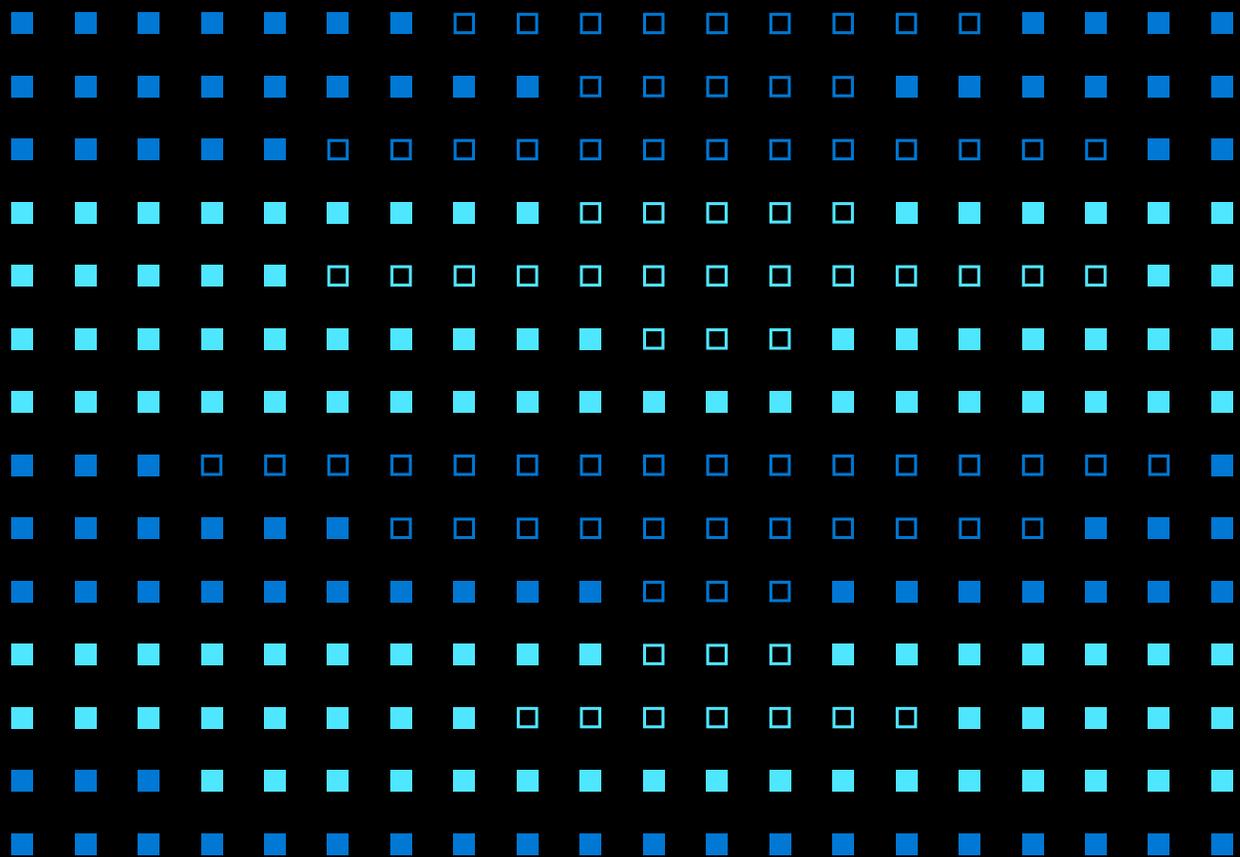
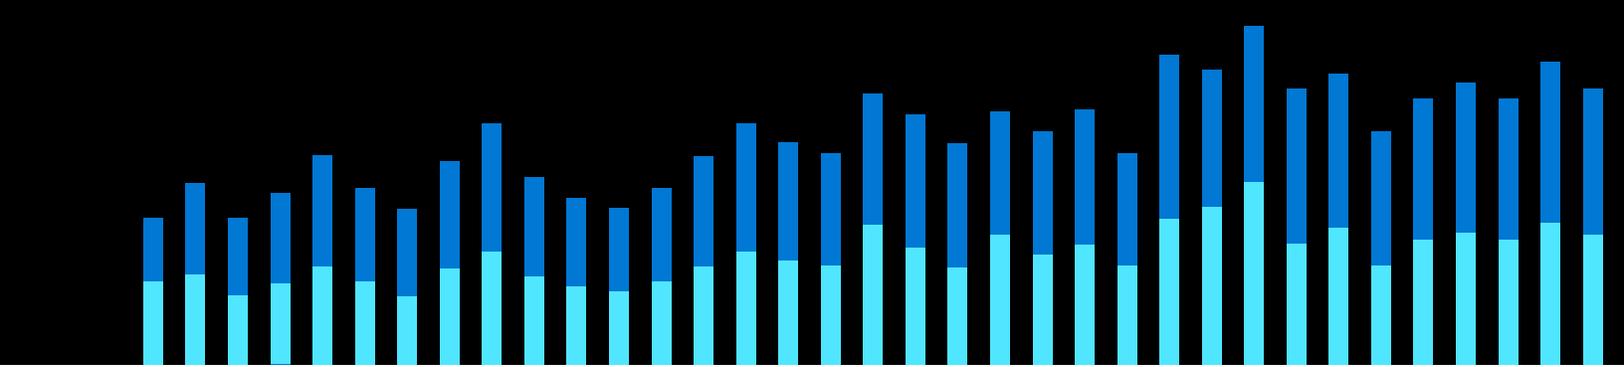


Making smart block storage decisions in the cloud

Guidance to minimize cost, maximize business growth





If you're an enterprise architect or business decision maker, your experience with storage might be limited or primarily focused around on-premises storage. This guide is for you.

Compared to on-premises storage, cloud storage is simpler to manage. It offers granular control and flexibility that enables you to optimize your environments.

For many, however, navigating through the broad array of cloud storage options can be difficult. It's easy to overspend or end up with poorly designed or configured environments.

Learning to make smart storage decisions will enable you to strike the right balance between simplicity, performance, scale, and budget.



This ebook focuses on Azure Disk Storage, a common type of block storage.

Azure application projects running in Azure Virtual Machines require Azure Disk Storage.

The guidance included here is for enterprise architects, business application architects, developers, and decision makers who are familiar with on-premises storage and want to make smarter decisions when they spin up an Azure VM.

How to think about cloud storage

First, let's walk through a few storage basics, which will provide context for the guidance that follows. It helps to understand the multiple storage options available in Azure and other public clouds. It also helps to understand the considerations required to meet performance goals. When you design cloud storage strategically, you can achieve a specific mix of storage attributes to optimize app performance. What follows is a quick overview.



Azure Storage offerings

[Core storage services](#) in Azure divide into three basic categories for different types of workloads. Let's take a look at the Azure services for block, file, and object storage before zooming in to talk about block storage for the remainder of the ebook.

Block storage is commonly used with virtual machines and containers to provide applications the right mix of storage throughput, capacity, and performance. The Azure service that offers block-level storage volumes for Azure VMs is [Azure Disk Storage](#).

File storage enables you to move on-premises file shares to the cloud. Use [Azure Files](#) for lift and shift migrations with zero downtime. Use [Azure NetApp Files](#) to run demanding Linux and Windows file workloads in Azure, including SAP HANA, high-performance computing, and virtual desktop infrastructure.

Azure also offers **object storage**—for massively scalable storage for data objects. In Azure, [Blob storage](#) is optimized to store large volumes of unstructured data in a single infrastructure, which makes it ideal for images and video. [Data lake storage](#) is optimized for use with analytics tools.

Focus first on your application

To get the most from Azure Storage options, start by examining your application's data, which may be structured, unstructured, or a mix of both. You can use any storage mechanism to store either structured or unstructured data, but some are better suited for one type versus the other.

Structured storage is best able to cope with data that follows a predictable format, with all the data following the same "shape." Think of a database of US telephone numbers: Every number will have exactly 10 digits, so a structured store can easily handle them.

Disk storage is commonly used with structured data.

Unstructured data, by comparison, has no guaranteed uniform structure and requires unstructured storage.

Imagine an e-commerce application that stores user comments—a single product might generate thousands of comments, with each comment consisting of some amount of text, links, pictures, and other contents, all with no predictable structure.

Blob or data lake storage is commonly used for unstructured data.



Combining disk storage attributes for your applications

Just as every taste can be described as a combination of salty, sweet, bitter, and sour, in much the same way, every storage mechanism offers varying degrees of four basic attributes: throughput, capacity, latency, and availability.

Your existing on-premises applications, and the storage systems that support them, deliver a specific combination of these attributes. For example, your application may require a certain amount of throughput enabled by an array of physical disks. Because each of those disks is a fixed size, the architect who designed the system for a specific throughput also gave you a specific capacity as a result of the storage design.

Delivering the right combination of attributes works differently in the cloud than on-premises. Let's walk through how this works with disk storage.

 **Throughput** is a measure of how much data can be written or read per unit of time. You may be familiar with the transaction-per-second (tps) metric often used to measure relational database (RDBMS) performance, or the I/O operations per second (IOPS) used to measure physical I/O speeds. IOPS is the number of requests that your application sends to the storage disks in one second. Think of throughput as the product of two factors: how many IOPS a storage system can deliver, and the amount of data moved in each operation.

For example, a storage system that can perform 100 IOPS each of 128KB has the same throughput as one that delivers 50 IOPS each of 256KB. The read or write throughput capabilities of a particular storage method will vary according to how the storage is provisioned by the cloud provider, how you use it (for example, performing one large write is often faster than performing several small writes with the same total size), and a host of other factors.

To choose the optimal disk type, start with a [thorough knowledge of your application's I/O patterns](#) (for example, the average read or write size, and the ratio of read operations to write operations). Different classes of disks will have more or less throughput.



Capacity is how much data you can store. In the physical-storage world, capacity is easy to measure: if you buy a 4TB disk, you can't magically turn it into a 10TB disk. In the cloud world, some storage methods are more scalable (and easier to scale) than others. Knowing how much your storage needs change over time will help you prioritize capacity and [scale](#).



Latency is the delay between a storage request and its response, imposed by the storage system. Depending on what the application does and how it works, you may find that whatever latency the chosen storage subsystem delivers is good enough, or you may find that your application requires a predictable amount of latency. For example, dev/test workloads will not have specific requirements for consistent or predictable latency. Other transactional applications, such as a video encoder or network processing application, are latency sensitive and will have requirements.

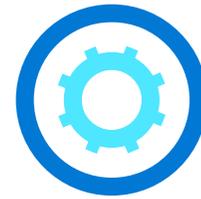


Availability. The fastest storage in the world won't do you any good if data isn't available where and when you need it.

Logical redundancy and geographic redundancy help provide guaranteed reliable access to your data no matter where in the world your application and its users are located. Of course, a line-of-business application used only in one location will have different requirements than a global application scaling to hundreds of thousands of simultaneous users.

Making smart disk storage decisions

Enterprises that seek to transform existing operations to take full advantage of the cloud typically focus first on application designs and workloads. In a world of inexpensive and flexible on-premises storage, most of us haven't had to worry much about the differences in storage behaviors and available capacities—we've had the luxury of just using whatever the existing storage hardware provided.



As you move to the cloud, you now have an opportunity to choose the right storage based on a cost-benefit ratio, instead of simply accepting the default option.

You're the expert on what your application does, its business value, and its technical requirements and constraints. When you use that knowledge to intentionally choose storage, instead of accepting the default option, you can achieve a better result. In the same way that an experienced chef knows how to mix flavors to get the desired results, you can combine the throughput, capacity, latency, and availability delivered by a specific Azure Disk Storage configuration to give your application exactly what it needs.

Your business may dictate that you try to optimize for one or more of these requirements at the same time. For example, requirements for a large-scale online transaction processing engine might prioritize throughput and availability over ease of use and cost, whereas a caching engine might prioritize speed and ease of use for developers over redundancy. Since you're the expert on your application's requirements, you'll want to evaluate each potential storage method against the unique needs of your application.

Driving growth with timely storage decisions

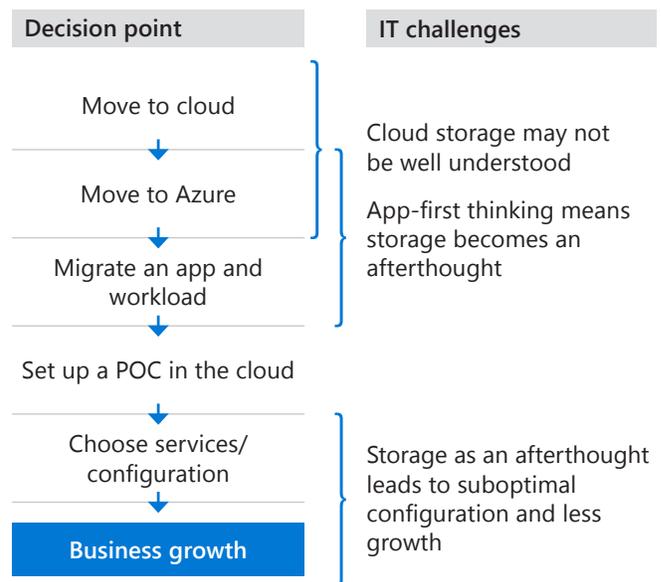


Figure 1: Organizations have opportunities to make cloud storage decisions at a variety of milestones. Focusing on storage upfront leads to smarter storage configuration, which improves app performance and drives business growth.

When lift and shift is not enough

Any time you modify the application's implementation, be ready to re-evaluate the storage you're using. Even a simple "lift and shift" to the cloud means you need to consider throughput, capacity, latency, and availability requirements for your application to see if you need to adjust your storage design.

A line-of-business application written as a single-threaded application ten years ago might have performed acceptably on contemporary storage hardware, and, thanks to Moore's Law, continued to run well on faster single processor cores

coupled with faster disk hardware. But if you want to move that application into a modern multi-core virtual machine (VM), the performance profile of the application may be very different. You'll probably need to re-evaluate the storage you provision on the VM to ensure that the storage and application are well-matched.

To meet performance requirements in a cloud VM, you can optimize your application's performance in multiple ways, from changing the storage configuration to taking advantage of Azure compute.



Azure Disk Storage tools give you a great deal of flexibility to reprovision, rescale, and re-envision your application requirements as your business and its needs evolve. While the range of choices you have available might seem a little overwhelming, [Microsoft and its partners](#) have a wealth of documentation and tooling to help you research deeply, choose well, and update as needed to ensure that your business runs as efficiently as possible on Azure Disk Storage.

Start with the application host

Every time you create an Azure VM, Microsoft will recommend a default disk type and size. The type of VM you choose and the type of disks you attach to it are intertwined: the more powerful the VM, the more work it can do, and the more work it can do, the more demand it may place on the disks you attach to it. The [more throughput and storage the disk can offer](#), the more application demand it can satisfy on demand. Note the use of “may”—whether or not the application actually places that kind of demand on storage depends entirely on what the application does and how it works. Ideally you’ll pick a disk type that matches the characteristics of your application well enough to ensure that your application has enough IOPS, enough storage capacity, enough throughput, and low-enough latency to perform well.

Optimizing for size

Here are some questions to ask as you consider how to optimize your storage design for your application’s size requirements:

- How much storage does your application use now?
- Is the amount of storage you have allocated to your application sufficient for the way the application runs right now?
- Are there features or capabilities you’d like to add that may require additional storage?
- What do you know about the changes in your application’s usage patterns driven by moving to the cloud?
- Could your application be more efficient if you changed the underlying mix of storage? For example, would moving to a smaller number of larger virtual disks give you improved performance for the same amount of storage?

Most of us are familiar with the ongoing process of trading off speed and size in application design and implementation. The good news is that Azure Disk Storage often reduces the need to trade off: You can have enough capacity and enough speed by choosing storage appropriately for your application.

Tale of two apps

Let’s look at the costs of under-provisioning and overprovisioning storage.

If you don’t provision enough storage capacity, in terms of size or throughput, your application won’t perform optimally. As an example, consider two apps running on different Azure VMs. Each application’s speed and scale will be limited by the storage.

If you under-provision App1 to save cost, you risk running into performance and customer satisfaction issues during peak times. I/O bottlenecks will throttle the application’s throughput as application processes wait for disk I/O to complete. This is an age-old problem, dating back to

the first hard disks. Disk storage offers a fixed amount of throughput, and without enough, your applications slow—but you can get more throughput by adjusting the type or number of disks you allocate.

If you overprovision App2, you risk wasting money. Spending more for a committed IOPS rate that your application never reaches because the application is compute-bound doesn’t improve performance. Reallocating that spending to buy more compute capacity might help.

In either scenario, you face unintended consequences, which reinforce the fact that making smart storage decisions is important.



Optimizing for performance

Performance optimization for storage can be very tricky because it requires you to get reliable, accurate data about the application’s performance under realistic conditions. The best way to do this is by measuring your application’s performance using appropriate tools to tell you how hard it’s working and what percentage of its workload is consumed by compute versus storage. Here are a few other points to consider:

- Your application’s actual performance requirements.
- Whether there are peak periods of demand when the application needs extra headroom to handle load
- Level of confidence in data reflecting baseline application performance.
- Whether your application has attributes that would make it faster, or slower with a particular storage type.

- Whether it is possible to isolate or divorce specific data items or types to improve the application’s speed. One common tactic, for example, is to split the files that the application reads from those it writes and put them on different storage types to take advantages of differences in the application’s read/write requirements.
- Whether the VM type you’ve chosen to host the application is sufficiently powerful to fully exercise the storage type you’re using and whether that storage has enough throughput and capacity to allow the application to run full speed. Learn more [about how disk storage](#) works in Azure.

Because you can [change the type of individual disks in Azure](#), it’s possible to tweak your storage configuration as you identify performance hot spots and bottlenecks. Services such as [Azure Monitor](#) help you spot these problems so that you can make adjustments.

Choose Azure Disk Storage type for performance

Primary Considerations:

- Latency
- IOPS
- Throughput

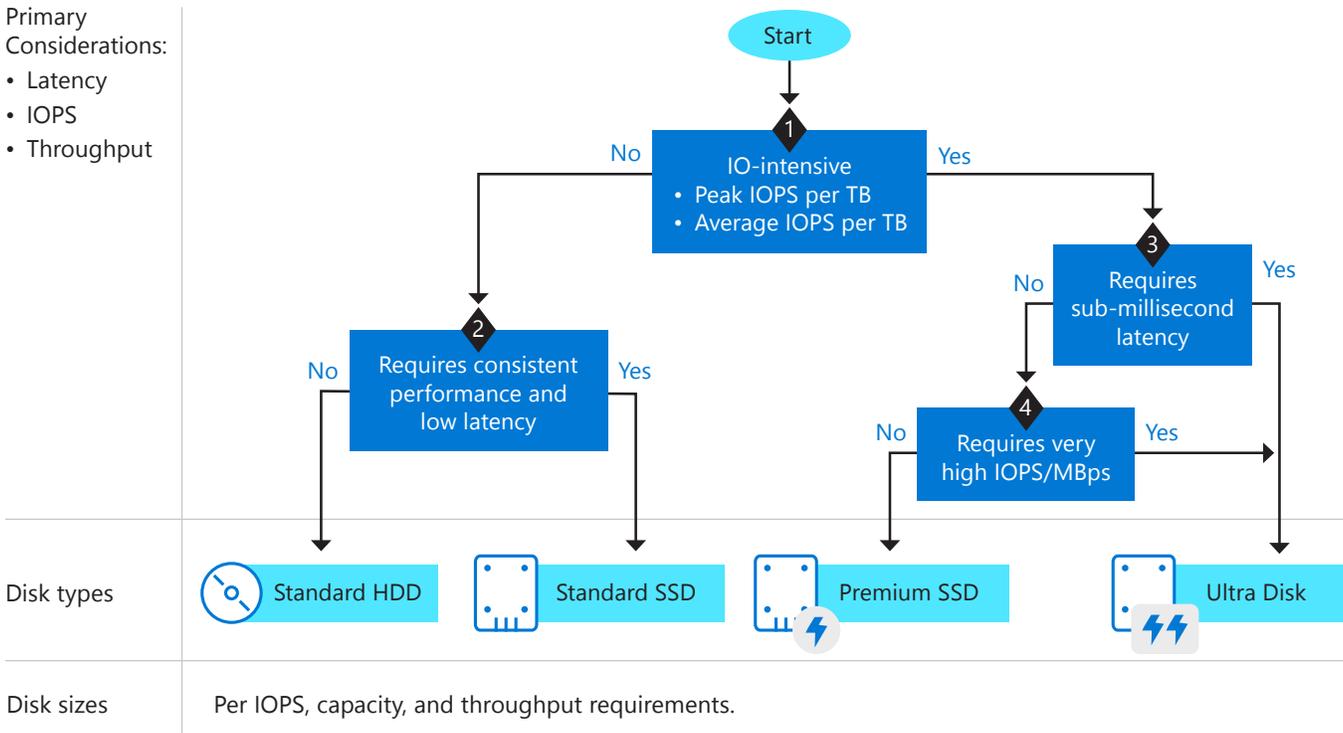


Figure 2. The diagram illustrates one way to determine which type of Azure Disk Storage resources to provision. Understanding the I/O patterns of your applications and its requirements for throughput, latency, and capacity helps you choose disks with the right set of attributes.

Optimizing for availability

Every application has availability requirements. These requirements can change dramatically depending on the user audience, what the application does, and other factors. As with so many other aspects of storage design, you—as the application expert—need to have a clear understanding of what your application availability requirements look like. Here are some questions to ask:

- Does your application have a defined or published explicit service level target that it must meet?
- Are there implicit service level targets based on user or business expectations?

- Does your application's data need to be replicated in a primary region or to a secondary region? Consider the tradeoff between lower costs and high availability and durability when determining the best [redundancy option](#).
- Does your application implement its own high availability strategies? Or is it necessary for you to add availability by distributing critical application storage?
- Does your application currently run with a single instance of any critical data items? You may be able to quickly increase its uptime and resilience just by implementing multiple instances of those data items.



Special consideration for VMs

VMs used to be much simpler to select and configure. Early generations of Azure VM support offered only a few machine and storage types to choose from. Now, of course, there are dozens of VM types from the smallest to the most compute-intensive, and the range of storage options has broadened too. Here are a few considerations when choosing storage for workloads that run on Azure IaaS virtual machines:

1. Understand what you get with storage for Azure VMs. Research [Azure managed disks](#), for example, to understand how they differ from unmanaged disks and on-premises storage.
2. Identify application bottlenecks. Is your application compute-bound or IOPS-bound?
3. Understand [availability options](#) for Azure Virtual Machines.

4. Is the disk storage associated with your VMs the main way the application stores data? For example, an image-processing application that stores incoming images in a blob or data lake may not need to rely on the storage attached to the VMs.
5. How easy is it for you to move application data? If you want to re-engineer your application to take advantage of serverless computing and your data is in conventional file or block storage associated with a VM, you'll probably need extra planning to move it.

You're the expert on your application, so use your knowledge when making storage decisions to choose the optimum storage mix.

Understanding Azure Storage pricing

You may have noticed that there is no section that suggests optimizing for price. Because you can easily see Azure costs upfront, some organizations will focus on price first. True costs, however, are not easy to see.

Calculating a total cost of ownership (TCO) for cloud storage is different than for on-premises storage. With on-premises storage, you simply figure out the capital investment required to buy it, then add the yearly maintenance and service cost. With cloud-based storage, you avoid capital expenditures but TCO is more elusive because of the way storage is priced.

Rather than focusing first on price, make sure the storage delivers the required performance and capability to your users. Only then are you ready to examine the [multiple pricing tiers](#) and cost options Azure offers.

When you buy premium Azure Disk Storage for high-performance workloads, for example, choose the option that best meets your needs for storage size, IOPS, and throughput. What you get for a monthly price is a set amount of capacity with a set throughput capability.

You can also lock in discounted pricing with Azure reservations and reduce your overall VM costs by [combining Azure Disk Storage reservations with Azure Reserved Virtual Machine Instances](#). The reservation discount is applied automatically to the matching disks in the selected reservation scope. [Find more details here](#).

After you've done the work required to get your application moved to a supportable mix of storage technologies in the cloud, you can start comparing the total TCO required for various types of storage and then start optimizing for cost. For example, you might find that provisioning with a specific disk type easily meets your performance targets by a factor of three. That's the right time to consider whether stepping down to a less performant, but less expensive, disk type will still meet your needs or whether you should retain the extra headroom to allow for growth in your application's scale.

Configure and estimate storage costs

Azure offers a range of storage solutions to fit your scale, performance, and budget needs—ranging from a low-cost Archive Storage to a consistent low-latency, high-throughput Ultra Disks offering for I/O-intensive workloads.

Use the [Azure pricing calculator](#) to estimate your monthly costs for Azure services.



A screenshot of the Azure pricing calculator interface. The title is "Your Estimate" with a plus sign icon. Below the title, it says "Your Estimate" followed by three icons: a link, a refresh, and a trash. Underneath, it shows "Storage Accounts" with a dropdown arrow, and "Managed Disks, Premium SSD, P6 Disk Type 5 Disks, Snapshot: 100 GB, Pay as you go" with a trash icon. The costs are listed as "Upfront: \$0.00" and "Monthly: \$64.24". The "Storage Accounts" section has three dropdown menus: "REGION:" set to "East US", "TYPE:" set to "Managed Disks", and "TIER:" set to "Premium SSD". Below that, "Disk size:" is set to "P6: 64 GiB, \$10.207". The "Number of Disks" section shows a text input with "5" and a "x" icon, followed by "\$10.21 Per month".

Other considerations

This ebook includes basic guidance to help you choose disk storage in Azure. There are other features and benefits worth mentioning, but not discussed because they are automatically included when you provision Azure Disk Storage. [Learn more](#) about these features and benefits, which help set Azure apart from other cloud platforms.

- Durability
- Security
- Simplicity and scalability
- Manageability

Next steps

It's time to apply the foundational knowledge contained in this guide to optimize storage to best support your application workloads.

To help you get the performance and capability your users need, keep in mind these steps:

- Carefully investigate and benchmark your application's needs, matching them to the Azure Storage types described earlier.
- Optimize first for capacity, availability, and speed and only then optimize for cost.

Choosing the right storage type for your applications will pay off now and into the future as your applications grow and evolve in response to changing business' needs.

Deep dive into disk storage videos

Build on your foundational knowledge by [learning directly](#) from Azure Disk Storage experts.



[Overview of Azure Disk Storage security features](#)



[Run clustered applications on Azure with shared disks](#)



[Understanding Disk IO and how to leverage disk bursting](#)



Get started on Azure. Start free.

Migrate workloads to the cloud with your Azure free account.

[Find an Azure partner for migration, backup, and disaster recovery](#)

[Keep up with the latest Azure Disk Storage updates and announcements](#)

[Explore Azure Disk Storage learning paths and modules](#)

[Learn more about core Azure Storage services](#)