local HA solution but still need to store data offsite for compliance

It is possible to achieve local HA and still maintain data off-site for compliance or disaster recovery. Depending on the IBM storage product, the intelligence to move the data either is found natively in the storage firmware or can be scripted to facilitate the remote copy of data.

In the example below, the utility model was used to deploy an HA solution on-site while replicating to an existing array at the disaster-recovery site.

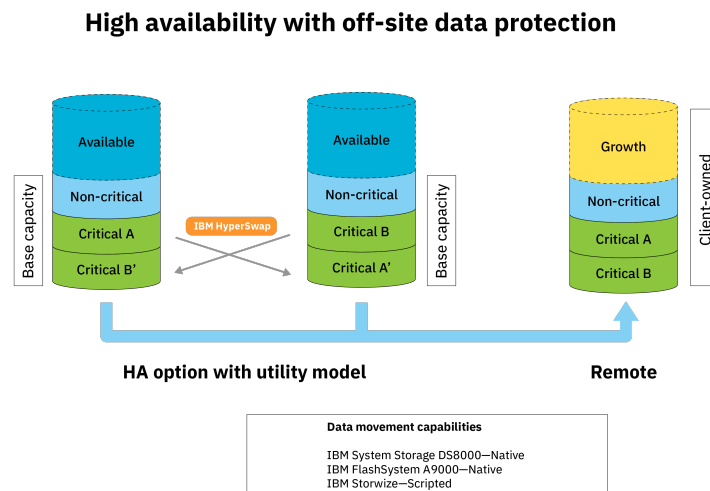


Figure 3: High availability with off-site data protection.

IBM FlashWatch details

The IBM High-Availability Guarantee is designed to eliminate IBM storage system downtime worries for enterprises deploying designated IBM all-flash solutions.

Avoid the costs and risks related to downtime. When you engage IBM Lab Services to deploy an HA configuration using IBM HyperSwap, IBM guarantees zero interruption to data availability during the guarantee period. (See ibm.com/it-infrastructure/storage/guarantee for the terms and conditions of this guarantee.)

The HyperSwap implementation must be performed by IBM Systems Lab Services, and the guarantee is only valid for eligible, newly purchased storage systems.

This new guarantee applies for new purchases of the following IBM Storage solutions:

- IBM Storwize V5030
- IBM Storwize V7000
- IBM FlashSystem V9000
- IBM FlashSystem A9000R
- IBM FlashSystem 9110 and IBM FlashSystem 9150

Do I have to use a utility model?

No, but leveraging the IBM Storage Utility Offering is suggested to address the primary objection some clients have about deploying an HA solution: cost. From a technical perspective, HA is about removing single points of failure from the entire architecture. The IBM Storage Utility Offering provides several key advantages over buying.

Advantages of the IBM Storage Utility Offering

The IBM Storage Utility Offering is a flexible procurement method that aligns your capacity costs to your business initiatives and moves you from a high up-front CapEx purchase to a monthly OpEx-based pay-as-you-grow strategy.

When combining HyperSwap and the High-Availability Guarantee with the IBM Storage Utility Offering, you expand the availability of typical five-nines or six-nines storage arrays to a guaranteed zero-interruption-to-data-availability solution, and—through the IBM Storage Utility Offering—achieve this advantage with a starting rate of only 20 percent higher than a normal single array lease.

In addition, the IBM Storage Utility Offering allows you to:

- Better align your capacity consumption and associated costs with your constantly changing business needs
- Optimize your annual IT spend, typically under constant cost-reduction pressure

- Improve capacity planning

The IBM Storage Utility Offering shines a new light on the storage refresh cycle for clients of all sizes and all industries. The entire three-year projected capacity for the full HA environment is delivered on day one (additional contracted timeframes are available to meet other business needs), ready for rapid on-demand consumption, so your business and IT staff can react instantly as business needs grow or shift. The benefit is that you only pay for what you use.

¹ Andy Lawrence, "Datacenter downtime: when it happens, it (still) really hurts," *451 Research*, March 24, 2016.

Why IBM?

IBM leads by offering a dedicated solution designed specifically for your HA needs. IBM differentiates its solution from the rest through thoughtful, forward-looking execution and by deploying the entire contracted capacity on day one. By comparison, many other vendors prefer to pass risk on to the customer by delivering just enough capacity to meet initial needs, plus a small buffer. This buffer is intended to keep your HA environment's data capacity slightly ahead of your needs. As the buffer fills, however, technicians may be needed to disruptively add incremental capacity to both environments (often separate and isolated from your initial storage pool) to true-up or add-on. This unfortunately often leads to a rebalancing, rebuilding and slowing down of the infrastructure.

For more information

To get started, contact your local IBM representative or IBM Business Partner, who can work with you to define your HA needs while building a plan for your future.

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: ibm.com/financing

© Copyright IBM Corporation 2019.

IBM, the IBM logo, and ibm.com are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at <https://www.ibm.com/legal/us/en/copytrade.shtml>, and select third party trademarks that might be referenced in this document is available at https://www.ibm.com/legal/us/en/copytrade.shtml#section_4.

This document contains information pertaining to the following IBM products which are trademarks and/or registered trademarks of IBM Corporation:
IBM FlashSystem®, IBM Storwize®, IBM HyperSwap®



All statements regarding IBM's future direction and intent are subject to change or withdrawal without notice, and represent goals and objectives only.