Private Clouds: Selecting the Right Hardware for a Scalable Virtual Infrastructure

Greg Shields

sponsored by

Realtime publishers

Copyright Statement

Chapter 2: Inside the Black Box: A Private Cloud Reference Architecture

This guide's first chapter is intentionally written to be visionary. In it, you learned about the many fantastic capabilities a well-designed private cloud infrastructure provides. Those capabilities center around improvements to

- availability for individual IT services
- flexibility in managing services, as well as deploying new services
- scalability when physical resources run out
- hardware resource optimization, to ensure that you're getting the most out of your investment
- resiliency to protect against large-scale incidents
- globalization capacity, enabling the IT infrastructure to be distributed wherever it is needed

That's an exceptionally strong set of benefits. Particularly so, considering that a private cloud at its core is little more than a virtualization technology, some really good management tools, the right set of hardware, and *business process integration*. You'll learn more about that last component throughout this guide. But recognize that for the enlightened, a private cloud's hardware infrastructure can be one of its most-compelling characteristics.

Why? Any organization can incorporate a minimal virtualization environment by obtaining a server and a hypervisor license. Yet a private cloud is so much more than simple virtualization. In fact, a private cloud *is more about the management of virtualization* than any virtualization technology itself. Private clouds embody the mechanisms to fold virtualization's technology into hardware as well as business processes and drivers. This quote from the previous chapter impresses this point:

> Although virtual machines are the mechanism in which IT services are provided, the private cloud infrastructure is the platform that enables those virtual machines to be created and managed at the speed of business.

As you can imagine, evolving a simple virtualization environment to that fully-realized private cloud requires a few steps. Chapter 1's conversation was framed around recognizing a private cloud's operational benefits. This chapter focuses on the architectural elements you need to construct a private cloud. This conversation may be technically complex, but it is absolutely important for understanding why virtual hardware can be as important as virtualization itself.

What's in the Black Box?

Users in an organization look to a black box of IT services for the applications and data those users require (for a refresher, take another look at Figure 1.2 in Chapter 1). Both the delivery infrastructure and the users connect to the Internet for further services. Those services can be traditional Web pages or offloaded Web services that perform some function of business.

Figure 2.1 peers into that black box to expose the details of its individual components. At the same time, the illustration shows an augmented box with a set of supporting capabilities needed by pretty much every business: data backup, data archival, disaster operations (which are essentially a mirror of your production environment), and overall environment management.

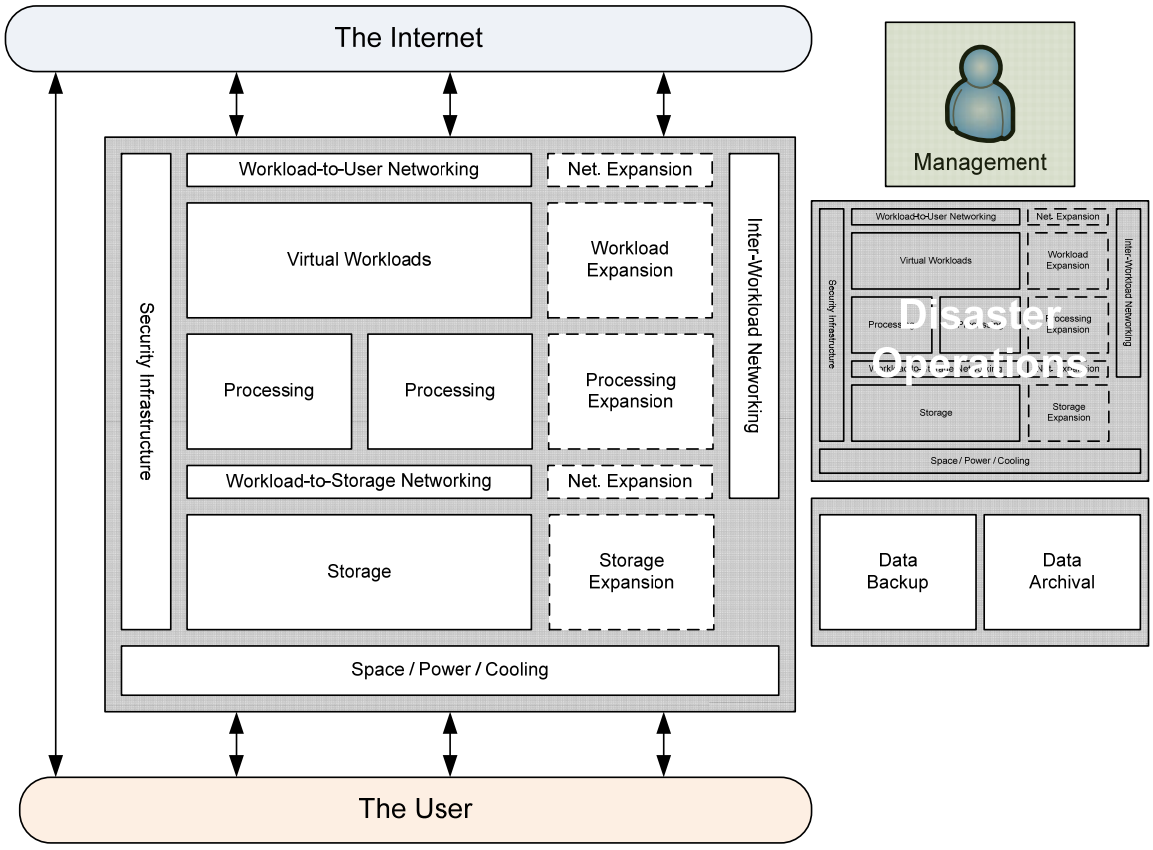

Figure 2.1: Peering into the IT services delivery infrastructure black box.

Let's take a look through each of Figure 2.1's elements in turn. Creating a private cloud obviously requires a hardware investment in each. Recognizing where and how to make that investment will help you make smart decisions about where to spend your money.

Private Cloud Processing

The individual server is probably the most-recognized component in the modern data center. That recognition makes it a perfect starting point for any private cloud construction story.

A private cloud's processing layer is comprised of a set of individual hardware servers (see Figure 2.2). Each server can be similar to the hardware that you already use in your data center today or it can be a component of a modular infrastructure. This flexibility means that in the transition from traditional IT operations to private cloud operations, you may potentially reuse much of your existing server infrastructure.

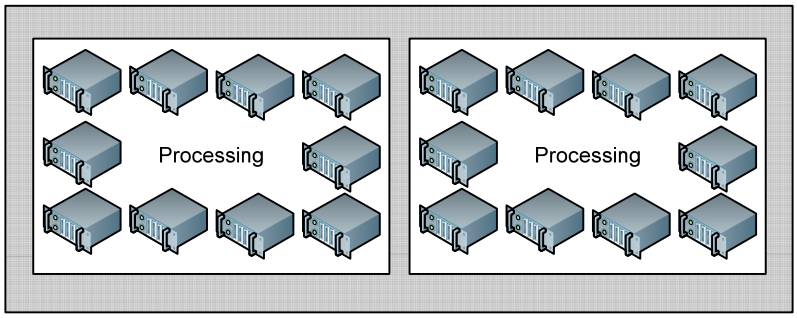

Figure 2.2: The private cloud processing layer.

Different here is the fact that each server does not run its own service workload. Instead, each server operates as an individual processing node in what is sometimes called a *resource pool*. In such a pool, a discrete level of processing capacity is available and can be assigned to virtual workloads.

Resource pool processing capacity is by definition additive. One server with 16 processors at a certain speed in gigahertz adds one times sixteen times that processing speed. Adding a second similarly-sized server into the pool doubles its capacity. Doubling the number of processors or speed of those processors further augments the pool.

The same holds true with onboard memory. IT workloads require both processing power and memory in order to accomplish their assigned tasks. Some require more than others, while some require far less. Each server in a private cloud's processing layer also contributes a discrete quantity of memory, which is also additive in nature. As with processing, add more servers or more memory to servers, and more memory becomes available to assign to virtual workloads.

The result of this resource pooling creates an abstraction over the top of each server (see Figure 2.3). IT workloads are no longer physically tied to individual server hardware (virtualization accomplishes this). More importantly, they are no longer logically tied to a specific set of resources (the private cloud accomplishes this). Further, IT workloads whose resource demands change can be assigned a new quantity either automatically or via administrator action inside the private cloud's management infrastructure.

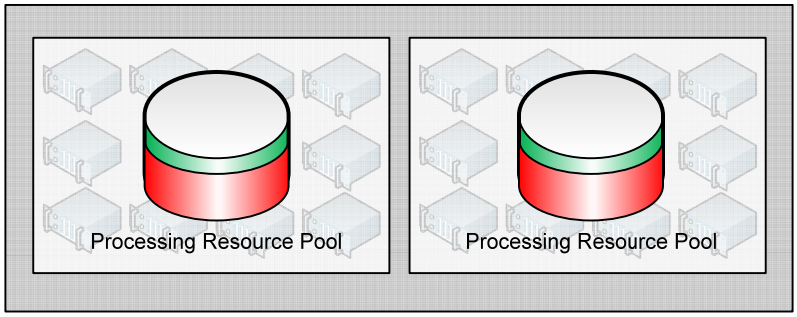

Figure 2.3: Abstracting physical processing to create a processing resource pool.

Particularly important to recognize in Figure 2.3 is that you aren't limited to creating just a single resource pool for the entire enterprise. You'll see that two are in fact present in the figure, which represents the fact that any number of resource pool subdivisions can also be created. These subdivided pools enable the ability to assign processing capacity to different teams, organizations, or business units *based not on technology but business rules*.

For example, consider an enterprise with two business units where the first business unit is more successful than the second. The first business unit may have a higher budget. That business unit may also have greater IT resource needs. In this situation, the business unit's size and success can drive the level of resources it requires—and, thus, how each resource pool is sized.

Resource pools can be also be defined by investment. Should the first business unit invest twice the money than the second into the environment, the first business unit can be given twice the resources. If a third business unit wants to participate at some point in the future, that unit can be assigned a proportionate level of resources commensurate with their investment.

> **Note**
>
> In a private cloud, all virtual machine resource assignments always start from a resource pool. This pooling provides the flexibility to re-assign resources as necessary to meet your business demands. Need to spin up a new service quickly? Just assign them from the pool. As a result, w*ith flexible resource pooling, IT and its technology are no longer a drag on business agility.*

Private Cloud Storage

Data center storage is sometimes considered one of virtualization's most-expensive costs. Important to recognize is that those costs relate both to the storage and the management of that storage. Left unmanaged, those costs can grow to become a major drain on budgets. This problem isn't new. Until the advent of Storage Area Networks (SANs), storage requirements for individual servers had grown unwieldy and wasteful: unwieldy from the perspective that IT needed to manage storage across individual servers; wasteful from the recognition that discretely sizing assigned storage to needed storage simply wasn't possible.

SANs in many ways changed the paradigm for data center storage assignment. Using a SAN, storage became aggregated into a central structure that could be provisioned to individual servers. Being centralized, storage could be discretely assigned to servers as needed. It could also be expanded (or, on occasion, shrunk) as needed.

Private clouds wrap that storage centralization mindset into the resource pool concept discussed earlier (see Figure 2.4). A private cloud resource pool is comprised of processing and memory. *It is also comprised of storage*. Storage inside the pool can be assigned to needy virtual machines as they require it, but not until then. This setup reduces wasted storage, enabling an IT organization to make more cost-effective purchasing decisions. It also ensures that storage is best optimized to the needs of IT workloads, and not necessarily its technology underpinnings.

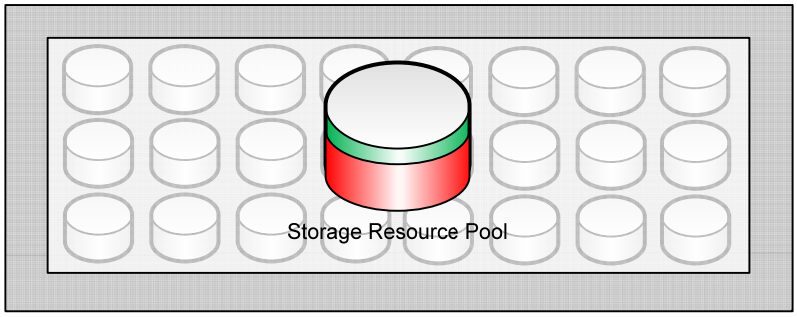

Figure 2.4: A storage resource pool.

An even more important facet of the private cloud approach is in its abstraction of individual SANs and/or SAN nodes themselves. In effect, a private cloud in combination with the right hardware can connect otherwise-separated SAN hardware to become a single unit of administration. Figure 2.5 illustrates a representation of how four individual SAN nodes can be abstracted, aggregated, and ultimately presented to the private cloud environment for provisioning.

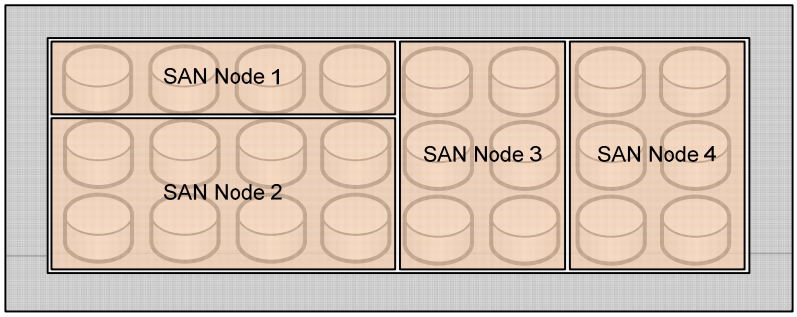

Figure 2.5: Abstracting storage and SANs within a private cloud.

This capability is particularly compelling when considered in light of some early SAN architectures and their intrinsic expansion limitations. In those early architectures, you might be able to expand your SAN but only to a point. Once your SAN reached its upper limit on capacity, your remaining options grew limited. Too often, your only choice was to purchase a new SAN frame and its connecting infrastructure—an expensive capital investment.

A private cloud in combination with modern SAN hardware—particularly SAN hardware with a node-oriented architecture—gets around these upper limits through node interconnection. As with servers, each individual storage node contributes a quantity of storage. Multiple nodes can then be connected when expansion is needed. *Thus, with node-oriented storage, you can grow your SAN whenever you want and to any level you want.*

> **Note**
> I'll talk more about this node-oriented architecture, both with SANs and processing nodes, in Chapter 3.

Private Cloud Networking

As I have said before, virtualization reduces complexity in many ways while adding complexity in others. That statement requires a bit of an explanation, although you might already agree with its premise. Operating a virtualized IT environment tends to be far less complex than its all-physical counterpart. The day-to-day steps are much simpler. Resources are available for assignment to workloads. Virtual machines can be created and destroyed with relative ease. These tasks are trivial in comparison with needing to unbox new equipment with every new server.

Yet although operating a virtualized environment can be of lesser complexity, creating one can often involve far more complication. One area where this has been the case is in the fabric of interconnections between a virtual machine and its needed storage. Figure 2.6 shows a representation of how those interconnections can relate.

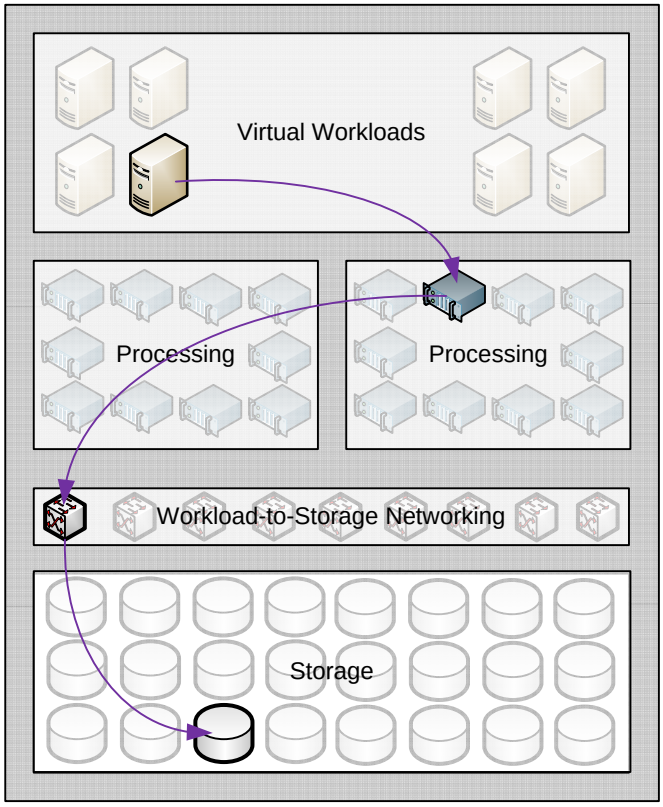

Figure 2.6: Connecting a virtual workload to its storage in a highly-dynamic private cloud environment.

You can see in this image how a virtual machine has a virtual networking interface. That interface leverages the physical adapter inside the server where it is currently running. The physical adapter in turn interacts with networking equipment that connects physical servers to storage. Logic inside the storage array must then securely complete the connection to the correct SAN volume.

If this picture exudes complexity, then add to it the fact that *virtual machines are constantly on the move*. A private cloud constantly shifts around workloads in order to balance loads or prepare for impending host downtime. These dynamics mean that a virtual machine will be using physical adapters from many hosts throughout its life cycle. Storage itself is also highly-dynamic, with storage virtualization technologies shifting around bytes for load balancing and other purposes. Nothing sits anywhere for long.

Complicating this discussion is the recognition that storage traffic *is only one type of networking that must be managed*. Figure 2.7 shows yet another setup with its representation of virtual workload-to-user networking (often called "production networking"). Best practices in virtualization demand that storage traffic be segregated into separate logical—and often physical—connections in order to preserve performance. This can mean lots of extra connections to manage that weren't there in the old physical infrastructure.

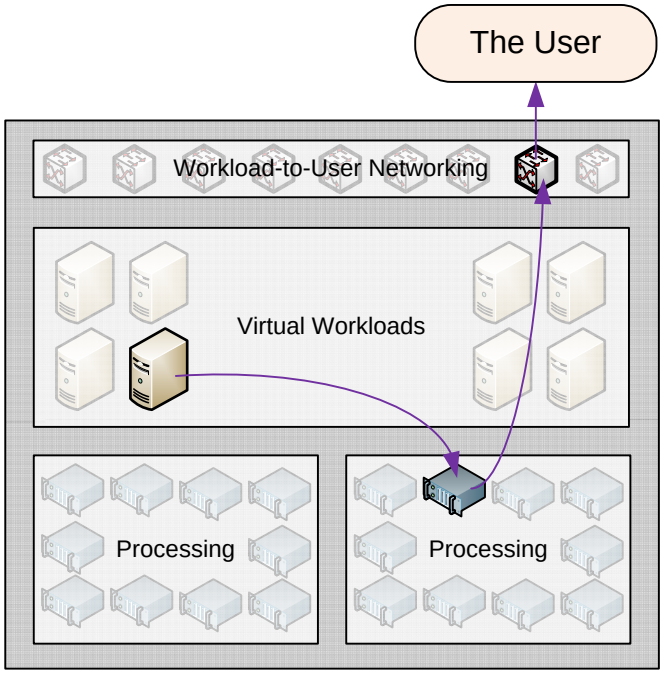

Figure 2.7: Workload-to-user networking.

As a result, your private cloud environment will be responsible for handling at least two—*and sometimes more*—completely isolated network paths. Although a discussion on management is best left for later chapters, your private cloud infrastructure should include the proper hardware and management toolsets that can prepare, monitor, and react to such dynamic situations.

Special Case: Inter-Workload Networking

Think this is getting complicated? This story on networking still isn't complete. Figure 2.1 also includes a third type of networking labeled *inter-workload networking*. This is a special type of networking that occurs between virtual machines in certain circumstances.

Most workload-to-workload connections occur through the same network path as workload-to-user networking. Network traffic leaves one virtual machine, hops around your production network, and eventually makes its way to another virtual machine. However, there are two important examples when this does not occur. In these cases, special consideration is necessary.

The first is when two virtual machines are collocated on the same host. In many virtualization architectures, this coexistence of workloads forces them to communicate across the host's bus rather than over the regular network. *This can be a good thing.* Forcing communication over the host bus can be a boon to network performance as the system bus is generally much faster than the network stack. However, routing traffic in this way *can eliminate some of the key network security and access control list protections* that are put into place on the regular network.

Technologies exist today that reach into server bus communications to extend those protections. Yet these technologies are often an added component that requires extra cost as well as extra management. Your private cloud, particularly in situations where virtual machine-to-virtual machine communication must be highly controlled, should consider the use of these technologies to protect that traffic.

There's a second situation where inter-workload networking needs special consideration. This is when workloads have been configured to communicate with each other over isolated networks that exist inside servers only. These internal networks won't necessarily be part of your production network.

Understanding these inside-the-server connections works best by way of example. Consider the situation where two virtual machines need a direct path to each other that doesn't route across the production network. Perhaps this direct path is needed because of special firewalling restrictions. Maybe the two virtual machines need not communicate over the production network. The same special technologies mentioned earlier also provide an assist here, enabling isolated networks to gain the same protection as regular or system bus networks.

These technologies should be well-known by your virtualization engineers but may not be known by your network personnel. The introduction of virtualization also introduces virtual networking, which has the tendency to relocate IT networking responsibility away from IT networking experts (and to the virtualization administrators). The special technologies described in this section return that responsibility back to your teams that live, eat, and breathe networking.

Private Cloud Space, Power, and Cooling

Space, power, and cooling (SPC) are particularly interesting topics in the private cloud conversation, due in part to virtualization's early hard-cost benefit promises. You know this story: By consolidating multiple virtual workloads atop a smaller number of physical servers, you'll need fewer pieces of hardware in operation. Fewer pieces of hardware mean less operational cost (which impacts power and cooling) as well as less capital cost (which impacts space and the cost to purchase the servers themselves).

The reality of business operations, however, introduces the fact that *not every business will actually realize SPC benefits from virtualization in the long term*. This statement probably violates the established dogma when lined up with many vendors' virtualization ROI calculations. But its assertion can often be broken down into a very few important points.

First, as has been said before, a private cloud is as much a business process entity as a technology entity. This infers that an un-optimized private cloud can impact the data center cost model as much as the un-optimized traditional data center. Where non-optimization truly kills the private cloud cost model is in an area I'll call *virtual machine life cycle governance*. In plainer terms, it means making smart decisions about when to create—and to not create—new virtual machines as well as when to decommission those you don't need any more.

You'll often hear the rhetorical question, *"When something is easy, what do we do?"* Its answer is, *"We do it."* That behavior becomes true when a virtual environment or private cloud is implemented without procedural governance. This problem is thus: W*hen creating new virtual machines becomes ridiculously easy, organizations begin creating lots of virtual machines*. This aggregation of those new virtual machines (the industry term is *VM sprawl*) can grow exponentially to the point where more hardware is required to support the new and faster rate of workload growth than was required back when IT kept its operations physical and slow.

SPC costs have a direct relationship to virtual machine concurrency. The more virtual machines you need running at the same time, the more hardware you'll need to run them. Complicating this fact is the recognition that a build-it-yourself private cloud adds its own un-optimizations. Namely, *without pre-engineered hardware and known specifications, it can be very difficult to build the right amount of SPC capacity while planning for growth.* Making any mistakes in that SPC build-out specification can be catastrophic to budgets.

> **Note**
> Chapters 3 and 4 go into further detail about this optimization, with Chapter 3 discussing it from a technology standpoint, and Chapter 4 analyzing its impact on costs and realized benefits.

Private Cloud Backup, Archival, and Disaster Operations

Another important topic relates to the supporting capabilities that pretty much every private cloud owner wants: data and server backups, data archival, and disaster operations. These elements are important for the environment, as they provide critical data preservation in the case of data loss as well as whole-environment preservation in the case of a large-scale incident.

You don't need to know that you need backups. But what you might not recognize is the high degree of interrelation between backups, archival, and disaster operations that virtualization's technologies enable. The purple arrows in Figure 2.8 show its compelling points. Those arrows illustrate the data flow between different elements in and out of the private cloud. Let's look at each individually.

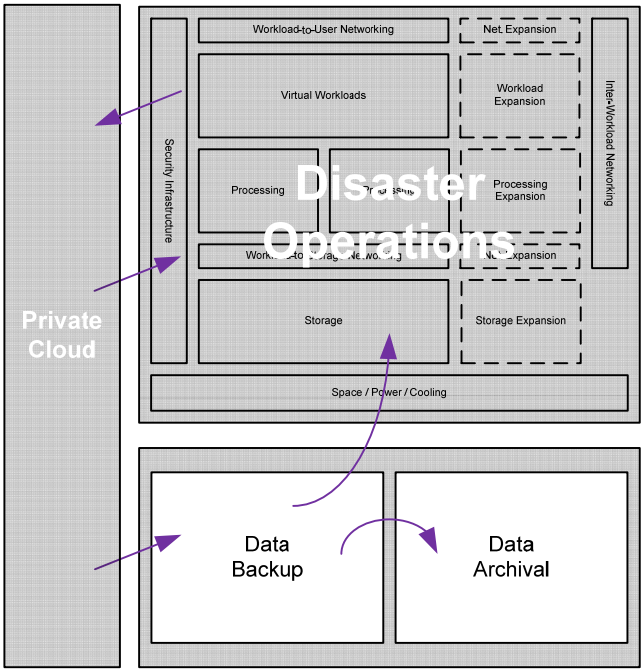

Figure 2.8: Backup, archival, and disaster operations in private clouds are highly interrelated.

Private Cloud-to-Backup

Your environment obviously needs backups. The problem is that without the right technology in place, *virtualization can actually complicate backups while it at the same time improving them*. Part of this dichotomy sources from a phenomenon some call *double backups.* Its story goes a bit like this: In a virtual environment, it is possible to install a backup agent either on the virtual host or inside the virtual machine. Agents installed on the virtual host back up the virtual machine's disk file as a single entity. Agents installed in the virtual machine behave much like agents in an all-physical environment, backing up all its little individual files.

Having options is great, but the dichotomy here also poses a restorability problem. Agents on the host back up a virtual machine's disk file all at once, enabling that whole computer's quick restore should it crash. Modern backup solutions can also peer into the disk file if you need to recover individual files. However, doing so often requires a full restore of the entire virtual machine first. That process takes time, sometimes a lot of time.

Conversely, an agent in the virtual machine sends that computer's individual files to tape one by one. This makes individual file restorability very fast. Unfortunately, its position inside the virtual machine means that it cannot restore an entire server by restoring one file. *Using classic backup technology, you can't have both at the same time.*

Important in that sentence are the words "classic backup technology." Techniques have evolved that enable disk-to-disk (as opposed to disk-to-tape) backups, which aggregate the in-the-virtual-machine with the on-the-virtual-host approach to realize dramatically speedier restore times for all kinds of restores. Disk-to-disk backup technologies also enjoy the random-read/random-write nature of disk-based storage over and above the linear-read/linear-write approach required by tape. Thus, backups can be accessed much quicker, restores complete with greater speed, and backup data can be compressed and deduplicated as it makes its way to disk. The result is a much lower cost of ownership for backups along with an improved time to restore.

A fundamental tenet of private cloud computing is constant workload availability, so your environment might want to look towards newer backup techniques in order to achieve the Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs) you require.

Backup-to-Archival

Even the very best of these advanced backup techniques still requires getting some data to alternate media such as tape or cloud-based storage. Every organization has a need for offsite data rotation for protection, compliance fulfillment, and peace of mind. The backup-to-archival purple arrow suggests separating private cloud backups from the later archival of important data. This data you might never need for disaster operations, but you might need it to satisfy an auditor or compliance officer at some point down the road.

Private Cloud-to-Disaster Operations

Modern backup techniques indeed present the ability to recover a workload with ease. Yet they have another not-well-understood benefit that makes true disaster recovery an affordable capability. Pairing modern virtual backups with a replication technology creates the backbone for a fully-realized disaster operations framework.

You'll assuredly need hardware, appropriate networking and storage, and the necessary management toolsets in that alternate location. But a replication solution that is paired with a virtual backup solution goes far into protecting your private cloud from the really big incidents.

Backup-to-Disaster Operations

Let's assume then that you've laid into place the concepts from this section's first three purple arrows. Most of your business' critical data will now be replicated over to disaster operations at the moment you need to fail over. However, there is the chance that some data might not make it. In this case, you'll need technologies that manifest this fourth arrow, which transfers critical data from backups to disaster operations.

This data flow is important because your backup environment may be a completely separate entity than your disaster operations environment. You might be leveraging a cloud-based solution rather than (or in addition to) disk or tape for backups. That cloud-based solution might be elsewhere on the Internet. As a result, ensuring that a connection exists between your backup environment and your disaster operations environment is an important failsafe measure.

Disaster Operations-to-Private Cloud

Lastly, and this goes without saying, any replication solution with offsite disaster operations must also include the capability to trivially get you back to production once the incident is over. The key word here is "trivially." The right solution won't make you hesitate to fail over your operations when the need arises.

> **Note**
>
> Chapters 3 and 4 will talk more about integrated management toolsets across every part of the private cloud stack to make these kinds of decisions very easy.

Scaling the Private Cloud

Chapter 1 has already discussed the high-level promise associated with private cloud scalability. There, you learned that scaling the cloud is an action that should require little more than connecting additional hardware and enabling it for use.

Figure 2.9 shows a representation of this scalability as an exploded view from Figure 2.1. There, you can see that scalability of the private cloud itself occurs through the addition of individual components. Need more processing? Just add more processing. Need expanded networking? Tie in more network hardware.

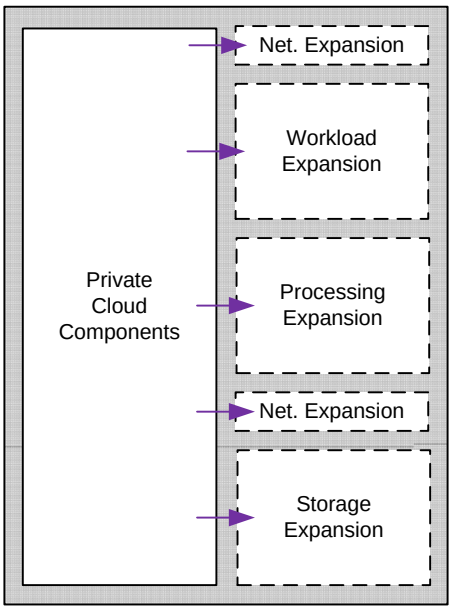

Figure 2.9: Scaling the private cloud's components.

This chapter has constructed its series of figures purposely to illustrate exactly how this component addition can occur. Take a look back through each of the figures for processing, storage, networking, and so on, and you'll see that the private cloud scalability conversation is baked right into its discussion on private cloud construction.

In essence, *you scale a private cloud in the very same way that you construct one, by adding additional components*. The management toolsets and virtualization architecture that is the foundation of a private cloud are the enablers for this trivial technology insertion. You and your virtual machines just enjoy the additional resources.

Private Clouds In-a-Box

This chapter has attempted to define the construction of a private cloud that will support the lofty promises outlined in Chapter 1. In this chapter, you've learned how the right combination of hardware and software in addition to a healthy insertion of business process is what really defines a private cloud's reference architecture.

But throughout this conversation, you should be saying to yourself, "This seems unnecessarily challenging." That's because it is. Correctly connecting each of a private cloud's individual pieces requires expertise and experience that you might not necessarily have in your organization. Correctly building it to your capacity requirements is even more difficult.

It is exactly this gap in experience and expertise where today's hardware vendors are bringing value. Those vendors see the complexity in constructing a private cloud out of whole cloth. They also see the opportunity to share their hardware development experience towards making the private cloud construction experience more like the server purchase experience. Essentially, today's hardware vendors are moving towards a kind of private cloud "in-a-box."

Their in-a-box approach represents a newer, less risky, more predictable, and fully modular approach to private cloud construction. You've already seen that meme teased in both this and the previous chapter. You'll learn more about how exactly to build your private cloud using this in-a-box approach in the next chapter.