

Chapter 1

Industry Trends and Perspectives: From Issues and Challenges to Opportunities

There is no such thing as a data or information recession.

– Greg Schulz

In This Chapter

- Moving beyond the hype
- Navigating the maze of cloud and virtualization stories
- The business demands of IT and data storage
- IT issues and challenges involving data storage
- The business benefit of cloud and virtual data storage networking
- Opportunities to address data storage issues and challenges
- The role of virtualization, cloud, and storage networking
- Maximizing IT resources without compromising quality of service
- What defines a public and private cloud service, product, solution, or paradigm
- The importance of information access, data consistency, and availability

2 Cloud and Virtual Data Storage Networking

This chapter looks at the big picture of business issues and demand drivers that set up the need for cloud, virtualization, and data storage networking. Key themes and buzzwords covered include cloud computing, cloud storage, public and private clouds, information factories, virtualization, business issues or challenges, barriers to productivity, technology tools and techniques, along with best practices. Additional themes and topics discussed include enabling agility, flexibility, scalability, resiliency, multitenancy, elasticity, managed service providers (MSPs), converged networks, Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and IT optimization.

1.1. Getting Started

You probably didn't wake up this morning thinking, "I need to have someone buy or implement a cloud, virtualization, or storage networking solution." Granted, if you are a vendor or a consultant, that may be your job (assessing, designing, selling, installing, or supporting data storage, networks, virtualization, or clouds). However, if you are not a vendor, consultant, analyst, or journalist, but rather someone responsible for information technology (IT)-related solutions for your organization, typically the need to buy and deploy a new solution is tied to solving some business issue or capitalizing on an opportunity.

A common challenge in many organizations is exploding data growth along with associated management tasks and constraints, including budgets, staffing, time, physical facilities, floor space, and power and cooling. Before going further into why you need or do not need a cloud, virtualization, or a storage network, let's take a step back and look at what is driving data growth and the consequent need to manage it more effectively.

1.2. The Importance of Data and Storage

We live in an information-centric world. As a society, we have a growing reliance on creating and consuming data (Figure 1.1), which must be available when and where it is needed. Data and related information services are enabled or provided via information technology services combining applications, facilities, networks, servers, storage hardware, and software resources.

More data can be stored in the same or smaller physical footprint than in the past, thus requiring less power and cooling per gigabyte (GB), terabyte (TB), petabyte (PB), or exabyte (EB). However, data growth rates necessary to sustain business activity, enhance IT service delivery, and enable new applications are placing continued demands requiring more processing, network, or input/output (I/O) bandwidth and data storage capacity.

As a result of this increasing reliance on information, both for home and personal use along with business and professional needs, more data is being generated, processed, moved, stored, and retained in multiple copies for longer periods of time. The

net result is that IT organizations of all sizes are faced with having to do more with what they have (sometimes with less), including maximizing available IT resources while overcoming common footprint constraints (available power, cooling, floor space, server, storage and networking resources, management, budgets, and IT staffing).




Figure 1.1 IT and data storage demand drivers.

1.2.1. The Business End of IT Data Storage Impact

Just as we live in an information-centric society which extends from home to the office, from the small office/home office (SOHO) to the remote office/branch office (ROBO), small/medium-size business (SMB), small/medium enterprise (SME), to ultra-large organizations or enterprises, there is another common theme, and that is economics. Economics are a constant focus, whether it is costs or expense, profits and margins, return on investment (ROI), total cost of ownership, or some other business specific measurement.

On the one hand, there is a need or reliance on having more information; on the other, there are the constants of economics, cause and effect, and supply and demand. You need or want information, but there is a cost to supporting or managing it. Yet information can also directly or indirectly drive profits, so a balancing act is necessary. Thus, to support or sustain economic (business) growth or manage the data necessary to maintain daily activity, there are associated costs (hardware, software, people, facilities, power, etc.) that need to be managed.

Innovation is doing more with what you have: supporting growth and enhancement of services without negatively impacting service-level objectives (SLOs), including quality of services, while reducing per-unit cost for service delivery (as shown in Figure 1.2). The trick is to find the balance among boosting productivity, reducing costs, and maintaining or enhancing customer service delivery.

Figure 1.2 sums up the balancing act of maximizing use of available IT resources while supporting growing business demands in a cost-effective manner. IT resources include people, processes or best practices, time, budgets, physical facilities, power,

4 Cloud and Virtual Data Storage Networking

cooling, floor space, server, storage and networking hardware, along with software and services. All too often, the approach has been to drive cost down by increasing utilization at the expense of quality of service (QoS) and SLOs. An example is leveraging consolidation or migration to a cloud service based on a lower-cost model that trades QoS and SLO for price.

Another variation is to boost QoS and SLOs along with performance to meet demand at the expense of cost or less effectively utilized resources. In other words, it's relatively easy to improve in one area while causing issues or aggravation in another. Innovation occurs when all three categories shown in Figure 1.2 are positively impacted.

Figure 1.2 identifies constraints or barriers to cost-effective service delivery while maintaining or enhancing the service delivery experience including QoS and SLOs. Cloud, virtualization, and data storage networking are tools and techniques that, combined with best practices, can be used to enable innovation and meet the objectives of Figure 1.2.

Clouds, virtualization, and data storage networks can be used to enable cost reduction and stretching of resources by supporting consolidation initiatives. However, these same tools and techniques can also be used for enabling agility, flexibility, and enhanced services that can improve both top- and bottom-line business metrics. For some environments or applications the focus may be on cost reduction while supporting little to no growth, while for others it may mean working with the same or a slightly increased budget while supporting business demand and SLOs. In some organizations this also means reducing costs or stretching available budgets and resources to do more with what they have.

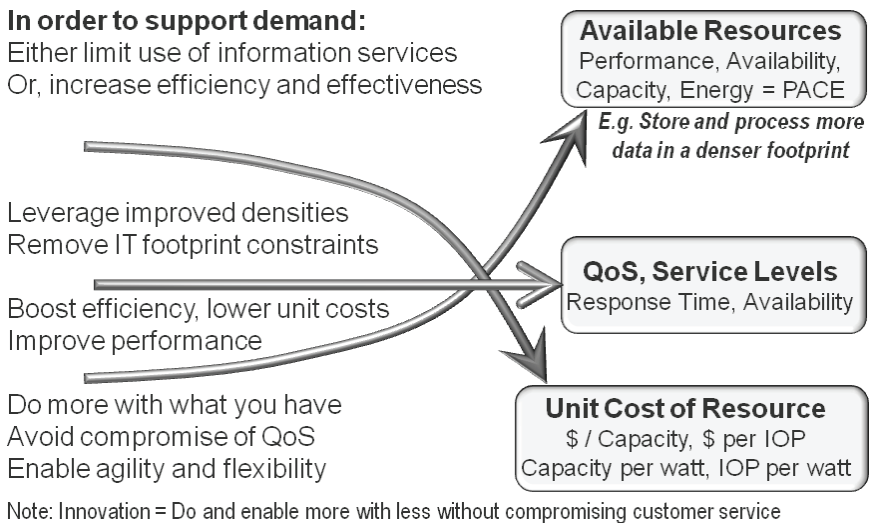


Figure 1.2 Supporting demand, maintaining quality of service (QoS), while reducing costs.

1.2.2. Addressing Business and IT Issues

Clouds, virtualization, and storage networks are tools, techniques, and best practices to help support or sustain growth while reducing per-unit costs, removing complexity, enabling flexibility or agility, and also enhancing customer experiences. Clouds, virtualization, and storage networks are not the objective themselves; rather, they are tools, vehicles, or mechanisms that can be used to help achieve broader business and IT objectives. They can be used for new, start-from-scratch environments; they can also be aligned with existing IT service delivery as well as help with a transition-over-time evolution of IT.

Thus, taking a step back from the technology, tools, and techniques, and keeping the bigger picture in focus, helps to understand what to use when, where, and why, as well as how to go about it in a more effective manner.

1.2.3. What Is Driving Data Growth and Information Reliance

The popularity of rich media and Internet-based applications has resulted in explosive growth of unstructured file data, requiring new and more scalable storage solutions. General examples of unstructured data include spreadsheets, PowerPoint, slide decks, Adobe PDF and Word documents, Web pages, video and audio, JPEG, MP3, and MP4, photos, audio, and video files.

Examples of applications driving continued growth of unstructured data include:

- Gaming, security, and other surveillance video or security
- Unified communications including Voice-over-IP (VoIP)
- Rich media entertainment production and viewing
- Digital archive media management
- Medicine, life science, and health care
- Energy including oil and gas exploration
- Messaging and collaborations (email, IM, texting)
- Internet, Web, social media networking, video and audio
- Finances, marketing, engineering, and customer relations management (CRM)
- Regulatory and compliance requirements

While structured data in the form of databases continues to grow, for most environments and applications the high-growth area and expanding data footprint along with associated performance bottlenecks is centered on semistructured email data and unstructured file data. Unstructured data has varying I/O characteristics that change over time—for example, data starting out with a lot of activity, then going idle for a time, followed by extensive reads, as in the case of a video or audio file becoming known and popular on a media, entertainment, social networking, or a company-sponsored website.

Data footprint is the total data storage needed to support your various business application and information needs. Your data footprint may, in fact, be larger than how much

6 Cloud and Virtual Data Storage Networking

actual data you have. A general approach to determine your data footprint is to simply add up all of your on-line, near-line and off-line data storage (disk and tape) capacity.

1.3. Business Issues and IT Challenges

I commonly get asked if virtualization and clouds are a passing fad, full of hype, or if they are real and being attacked by fear–uncertainty–doubt (FUD). Granted, and unfortunately, there is a lot of hype along with FUD, leading to confusion about both cloud and virtualization—tending to set them up as popular fads, much as compliance, “green” IT, information lifecycle management (ILM), client server, and storage networking were initially viewed.

Common business issues, challenges, and trends pertaining to IT include:

- Increased reliance on information services being accessible when needed
- Competitive and other market dynamics causing financial constraints and focus
- Regulatory compliance and other industry or corporate mandates
- Stretched resources (staffing levels, skill sets, budgets, facilities)
- The need to reduce costs while increasing services and productivity
- A shift from cost reduction or avoidance to efficiency and effectiveness models

How often do you use data storage or information services? Perhaps you use data storage without realizing it, at home, at work, and elsewhere. Data storage is in play all around us, used for different purposes and in various forms. Some might say that data storage is the most important IT resource compared to servers or computers, networks, desktop, laptops or workstations, or application software tools. On the other hand, some would say that networks are the most important, or servers, or whatever is that individual’s specialty. For the sake of argument I will position data storage as equal to servers, networks, hardware, and software, as all are needed to be effective.

Common IT issues, challenges, problems, and trends include:

- More data to process, move, manage, store, and retain for longer periods of time
- Increased reliance and expectation that information services be available 7x24
- Limited or strained resource constraints causing bottlenecks or barriers
 - People or staffing and applicable skill sets
 - Hardware, software, and networking bandwidth
 - Budgets (capital and operating)
 - Power, cooling, floor space
 - Time for backup or data protection windows
- Regulatory, compliance, and other regulations
- Demand causing performance, availability, capacity, and energy (PACE) impacts
- Software or hardware licensing and maintenance, support as well as service fees
- Aging IT infrastructures along with related interoperability and complexity
- Time involved in aligning IT resources to business or service needs
- Speed and accuracy of IT resource provisioning

When I talk with IT professionals or customers, I ask them if they have a mandate to reduce costs, which is a common industry messaging or rallying theme. Surprisingly, a majority of them indicate that it is not costs per se that they have to reduce (though some do). Rather, they have to do more with what they have with their current budget to support business growth, new applications, and functionality.

1.4. Business and IT Opportunities

Now, back to the question you woke up with this morning: “Do I need to have someone buy or implement a cloud, virtualization, or storage networking solution?”

Or maybe you woke up wondering how you are going to support business growth, demands for more data, flexibility, reduce cost, and enhance service delivery. Or perhaps you need to figure out how to defend your environment or market your environment to the rest of your business as opposed to the business going to external resources.

For some, efficiency and optimization can be avoidance or simply increasing utilization to reduce or spread costs over more work being done. However, another form of efficiency and effectiveness is stretching resources to do more while boosting productivity or removing barriers and constraints.

1.4.1. Traditional Information Services Delivery/Model

Delivery of information services continues to evolve. As this evolution of techniques, technologies, best practices, and new products continues, there are decades of legacy applications and data that need to be supported.

The basics of any information services delivery model can be summarized (Figure 1.3) as users accessing business applications running on a server with information stored somewhere, all accessed via some device over a network. The device could be a dumb terminal cabled or networked to the server with dedicated direct attached disk storage or a smart phone via a wireless connection.

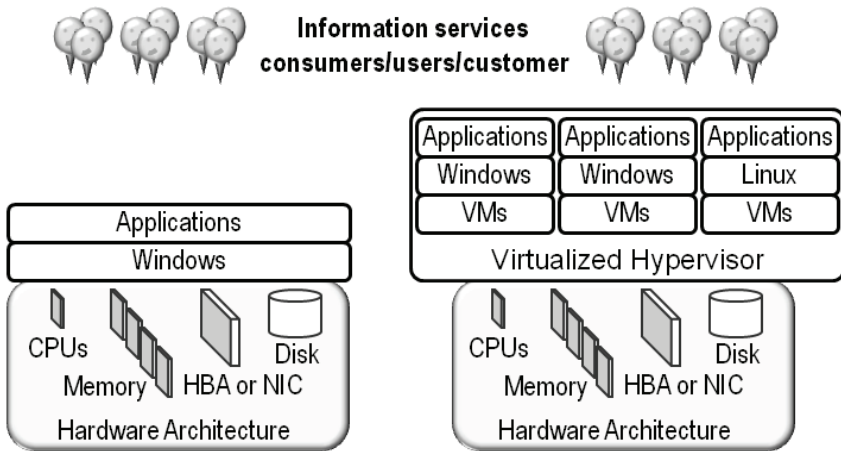
Similarly, the server could be virtualized or nonvirtualized with an operating system, database, and other tools to support and manage the applications and storage. From these basic information services delivery models, additional deployment scenarios can be established, including dedicated or shared applications, services, time sharing, or what we call today cloud and managed services, among others. Likewise, there can be different types of tiered servers, storage, and networks, which will be discussed in more detail over the next several chapters.

A common expression is that what is old is new and what is new is old. In the case of cloud and virtualization, for some these are new or perhaps repackaged open variations of what they have seen, heard, or read about previously. For example, IBM mainframes have had propriety virtualization for several decades. Those same platforms—which have been declared dead or dinosaurs—with their tenets of being highly optimized, metered or instrumented with metrics and reporting, scalable, and resilient—are what some cloud attributes seek to emulate from an open perspective. This has led to some

8 Cloud and Virtual Data Storage Networking

interesting discussions about why not keep things on the mainframe (clinging to the past) versus moving everything to the cloud (racing to the future).

Depending on the types of applications, the answer may be to move some or all to the cloud. On the other hand, leveraging a cloud managed service provider, hosting, or outsourcing can be the answer for other applications to coexist with your environment. For example, if you still have the need for an IBM zOS class mainframe, but it has become a small part of your environment, then outsourcing it or leaving a hosted or managed service can be an option.



Measuring, Monitoring and Managing Service Delivery
 Metrics and measurements, SLO, SLAs, Performance and QoS, Availability and reliability, RTO, RPO, uptime, Space capacity, Energy and economics

Figure 1.3 Information services delivery basics.

1.4.2. Information Factories

Most IT organizations or infrastructures exist to support the business applications and information needs of an organization. In some cases, the business applications services provided by IT include supporting factories, accounting, marketing, and engineering, among others. However, IT or information providers also often suffer from “shoemaker’s children” syndrome in that they may not have adequate insight or management tools for themselves. For example, an organizations may have accounting and tracking systems supported by IT, but does IT have accounting or metrics on performance, availability, capacity, configuration, energy, and economics for a given service delivery?

Traditional factories (Figure 1.4) leverage different tools, techniques, metrics, measurements, best practices, resources, and people skill sets to build and deliver goods or services to a prescribed service level and price point. Factories can be dedicated or

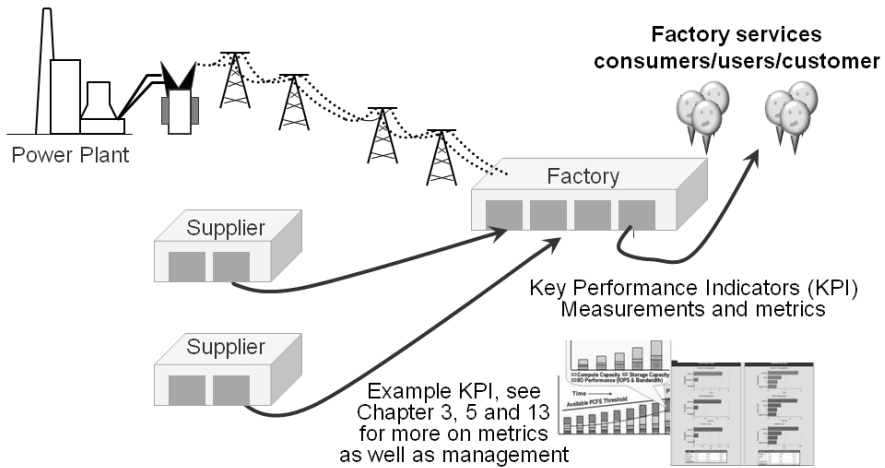


Figure 1.4 Traditional factory.

private, they can be shared or supplier-based, or they can be hybrid, similar to how IT services can be sourced, used, and delivered. An organization may have its own factory, or its factory could be a virtual or third-party jobbing or other service. Goods/services may be produced by someone under contract. IT services may be delivered via an organization’s own factory, via a third party, or virtualized.

Basic characteristics of factories include:

- Reliable, to meet demand, avoid downtime, avoid mistakes and rework
- Scalable, to meet changing workload demands
- Efficient, reduce waste, customer SLOs met in an economical manner
- Work is done quickly, yet reliably, with good quality
- Flexible capacity and ability to retool to meet changing needs
- Factories may be wholly owned, shared, or owned by a third party
- Factories consume materials and resources to create/deliver goods and services
- Those goods and services may in turn be consumed by other factories
- Factories produce product to a blueprint, template, or run book specifications

The notion of the information factory (Figure 1.5) sets up the discussion around cloud, virtualization, and storage networks on either a public, private, or hybrid basis.

For some, the idea of an information factory and cloud may bring *déjà vu* experiences of the information utility model of the late 1980s and early 1990s.

Additional characteristics of factories include:

- Rely on suppliers or secondary and tertiary factories (subs)
- Have bill of materials, metrics and measurements, costing information
- Quality assurances programs to ensure that QoS and SLOs are being met

10 Cloud and Virtual Data Storage Networking

- Focus on reducing defects and waste while boosting productivity to reduce cost
- Build templates for optimized information service delivery
- Best practices, processes, policies and procedures
- Balance of productivity, cost control, waste reduction, utilization, meeting SLOs
- Leverage new technologies that have good payback for enabling goals
- Cost-effective habitats for deploying and using technologies
- Efficiency gained with repeatable processes, and increased workload activity

Information factories can be

- Private
- Public
- Hybrid

Information factories (or clouds) should be or enable:

- Multitenancy, measurability, and accountability
- For service providers, this can include chargeback
- Secure, flexible, dynamic, scalable, and resilient
- Able to relocate services as needed
- Rapid deployment and provisioning of resources
- Efficient, cost-effective resource usage that meets QoS and SLAs
- Automate and guide users or customers to best-fit services selection

The similarities between factories, information factories, clouds, and information services delivery should be clear.

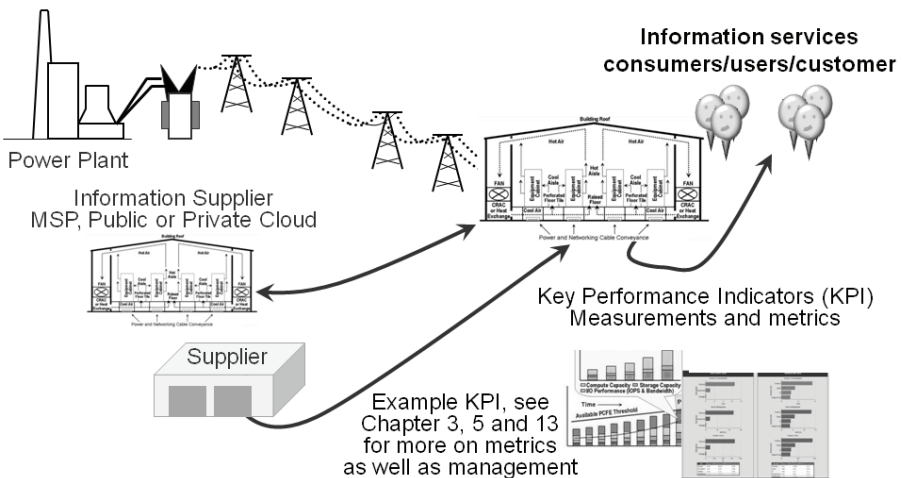


Figure 1.5 Information factory.

1.5. Opportunity for Cloud, Virtualization, and Data Storage Networking

Like a physical factory, some of an information factory's work is done on the premises and some off-site at other locations, including those of subcontractors or suppliers. In the case of information factories, the product being produced is information services, with the machinery being servers, storage, and I/O networking managed with software, processes, procedures, and metrics. Raw materials include data, energy to power and cool the physical facility, and technologies, all operating to deliver the services at a low defect or error rate while meeting or exceeding QoS, performance, availability, and accessibility requirements in a cost-effective manner.

For some cloud compute or storage providers, the value proposition is that they can supply the service at a lower cost than if you use your own capabilities. Similar to service bureaus, out-sourcing, managed service, or hosting facilities of the past, cloud-based services are a means of shifting or avoiding costs by moving work or data elsewhere to be processed or stored.

However, it is a mistake to consider clouds for just for their cost-saving abilities while ignoring performance, availability, data integrity, ease of management, and other factors that can impact service delivery and expenses. Clouds should be looked at not as a replacement or competing technology or technique, but rather as a complementary approach to existing in-house resources.

Cloud computing and storage are simply additional tiers of servers and data repositories that may have different performance, availability, capacity, or economics associated with them to meet specific business and/or application needs. That is, cloud computing and cloud storage coexist and complement what is currently being done, with the objective of boosting quality of service, availability, or customer satisfaction while supporting more data being processed, moved, and stored for longer periods of time at a lower unit cost.

1.5.1. IT Clouds and Virtualization: Not If, Rather When, Where, Why, and How

There are many different types and definitions of clouds, including those of the National Institute of Standards and Technology (NIST) and the Data Management Task Force (DMTF). Cloud computing is a paradigm, and thus its definition is still evolving along with use cases and the underlying technologies, techniques, and best practices.

Some see clouds as the wave of the future, even if they're not sure what that future may be. To others, a cloud is a cloud if, and only if, it is outside of what you currently are doing or have done with IT. Some will argue that a cloud is only a cloud if new hardware or software is involved, while others will assert that a cloud is only a cloud if your applications and data exist outside your location.

Consequently, different people will have different thoughts or perspectives about clouds, depending on their perception or definition. For example, in Figure 1.6, thoughts and opinions based on an ongoing StorageIO blog research poll of a mix of

12 Cloud and Virtual Data Storage Networking

vendors, IT professionals, and others shows at two extremes; those who see clouds as the solution to everything, and those who see no chance or place for a cloud whatever it happens to be. In the middle are about 81–82% (the poll is ongoing, so results may vary) of the respondents, who vary from seeing a place for clouds depending on the definition or use case to others who are skeptical but want to learn more about what to use when, where, and why.

It’s about delivering information services in a cost-effective manner that supports demand while meeting service objectives (Figure 1.2). Figure 1.7 shows various information services delivery models that rely on different techniques, technologies, and best practices that can be competitive or complementary. Cloud metaphor has been used in and around IT for decades as being a means to abstract underlying networking details or an applications architecture.

A key attribute of clouds is that of abstracting or masking underlying complexities while enabling agility, flexibility, efficient, and effective services delivery. This leads to some confusion, which for some creates opportunities to promote new products, protocols, standards, or services while for others it means repackaging old initiatives. For example, some may have a *déjà* moment when looking at a cloud presentation back to the late 1980s during the information utility wave that was appearing with the advent of the x86-based PCs along with client servers. For others, that moment could be time sharing or service bureau, and for others the cloud’s Web-based and highly abstracted virtualized environments.

What are IT clouds? Where do they fit? How does tape coexist with clouds? Like many IT professionals, you may already be using or leveraging cloud-based computing or storage techniques, either as a product or as a service, without realizing it.




Figure 1.6 IT cloud confusion and opportunity. (Source: StorageOblog.com.)

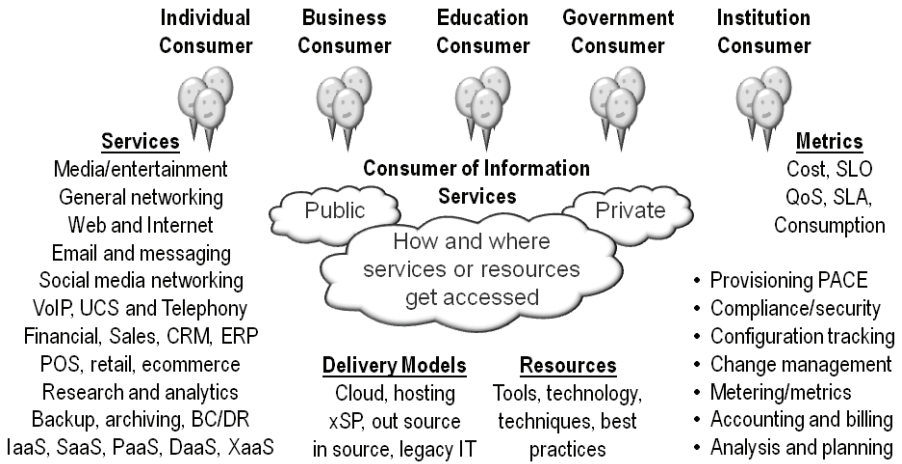


Figure 1.7 Various information services delivery and resource models.

Common cloud-based functions or services include:

- Remote or off-site backup, replication, vaulting, or data copy
- Remote or off-site storage on-line, near-line, or off-line
- Email and messaging services including social media networking and Web 2.0
- Archive, fixed content, and other reference or look-up data
- Website, blog, video, audio, photo, and other rich media content hosting
- Application hosting (e.g., salesforce.com, concur expense, social media)
- Virtual server or virtual machine (VM) hosting (Amazon, VCE, etc.)
- General on-line storage or application-specific storage such as Google Docs

Does this mean that if backup, business continuance (BC) or disaster recovery (DR), or archive data is sent off-site to a storage or hosting facility, it has been sent to the cloud? Some say no unless the data were transmitted electronically to on-line disk at a service provider location leveraging programmatic interfaces and other cloud ware (technology, services, or protocols developed, optimized, or packaged for public and private clouds). That might also be a product- or services-based definition. However, in theory, the concept is not that far off, as clouds, in addition to being a product or service, are also a management philosophy or paradigm to do more with what you have without negatively impacting service delivery.

Characteristics of clouds include:

- Ease of service access (self-service)
- Ease of service deployment or provisioning
- Elasticity and multitenancy
- Safety, security, with data integrity

- Flexibility, scalability, and resilience
- Cost effectiveness and measurability
- Abstraction or masking of underlying complexities
- Can move or change the focus and presentation
- Leverage repeatable processes, templates, and best practices
- Efficiency as a result of scale and increased workload or usage

Confusion exists in that there are many different types of clouds, including public and private, products and services, some that use familiar interfaces or protocols with others using different technologies. Clouds can be a service or a product, an architecture, or a management paradigm, similar to previous generations such as the information utility or service-oriented architectures (SOA), client server computing, and others.

What this means is that some of the tenets of cloud storage and computing involve shifting how resources are used and managed, thus enabling the notion of an information factory. They can be external or internal, public and private, housed at a hosting or co-location facility as well as at traditional out-sourcing or managed service providers. Thus a hosting site may or may not be a cloud, and a cloud can leverage hosting services but does not require them. Various information services delivery models are shown in Table 1.1.

1.5.2. Private Cloud: Coexistence vs. Competing with Legacy IT

Clouds and virtualization should be seen for what they really are as opposed to what they are often portrayed to be so that a solution can be sold. In other words, take a step back, look at the business issue, then apply the applicable technology or task at hand to the situation. Instead of clouds being a solution looking for a problem, they become a tool and technique that can be used in different complementary ways. Cloud computing and storage are another tier of traditional computing or servers providing different performance, availability, capacity, economic, and management attributes compared to traditional technology delivery vehicles.

If IT is a core piece of the business, it probably makes more sense to retain tighter control. For example, a manufacturing company may out-source or rely on suppliers for key components, or perhaps even provide contract-under-supervision manufacturing, leveraging proprietary processes and techniques. Than if the related IT functions are also essential, they too would be retained and kept close to the vest while other functions might be out-sourced or sent to the cloud.

1.6. Common Cloud, Virtualization, and Storage Networking Questions

Does cloud storage require special hardware or software, or is it a matter of how those resources are used, deployed, and managed? As with many things, the answer is, “It

Table 1.1 Information and Data Services Deployment Models

Model	Characteristics and When to Use	Examples
Co-location ("colo")	Shared facilities with dedicated space for your equipment. Power, cooling, security, networking, and some monitoring or other optional services provided. Primary or supplemental space for your hardware.	Iphouse, Rackspace, Sungard, Timewarner, visi, and many others
Hosting services	Services and or application hosting. These could be email, Web, or virtual machines. In addition, these could be Applications as a Service (AaaS). Many colos provide application hosting services. Instead of renting space for hardware, you rent time and use of software.	VCE, ADP, Amazon, Bluehost, Google, HP, IBM, Iphouse, Oracle, Rackspace, Salesforce, and others
Legacy IT	Hardware (servers, storage, and networks) plus software (applications and tools) are bought or leased, operated, and managed by IT staff.	Hardware and software in your existing environment
Managed service provider	Similar if not the same as a hosting service. Provides some service, which might be applications, archiving, backup, storage space, backup, replication, email, Web, blogs, video hosting, business continuance/disaster recovery, among others. Instead of you running or hosting the application, you use a service provided to you that runs on someone else's shared infrastructure. Some may have déjà with service bureaus or time sharing.	Amazon, AT&T, Carbonite, Campaigner, EMC Mozy, GoDaddy, Google, Iron Mountain, Microsoft, Nirvanix, Seagate i365, Sungard, Terremark, Wells Fargo vSafe, among others
Out-sourcing	Could be on- or off-site, where you either move your applications and possibly equipment to a third party who operates and manages to specific service-level objectives (SLOs) and service-level agreements (SLAs).	Dell/Perot, HP/EDS, IBM, Lockheed/Martin, SunGard, Terremark, Tata, Xerox/ACS, and Wipro, among others
Private cloud	Dedicated to an organization need. Could be managed by IT staff or third party on-site or off-site. May use cloud-specific technologies or traditional technologies managed with cloudlike premises or paradigms. May be called in-source or IT 2.0.	Instrumented or metered IT environment for effective service delivery. Many different products available
Public cloud	An IT infrastructure that supports shared computing, storage, and or application services provided by an organization. The services may be available free or for a fee. They can be used to replace or complement existing IT capabilities. Access shared applications such as salesforce, email, backup or archive destination, virtual servers, and storage, among others. Buzzwords include applications as a service (AaaS), infrastructure as a service (IaaS), storage as a service (SaaS), and platform as a service (PaaS), among many other XaaS variations.	AT&T, VCE, Amazon E2C or S3, Google, Iron Mountain, Rackspace, Salesforce, Terremark, HP, and IBM, among many others

depends.” For example, a vendor of a particular product will try to convince you this product is needed, and that a deployment must incorporate this solution.

Are clouds real, or just hype? Certainly there is a lot of hype from vendors and the industry in general, as is often the case with new or emerging technologies, techniques, paradigms, or movements. Likewise, some solutions or vendors trade more on hype and FUD compared to others. However, there are also plenty of real and viable techniques, technologies, products, and services that organizations of all sizes can be leveraging today as well as planning for in the future. The trick is sifting through the fog of cloud fodder, FUD, and hype to determine what is applicable for your environment today and into the future.

Of course, your definition of a cloud will also have a bearing on the above; however, I hope that after reading this book, you will also see that there are many different approaches, technologies, techniques, services, and solutions that can be applied to different situations. In other words, let’s move from a solution looking for a problem to problems that need solutions, and what to use when, where, why, as well as how.

What is virtualization life beyond consolidation? The next major wave (trend) of virtualization, including from applications to desktop, servers to storage and networking, will be an expanded focus on agility. What this means is that there will continue to be an expanding market capability for consolidation, which is the current focus of the virtualization wave.

However, the next wave shifts to expand in another dimension that is less focused on how many virtual machines (VMs) there are per physical machine (PM) and instead around agility, flexibility, as well as ease of management. In other words, for those servers or storage systems that cannot be consolidated and hence are thought to be unsuited to virtualization, break down those myths and virtualize for agility instead of consolidation.

Should everything be consolidated? Generally speaking, I would say no. However, many, if not most, things, can be virtualized, assuming that for some servers or storage there may be fewer VMs per PM or even a single VM per PM. Some might wonder why you would virtualize with only one VM per PM, as this seems to defeat the purpose of virtualization. This is only the case if your view of virtualization is that of consolidation. However, many things can be virtualized for agility, emulation, transparency, or abstraction, keeping performance, quality of service, and other constraints or concerns in mind.

Virtualization of Windows and x86 environments is understood, but what about UNIX and other systems or environments? Some UNIX as well as Linux distributions that rely on or support x86 platforms can run on popular hypervisors such as vSphere, Hyper-V, and Xen. However, for other UNIX systems, such as HP-UX, Solaris (non x86), and AIX, those systems have features as part of the operating system or underlying hardware platforms for virtualization. Examples include Logical Domains (LDOMs) for Solaris, HP-UX partitions, and hypervisor in the IBM pSeries that supports AIX, among other systems.

Do clouds have to be physically off-site and hosted by a third party? No, clouds can be implemented internally (known as a private cloud) at your premises using existing technologies as well as leveraging off-site or third-party-provided services.

Do you have to start from scratch to leverage clouds or virtualization? Generally speaking, you do not have to start with a fresh sheet to leverage clouds and virtualization. In some cases you can leverage different technologies or services to coexist and complement what you have, while others may require changes. You may also want to leverage the opportunity to change the way things are done as part of deploying cloud and virtualization technologies.

Are clouds, virtualization, and storage networks only for large environments? No, those are common myths in that many solutions, services, or products are marketed toward higher-end environments. However, there are also solutions applicable across different size organizations, with some, such as cloud-based backup, available all the way down into the consumer space. Cloud-based backup is a market “sweet spot” for small to mid-sized businesses, given the relatively smaller amount of data that can be moved along with opportunities to leverage cloud or managed service providers.

Are clouds a replacement for traditional IT? Depending on what you believe or want to believe defines a cloud, the answer is maybe or no. Some cloud services are aimed at replacing all or portions of traditional IT, while others are positioned to complement existing IT. Likewise, there are cloud services and products that can be used in either public or private that coexist with traditional IT.

Is chargeback a requirement for having or using a cloud? For public cloud and other fee-based service models, some form of invoicing, charging, and billing is needed. However, it is a popular myth that all private clouds need chargeback. What they do need is instrumentation, metrics, and measurement, including bill-of-material (BOM) information about the cost to deliver a given level of service. For some organizations that currently implement formal chargeback with real or virtual invoices, it would make sense to be continued candidates for chargeback. The important thing is that there are metrics to show how resources are being used and accounted for. If your view of accounting, merging, measuring, and reporting is chargeback, then yes, you need chargeback. However, for many other scenarios, the emphasis should be on accounting and metrics that matter.

Do clouds and virtual environments automatically guide users or customers where to place data? Various cloud services, along with some products, have tools and wizards to help guide users or customers to what resources are best for a given need. Some solutions also help to automate or support scheduled polices to perform recurring tasks or functions. Automation and tools can help shift recurring tasks from having to be performed by IT staff, enabling skilled workers to focus on service enhancement, analysis, and other value-adding functions.

1.7. Cloud, Virtualization, and Storage Networking: Bringing It Together (for Now)

Additional examples of technologies to use for addressing various problems or enabling opportunities are shown in subsequent chapters. Figure 1.8 shows as an example of how clouds, virtualization, and storage networking technologies and techniques combine in complementary manners for different IT and business purposes. For example, server

and storage virtualization are shown for both consolidation as well as to enable agility, simplified management, and emulation to bridge old technologies to new roles.

Other items shown in Figure 1.8 include leveraging public cloud-based backup and archive services along with solutions for enabling private cloud and virtualized environments. Various types of I/O networking technologies are used to attach servers (physical and virtual) to local and remote storage. Different types of storage including high-performance on-line or primary, secondary, near-line as well as off-line and removable technologies are also shown.

Although not shown explicitly, various management tools, protocols, and interfaces also combine in Figure 1.8 for enabling high availability and business continuance/ disaster recovery, including routine backup as well enabling data footprint reduction (DFR). Also shown in Figure 1.8 are metrics for gaining situational awareness of resources, including cost of service delivery and service-level objectives, among others. These and other topics, technologies, tools, and techniques will be discussed in more detail in subsequent chapters.




Figure 1.8 Public and private cloud products as well as services can coexist.

1.8. Chapter Summary

Don't be scared of clouds: Learn to navigate your way around and through the various technologies, techniques, products and services, and identify where they might complement and enable a flexible, scalable, and resilient IT infrastructure. Take some time

to listen and learn, and become educated about the different types of clouds (public, private, services, products, architectures, or marketecture), their attributes (compute, storage, applications, services, cost, availability, performance, protocols, functionality), and their value propositions.

Look at how cloud technologies and techniques can complement your existing environment to meet business objectives. You might find there are fits, or you might find there are not, but it's important to do the research and know where you stand.

The subsequent chapters will look at the what, why, where, when, and how to use various techniques to address or enable business and IT objectives. Given the diversity of readers' backgrounds, feel free to jump around different chapters as you see fit. Likewise, I need to put in a shameless plug to read my other books, *The Green and Virtual Data Center* (CRC) and *Resilient Storage Networks: Designing Flexible and Scalable Data Infrastructures* (Elsevier), as companions to this book, as well as my blog and website.

General action items include:

- Avoid simply moving problems or bottlenecks—find and fix them instead.
- With clouds and virtualization, the question is not if, but rather when, where, with what, and how.
- Prepare for next wave of virtualization: Life beyond consolidation enabling agility.
- Cloud services, products, and solutions are complementary to existing IT infrastructures.

