IBM

Redpaper

Elisabeth Stahl
Lydia Duijvestijn
Avin Fernandes
Pamela K. Isom
Dave Jewell
Martin Jowett
Todd R. Stockslager

Performance Implications of Cloud Computing

In this IBM® Redbooks® publication, we present the performance and capacity of the evolving cloud computing model. We provide an overview of cloud computing and its potential benefits from IT and business perspectives. We also present different cloud computing deployment models and roles involved in delivering cloud solutions.

We present design and delivery solutions, highlighting the key performance and capacity considerations, and outline performance engineering considerations. We conclude by reviewing current cloud offerings and briefly evaluate their performance. We provide use case scenarios for clarification and elaboration of real-life cases where performance considerations are part of the cloud solution strategy.

This paper is intended for IT professionals who are interested in adding cloud computing solutions to their data management systems. This paper also is intended for business professionals who seek a basic understanding of the performance implications of using cloud computing to deliver their business capabilities.

In follow-on papers, we detail performance considerations in depth, from the perspectives of those who are responsible for creating cloud solutions and those who are responsible for managing the solutions after deployment. For more information, see *Performance and Capacity Themes for Cloud Computing*, REDP-4876.

About cloud computing

Cloud computing is a new way to use and deliver IT services that is driven by consumer Internet trends. The cloud computing model builds on the maturation of the World Wide Web, combining rapid scalability, proliferation of the Internet and Internet-connected devices, unprecedented self-service, and the emergence of elegant web-based applications. The cloud allows users to execute complex computing tasks without the need to understand the underlying technology.

Cloud computing is emerging at a critical time for the IT industry. There is a growing realization that physical and IT assets, systems, and infrastructure are fast reaching a breaking point. As the pace of business and society in general continues to accelerate, the physical and digital foundations on which progress depends are straining to keep up because of the following factors:

- The explosion of data, transactions, and digitally aware devices is straining existing IT infrastructure and operations.
- Exponential growth in communications subscribers and services is showing limitations in network bandwidth and storage capacity.
- Supply inefficiencies and demand spikes are putting pressure on energy and utility systems.

The emerging cloud computing model, a key part of the Smarter Planet™ vision articulated by IBM and embraced by others, uses improved technologies. This cloud model is characterized by innovative Internet-driven economics that use the massive underlying ability of the technology to scale. Cloud computing offers the vision for improved service (not just high availability and quality of existing services) but also meeting expectations for real-time access along with dynamic access to innovative new services. Cloud computing also allows for reducing cost, not just containing it, but achieving breakthrough productivity gains through virtualization, optimization, energy stewardship, and flexible sourcing.

However, cloud computing does have its challenges. The complexity of cloud computing, with its increased dependence on Internet, virtualization, and on-demand enabling technologies, requires critical attention to performance and capacity considerations throughout the design, delivery, and management of cloud-based solutions.

Cloud computing benefits and challenges

Cloud computing is a flexible and cost-effective delivery platform for providing IT services over the Internet and is useful for numerous applications. You can rapidly deploy and easily scale cloud resources, with all business processes, applications, and services provisioned on demand, regardless of the user location or device. Cloud computing gives your organization the opportunity to increase service delivery efficiencies, streamline IT management, and better align IT services with dynamic business requirements. It provides support for core business functions along with the capacity to develop and deploy new and innovative services, making them more accessible than ever before.

Business factors for cloud computing

Why is cloud computing a viable option for meeting storage and IT needs? Presented here are some key business factors that can affect cloud solution considerations.

The number one factor is cost. Clouds offer a pay-as-you-go model that allows your company to invest in resources as needed rather than in anticipation of the need. This ability to pay for only what you need is especially important when you consider that the investment is considered by the business as an operating expense rather than as a capital expenditure.

The second key factor is speed to value, or “I can deploy resources in this environment”, a key factor in environments that need to grow or shrink quickly.

Combine these two factors and you get a service that grows and shrinks with the needs of the users and you pay only for service usage. For example, one day, cloud service might cost you USD10.00; another day, when the demand was higher, it might cost USD20.00. Another advantage is the portability of services. This portability ensures that the service is available regardless of the operating platform.

Overall, clouds dynamically offer you a service that meets your availability and performance needs and keeps operating costs low. Clouds also limit the expenses for what was used rather than capital investment based on projections of what might be needed.

Cloud benefits and advantages

Implementing a cloud provides advantages in architecture, business solutions, and cost reduction.

Infrastructure architecture advantages

The cloud computing model offers many advantages to IT organizations that are deploying infrastructure architecture (that is, infrastructure provisioning). The model has the potential to expand and automate resource virtualization, data center resource billing and metering, and self-service catalogs and requests. Cloud computing also enables workload-optimized solutions with efficiencies and innovations in areas, such as development, testing, analytics, infrastructure, and storage. These efficiencies and innovations can be quantified in a number of areas:

- Reduced capital expenditures and labor costs.
- Rapid provisioning and de-provisioning of services.
- Enhanced resource pooling as computing resources are pooled to provide multiple capabilities to users. Resource pooling allows virtual and physical resources to be dynamically configured and assigned based on service-level agreements (SLAs) and demand.
- Superior service management with visibility, control, and automation across IT and business services.
- New deployment choices over the cloud, behind the firewall, or as an integrated service delivery platform.

Business solutions benefits

Business solutions drive technology solutions, especially for cloud computing with respect to some recent trends in information storage:

- Up to 80% of data today is unstructured content, including email, video, and images.
- Medical images take up 30% of the available worldwide storage.

  The field of medical imaging is driving breakthroughs in diagnosis and treatment in medicine. As a result, there is exponential growth in the number and size of digital medical images. Medical images that were two-dimensional and 1 MB in size a few years ago are now typically four dimensional and 1 TB in size.

Cloud computing as a technology is viable because it offers many of the following benefits to business organizations that are deploying business solutions:

- Optimized systems resources to keep developers productive
- Reduced capital and licensing expense (as much as 50 - 75%) by on-demand provisioning of virtualized test resources [1]
- Decreased operating and labor costs (as much as 30 - 50%) by automated provisioning and configuration of test environments [1]
- Facilitated innovation and shorter time to market by improving test provisioning from weeks to minutes and reducing test cycle time [1]
- Improved quality by reducing defects that result from faulty configurations and poor modeling, as much as 15 - 30% [1]

Reducing costs

Just as it is true that business solutions always drive technology solutions, it is equally true that costs drive both business and technology solutions. Cloud computing identifies and attacks costs as a factor for the IT advantages and business benefits listed previously in this publication. Hardware virtualization reduces costs by boosting hardware utilization by stacking multiple virtual servers in a physical server. This virtualization reduces software licensing costs because charges for operating systems and other software are typically based on the number of physical servers rather than on the number of instances. Therefore, fewer physical servers require fewer licenses.

System administration and operation costs are also reduced in cloud infrastructures with fewer physical servers. Labor savings are realized in provisioning processes that automate and standardize service request management and fulfillment. These reduced administration and operations costs also extend to reduced development and test team costs. Providing ready-to-use systems quickly reduces development and test team idle and waiting times and increases their flexibility in planning, execution, and deployment.

Cloud computing challenges

A number of challenges need to be overcome before the benefits of cloud computing can be realized. Deploying a set of cloud constructs requires thorough planning. The areas of concern noted here are all critically important. Each area can incrementally improve overall operations. However, improvements in one area can cause strain in another:

- Demand services provisioned on demand for a better enterprise-wide information infrastructure can stress the issues of security and business resiliency.
- Creating highly virtualized resources demands a stronger, more integrated service management approach.
- Consolidating to optimize systems can drive up the density of the systems, which strains the environment issues.
- Converging business and IT infrastructures can be a daunting task if not handled in an integrated way.

---

[1] *Cloud Computing: Save Time, Money, and Resources with a Private Test Cloud,* REDP-4553, at this website: http://www.redbooks.ibm.com/redpapers/pdfs/redp4553.pdf

Therefore, it is critical to consider all of the elements together. How do things interrelate so that improvements in one area are matched with tools and techniques to support them in another? Although taking a view of each of these areas is important, even more important is a plan to integrate them, thus creating the backbone, or DNA, needed to thrive in a Smarter Planet.

For more information about obstacles for Cloud computing, see *Above the clouds: A Berkeley View of Cloud Computing* found at:

http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

Cloud perspectives

Cloud is a new delivery and acquisition model of IT and the enabled services of IT that is inspired by Internet consumer services. Through this model, users consume resources in a fundamentally different way. Previously, usage was based on a limited set of resources that were well-delineated and identified beforehand. With cloud computing, services can be rapidly provisioned and released with minimal human management effort or service provider interaction. The cloud might be considered from different perspectives. Clouds might be deployed by way of several different models and delivered as varied services. Roles for cloud consumption and delivery are important to understand to evaluate performance implications. The cloud perspectives presented in this paper are shown in Figure 1.

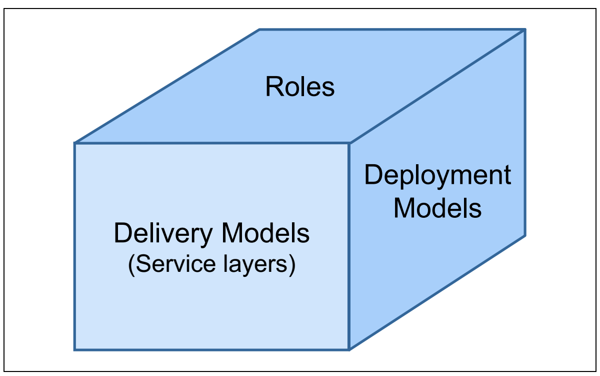

Figure 1 Cloud perspectives

To successfully consider these different perspectives, you need to be aware of three fundamental questions that pertain to cloud service offerings:

1. What types of services are provided?
2. How and where are the services deployed?
3. Who is involved?

Types of services and cloud delivery models

The following service layers are available in the cloud arena, depending on the level in the solution stack where the service is being delivered:

- The top layer, corresponding to the business application viewpoint, hosts Software as a Service (SaaS).
- The middle layer, corresponding to the middleware services viewpoint, hosts Platform as a Service (PaaS).

- The lowest layer, corresponding to the physical environment viewpoint, hosts Infrastructure as a Service (IaaS).

The cloud is defined as the virtualized infrastructure found on the lowest level of the solution stack[2]. The higher service layers depend on the underlying, supporting service layers. Service providers can be service users. For example, a SaaS provider might be a SaaS user; a SaaS provider might not be a PaaS user while SaaS and PaaS providers are directly or indirectly IaaS users.

The QoS characteristics of the upper service layers depend on those characteristics in the underlying layers. SaaS provides user applications as a service and contains the following elements:

- Business process support
- Enterprise applications
- Collaboration services

Although the SaaS classification existed before cloud computing, we accept that massively scalable SaaS offerings can be considered cloud computing service models. The applications are generally accessible through the web; however, user configuration is limited to application configuration. Two examples of SaaS are Fidelity.com (business processes) and IBM LotusLive™ (collaboration services).

At the level of SaaS, QoS is directly perceived by users and defined by business transaction response times and throughput, business service reliability and availability, and by scalability of the applications.

PaaS contains the following elements:

- Databases
- Middleware
- Development tools
- Java™ and Web 2.0 run times

PaaS provides a development and deployment environment for applications. An example of PaaS is Google.

At the level of PaaS, QoS is indirectly perceived by users and defined by technical transaction response times and throughput, technical service reliability and availability, and by scalability of the middleware.

IaaS features the following functionality:

- Server functionality
- Networking functionality
- Data center functionality
- Storage functionality

An example of IaaS is Amazon Web Services.

The quality of services, provided at the IaaS, is defined by infrastructure *performance, capacity, reliability, availability,* and *scalability.*

The three cloud service layers are shown in Figure 2 on page 7.

[2] *Above the clouds: A Berkeley View of Cloud Computing*, at this website:
http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

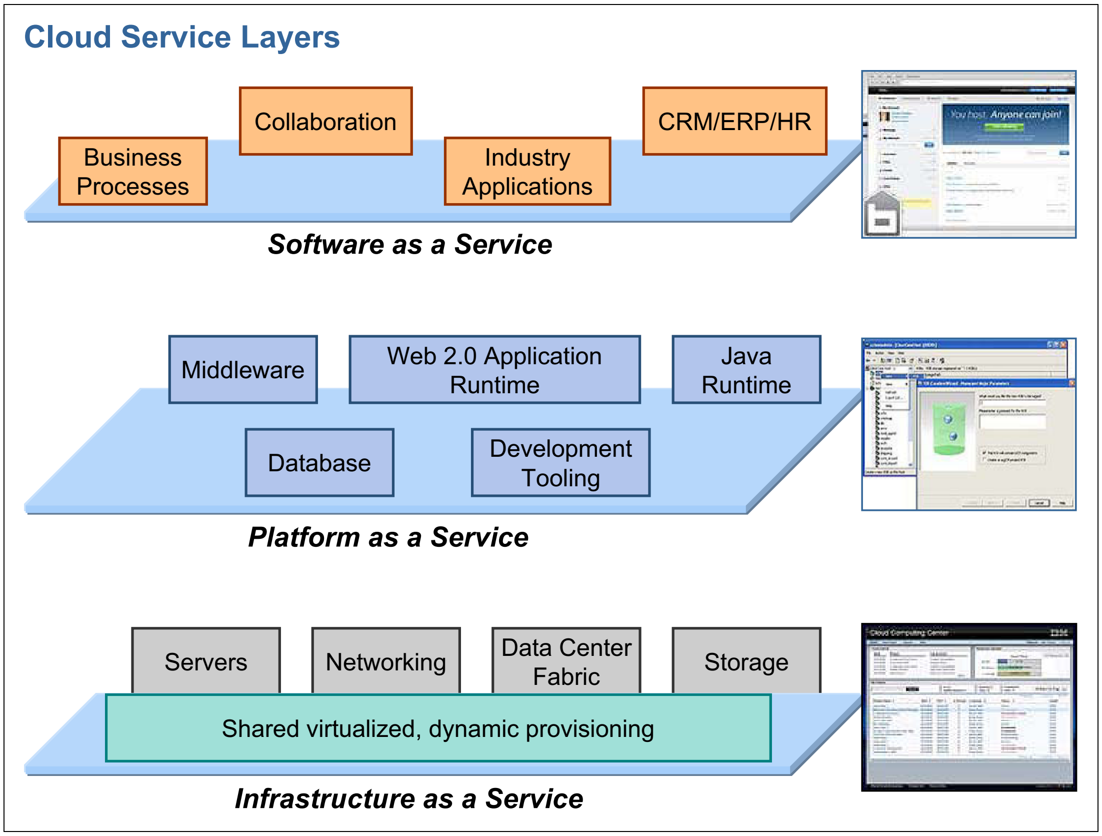

Figure 2 Cloud service layers

How and where are the services deployed

How and where cloud services are deployed influences the type of cloud services model used and the roles considered for cloud service consumption.

Cloud deployment models

Cloud deployments are based on workload characterization and nonfunctional business requirements.

A *public cloud* is owned and managed by a third-party service provider, and access is by subscription. A public cloud offers a set of standardized business process, application, and infrastructure services on a flexible, price-per-use basis. Advantages of a public test cloud include standardization, capital preservation, flexibility, and a shorter time to deploy environments and, as a consequence, the applications under test.

Private clouds are owned and used by a single organization. They offer many of the same benefits as public clouds, and they give the owner organization greater flexibility and control. A private test cloud provides greater ability to customize, drives efficiency, and retains the ability to standardize and implement organizational best practices.

Other advantages of private clouds include availability, resiliency, security, and privacy. Private clouds can also provide lower latency than public clouds during peak periods. This feature is essential when guaranteeing performance as a QoS is key.

Many organizations embrace both public and private cloud computing by integrating the two models into *hybrid clouds*. According to the National Institute of Standards and Technology (NIST), hybrid clouds are bound together by standardized or proprietary technology that enables data and application portability. A hybrid cloud can also contain multiple services and any combination of providers and consumers, which is described next in “Cloud roles”. In general, hybrids are designed to meet specific business and technology requirements, helping to optimize security and privacy with a minimum investment in fixed IT costs.

Cloud roles

Several roles are considered for cloud service consumption and delivery:

- *Consumer*: User or users of cloud services. In private clouds, the users are within the same enterprise boundary as the supplier. However, in public clouds, the consumers can be either within the same enterprise as the provider or (more commonly) external to the provider.
- *Provider*: The party responsible for providing access to cloud services for registered consumers and for maintaining the QoS attributes for the individual service access requests, including performance. The providers own the physical assets that are required to produce and deliver cloud services to the consumer.
- *Integrator*: The party responsible for the design and construction of a composite cloud service offering. This role might be within the same enterprise as the provider or it can be separate. In the latter case, the service integrator can retain some level of accountability for the cloud services provided by the third-party consumer.

Each of these roles includes a number of sub-roles, as shown in Figure 3 on page 9.

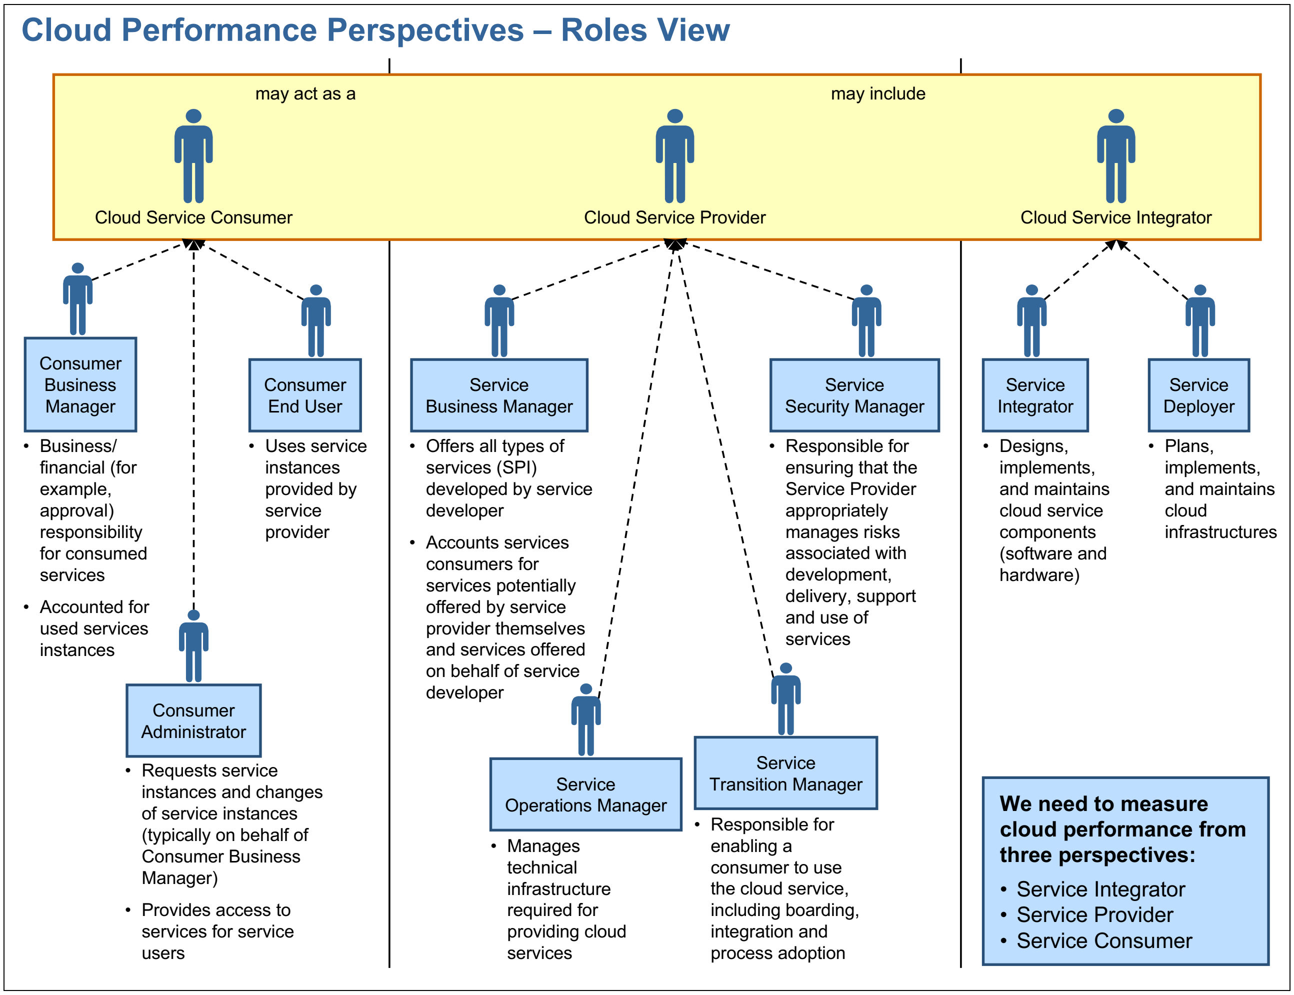

Figure 3 Cloud roles

The actual deployment and delivery models for clouds and the perspective of the role involved, consumption, and delivery are integral to the performance analysis of these solutions.

Cloud service consumer roles

Do not assume that cloud service consumers are merely users, or that they are not interested in performance.

By virtualizing and extending computing resources into the business, cloud services make consumers active and interested participants in the performance of cloud services.

The *consumer business manager* is most interested in performance in terms of the benefit per dollar spent in obtaining cloud services. In addition, the business manager is keenly interested in ensuring that the cloud performs as expected so that the cloud services purchased are able to add to the profitability of the business.

The *consumer user* is most interested in performance because they are the person most directly affected by cloud performance issues. The user is responsible for integrating public cloud services into the processes of the consumer, and private or hybrid cloud services into the computing services of the consumer.

The person in the *consumer administrator* role can be a technical staff member responsible for integrating a private or hybrid cloud into existing data center services. They are interested in performance as a component of the usability, manageability, and sustainability of the cloud service.

Cloud service provider roles

The *cloud service business manager* works with the consumer business manager to ensure that the cloud service meets the business needs for the consumer business manager. For a private cloud, both business managers work for different organizations within the same company.

The *service operations manager* is focused on the performance aspects of the technical infrastructure that supplies the cloud services. This role is most directly focused on traditional system performance, in this case as it relates to delivering cloud services to integrators and consumers outside of the managed data center.

The *service security manager* is concerned with the performance of security integration of private and hybrid cloud services with the user directory services of the consumer. Ensuring overall security of the cloud services offerings without reducing performance is also a key component of the job description of the security manager.

The *service transition manager* brings the services consumer on board, provides support and training, sets performance expectations, and works as a liaison between the organizations.

Cloud service integrator roles

The increased technical complexity of integrating disparate components into deployable cloud services increases the performance risks and challenges for the cloud service integrator roles.

The *service integrator* performs traditional development to design, implement, and maintain cloud services, and the role can include the full range of hardware, software, or services configuration. The role also can include the integration or assembly of subcomponents to enable cloud services. Performance and scalability of the integration are key concerns of the service integrator.

The *service deployer* plans, implements, and maintains cloud infrastructure and deploys and maintains cloud services to the infrastructure. The service deployer also focuses on system performance and resource utilization of the cloud service.

Cloud performance engineering and capacity management

The disciplines of performance engineering and capacity management are not obsolete as a result of the introduction of the cloud and of services provided through the cloud. On the contrary, cloud deployments are more dependent on having professional performance engineering and capacity management in place.

Five of the 10 obstacles and opportunities for cloud computing are related to QoS aspects, such as availability, performance, capacity, or scalability.[3]

- Obstacle 1, “Availability of service”, presents availability risks for cloud computing as a result of factors, such as programming errors, overload of common services, or Distributed Denial of Service (DDoS) attacks.

- Obstacle 4, “Data transfer bottlenecks”, presents the growing data intensity of applications and how this affects data transfer rates and costs in the cloud.
- Obstacle 5, “Performance unpredictably”, presents performance risks caused by factors, such as inefficiencies in I/O sharing and by high performance computing.
- Obstacle 6, “Scalable storage”, presents the difficulties of applying cloud computing to solutions requiring highly scalable persistent storage.
- Obstacle 8, “Scaling quickly”, presents the difficulties of quickly scaling up and down in response to load without violating SLAs.

Potential cloud solutions to overcome these obstacles need to be assessed for their feasibility in real-life situations.

Performance engineers who are advising cloud computing users and cloud computing providers need to gain a thorough understanding of the technical transactions that underlie cloud services. The degree to which cloud services can meet agreed service-level requirements for availability, performance, and scalability can be estimated by using performance modeling techniques. Potential performance anti-patterns can be detected before they happen.

The automatic provisioning and usage-based costing facilities, two of the major features of cloud computing, rely heavily on fine-grained capacity management. Until more sophisticated tooling for automated monitoring, data collection, analysis, and forecasting are in place, capacity management is more opportune than ever.

But even in an ideal setting where capacity management is fully automated, cloud computing users still have to analyze their demand for capacity and their requirements for QoS. In their negotiations with cloud computing providers, users have to accurately formulate their service-level requirements.

It is to be expected that the work of the capacity manager shifts toward translating the capacity demands of cloud computing users into capacity plans on which service-level agreements (SLAs) with cloud computing providers can be based.

Service-level agreements in the cloud

An SLA is a reciprocal agreement between the provider of an IT service and the consumer of that service about the level of service, or QoS, to be delivered. Examples of QoS attributes are performance, which can cover both response time and throughput, availability, security, and so on. The notion of reciprocity is important. It is true that SLAs might address the manner in which the service is supplied (for example, achieving specified response time or availability targets). However, a complete SLA also addresses expected demands for services (for example, transaction mix and volumes).

In the classic data center environment, SLAs are formalized as the service provider organization has sufficient experience with the application in question to understand its performance and capacity characteristics relative to the delivery capabilities of the provider. Cloud delivery models can create multiple levels of SLA dependencies, as shown in Figure 4 on page 12.

[3] *Above the clouds: A Berkeley View of Cloud Computing*, Michael Armbrust, et al: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf

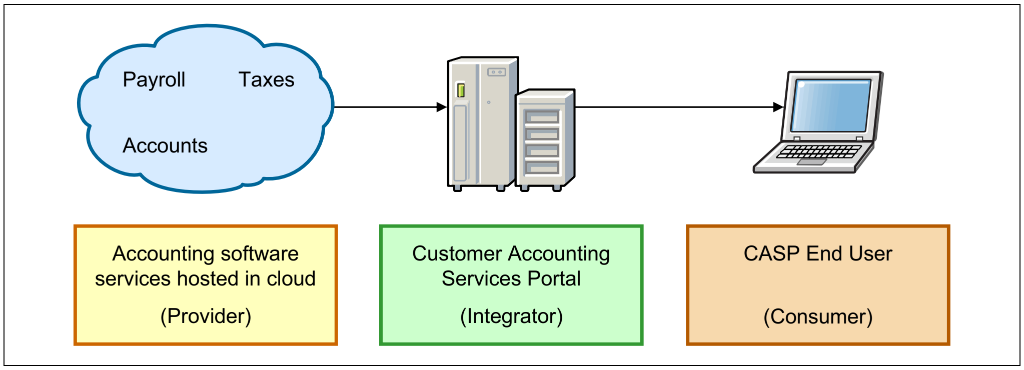

Figure 4 Cloud SLAs

In this example, the integrator developed a commercial offering that allows small businesses to obtain accounting services from the cloud through a portal. Consumers might choose to subscribe to one or many accounting services through the portal. The portal in turn funnels requests from multiple consumers to the cloud-hosted accounting services of the provider. This scenario creates the opportunity for at least two levels of SLAs:

- The level of service to be provided by the integrator to the consumer
- The level of service to be provided by the provider to the integrator

To guarantee a level of service, the provider must consider the demands made upon the cloud by this integrator (and possibly others) from a volume and processing perspective, and the architecture, design, and infrastructure used to implement the cloud solution. During the assembly of the cloud solution, estimation, testing, and measurement activities are required on the part of the integrator, provider, or both to finalize cloud SLAs. Ongoing monitoring also is required to ensure that the cloud SLAs continue to be met after deployment.

Typically, the integrator is responsible for guaranteeing a level of service to one or more consumer groups. However, except for the user connectivity and portal technologies used, in this example the integrator is dependent on the cloud SLAs being met by the provider to meet the portal SLAs guaranteed to the consumer. This concept of dependent SLAs is not new to the IT industry as it first emerged when enterprises started outsourcing network management to third parties.

Although the service-level responsibilities flow left to right in Figure 4, requirements information must flow right to left. The portal integrator must collect or estimate usage information (volumes, transaction mixes, usage patterns, and required responsiveness) from the consumer community. Likewise, the cloud provider must understand the demands on the cloud, which can come from multiple integrators or consumers.

As the SLAs are defined, the interactions between cloud components and their effects on performance must be considered both during solution integration (to confirm that SLAs can be met) and after solution deployment (to confirm that SLAs are met). In our previous example, if some of the SLAs also included response time, interactions such as those shown in Figure 5 on page 13 must be considered by the integrator when devising performance budgets. Interactions also must be considered by the provider when setting up performance monitoring.

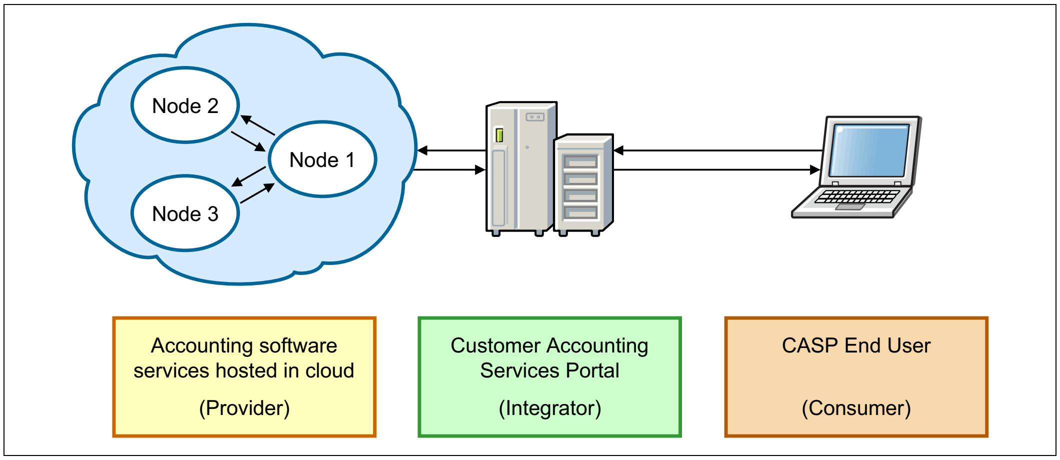

Figure 5 Cloud interactions that affect SLAs

The classic principles of performance engineering apply to the end-to-end cloud solution. However, the responsibility for applying performance engineering methods to the solution must be divided based on the respective SLA responsibilities of the provider and the integrator.

Service-level requirements for performance usually address acceptable response times or elapsed times, requirements for throughput, and capacity requirements; for example, for network bandwidth or storage.

Service-level requirements for availability are related to performance, because lack of availability can be considered as the worst possible performance. Also, required availability within the service window must be met.

To ensure stable performance in the long term, requirements for scalability and reliability are captured in the SLA.

Table 1 on page 14 shows each of the service-level categories and the key performance indicators (KPIs) that are used to assess SLA attainment.

Table 1 Cloud service-level categories and key performance Indicators

Service-level category	KPIs	Definition	Unit of measurement
Availability	Service window	Time window within which KPIs are measured	Time range
	Service/System availability	Percentage of time that service or system is available	%
	MTBF	Meantime between failure	Time units
	MTTR	Meantime to repair	Time units
Performance	Response time	Response time for composite or atomic service	Seconds
	Elapsed time	Completion time for a batch or background task	Time units
	Throughput	Number of transactions or requests processed per specified unit of time	Transaction or request count
Capacity	Bandwidth	Bandwidth of the connection that supports a service	bps
	Processor speed	Clock speed of a processor	MHz
	Storage capacity	Capacity of a temporary or persistent storage medium, such as RAM, SAN, disk, or tape	GB
Reliability	Service/System reliability	Probability that service or system is working flawlessly over time	%
Scalability	Service/System scalability	Degree to which the service or system can support a defined growth scenario	Yes/No, or description of scalability upper limit

The cloud delivery model determines what performance indicators must be covered in an SLA:

- If the delivery model is Business Process as a Service, the service consumer is likely to have requirements regarding response times, throughput, availability, reliability, and scalability of business processes. The service consumer leaves the derived capacity requirements to the provider.
- If the delivery model is SaaS or PaaS, the service consumer is likely to have requirements for response times, throughput, availability, reliability, and scalability of transactions supported by the software. The service consumer leaves the derived capacity requirements to the provider.
- If the delivery model is IaaS, however, the service consumer is likely to have requirements for throughput, capacity, availability, reliability, and scalability of infrastructure components. In this instance, the services provided on top of the cloud infrastructure are not in the scope of the contract. Therefore, service response times or elapsed times cannot be addressed in the SLA.

The cloud deployment model determines how the SLA is offered to the consumer:

- If the cloud deployment model is *public*, the cloud provider might choose to offer a fixed service level or a simple set of service levels, ranging from, for example, bronze to gold, where better service is offered to consumers who are willing to pay more.
- If the cloud deployment model is *private*, the cloud consumer and provider might choose to agree upon a custom SLA.

- If the cloud deployment model is *hybrid*, the SLA offering might be determined by a combination of the two methods.

In addition to making arrangements on the performance indicators that are monitored, the SLA must address the financial implications of cloud usage. Cloud capacity is used simultaneously by multiple consumers and all consumers pay for their own usage. As a result, the cloud is equipped with usage-based costing capabilities. The consumer, provider, and integrator must provide data to ensure that these SLAs are in compliance.

The role of the Service Catalog in the cloud

The service catalog documents the services available for delivery through the cloud. The service catalog consists of a list of standard services that can be requested from and provisioned by the provider, along with a standard price list for the services offered. The cost is typically based on QoS or usage metrics, such as:

- Level of response
- Number of accesses per time unit
- Units of system processing or application usage time
- Number of named or concurrent users
- Any other terms and conditions for usage of the services in the catalog

The service catalog is part of a cloud management service that provides the user interface, business support systems that provide the metering, billing and reporting and operational support systems that provide the system provisioning, monitoring, and support. For example, the IBM Tivoli® product family includes tools, such as Tivoli Service Automation Manager, Tivoli Provisioning Manager, and IBM Tivoli Monitoring, which provide service catalog features as a central part of cloud management.

By documenting exactly the services the cloud has available for an individual consumer, the level of each QoS attribute available and requested, the price at which the service might be obtained, and the terms and conditions of using the services on the cloud, the service catalog holds information uniquely defining the performance and capacity commitments within the contractual SLA for each consumer. Ensuring that these contractual SLA commitments can be met from a performance and capacity perspective by the cloud computing infrastructure requires a structured method of performance engineering and capacity management that we introduce in the next section.

As shown in Figure 6 on page 16, the service catalog is built by the cloud administrator and integrates the infrastructure, platform, and software services that exist within the cloud.

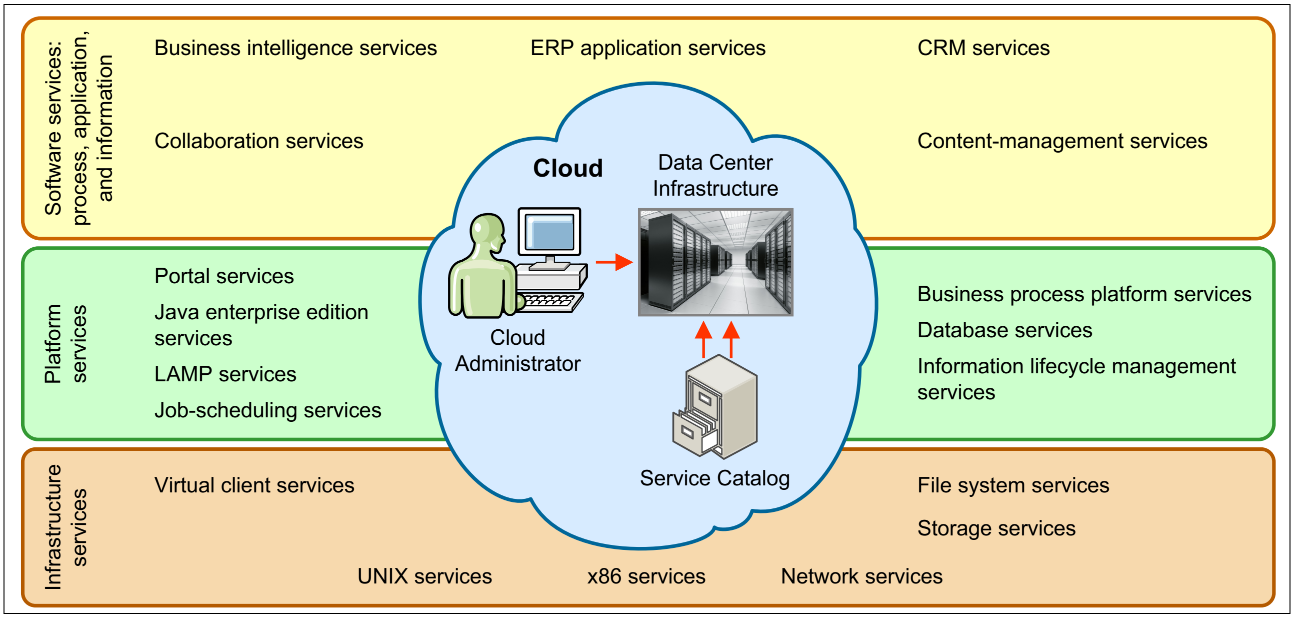

Figure 6 Service Catalog in the cloud computing infrastructure

As services are added to the catalog, they are published and made available to the services consumers. This availability puts the service catalog at the core of the metering, billing, reporting, and operational support systems of the computer cloud.[4]

Activities to ensure SLA compliance

Specific profile information is required from the consumer about how the cloud is intended to be used to ensure SLA compliance:

- A profile of the user population that uses the services within the cloud, such as the number of users, user types, and the number of concurrent users and hence service requests expected at the peak times.
- The number of Business Service executions or transactions expected to be performed at the peak times. This information can be expressed as a projection or as actual numbers aggregated daily, monthly, or annually and is used in calculating the peak hour volumes.
- If requesting a Test Cloud, a profile of the test environments that details the required system components, how many environments are required, the number of users that are accessing them, and the times when these test environments is required.
- Security information as it relates to systems hardware components, data, and application system components that include the requirements on shared versus dedicated access.

The provider needs to assess the requirements of the consumer and develop a solution that meets those requirements over the course of the project life cycle. From a performance and capacity perspective, this solution includes the following considerations:

- Creation of the appropriate test environments from the virtualized pool that meets the required hours of operations.
- Execution of performance tests to establish the behavior of the application and the footprint it has in the cloud production environment.

[4] G. Breiter and M. Behrendt, *Life cycle and characteristics of services in the world of cloud computing*, IBM Journal of Research and Development, 2009.

- Estimate of planned resources required within the Cloud along with performance objectives, which are confirmed during the lifecycle process in the performance model initially and, ultimately, in the deployment model.
- Determination of the system resources required in production based on performance modeling. Testing must be identified in terms of a graduated scale based on a baseline resource unit and the required number of additional resource units that coincides with horizontal increments of resources at each tier. The deployment model contains sizing and configuration information that is required when the application is migrated to the production cloud. The information also provides initial confirmation on the resource and capacity selections from the services catalog offerings.

The integrator must confirm that all system components required for the solution are included. The integrator also must verify that the proposed solution is feasible and meets the performance and service-level objectives:

- Proposed test systems environments satisfy the needs of the requester
- Contents of the Deployment Model and the proposed solution for Production are validated
- Activities for the production migration are coordinated

The next section of this paper outlines several use case scenarios to clarify real business situations and highlights when performance is considered as a part of the cloud solution strategy.

Cloud computing scenarios and their performance implications

Use case scenarios help to clarify real business situations and highlight when performance considerations are part of the cloud solution strategy. Scenarios involving private, public, and hybrid cloud deployment models are described next.

Scenario 1: Private cloud deployment

An Internet company must implement a strategic new business model without increasing operational costs. The current manual processes led to unexpected downtimes and system outages. The company is anticipating storage capacity increases from hundreds to thousands of servers to accommodate anticipated business growth, volume increases, and automation.

Solution strategy

The solution strategy is to introduce cloud gradually and maintain an open and adaptable cloud management platform. The guiding principles are to begin with an internal and private cloud, grow into a public cloud, and then generate an open and adaptable cloud management platform with support of heterogeneous hardware.

Solution details

The solution strategy includes the following details:

- Implement Tivoli Service Automation Manager to provide an end-to-end service management layer, and automated deployment and management along with a virtualized infrastructure.
- Deploy virtual servers with Linux, Apache, Middleware, and Perl (LAMP) to a specified hardware, operating system (OS), network, and storage configuration.
- Provide virtualization hypervisor on bare-metal hardware.
- Manage and scale infrastructure and platform services (virtual servers or LAMP).

Performance considerations

There are several performance considerations that must be addressed as a solution strategy is developed and implemented:

- Consider service-level agreements and workload optimization to manage capacity and prevent bottlenecks from occurring within the private cloud.
- Organize workloads so that they can scale to support anticipated and unanticipated growth by using techniques that optimize use of storage, such as virtualization.
- Consider the skill set of teams and of the systems management capabilities to help drive service-level performance.
- Consider security implications.

Scenario 2: Public cloud deployment

A sales organization wants to reduce the management and support costs associated with remote notebook users. New sales team members need to be added quickly with a standard suite of office productivity and sales tools.

Solution strategy

The solution strategy is to contract with a desktop cloud vendor to provide the following services:

- Standard desktop image
- Standard tools plus custom software built into the image
- Desktop data storage options

This solution is a pre-priced, prepackaged subscription service with some limited customization to keep costs down. The total cost is based on metered usage, so the sales organization pays only for the time they use.

Solution details

This solution strategy includes the following details:

- The desktop image is accessible from any remote device over Remote Desktop Protocol (RDP) or Independent Computing Architecture (ICA) protocol.
- The desktop image is accessible from corporate network or public Internet.
- The public Internet access connects to the desktop cloud through a dedicated security appliance to ensure data security for corporate data.
- The desktop cloud service, infrastructure, and availability are managed by the provider.
- The standard desktop image can be modified by the consumer.

Performance considerations

There are several performance considerations that must be addressed as this solution strategy is developed and implemented:

- The performance of remote desktops is dependent on the software installed and used on the remote desktop.
- The performance of remote desktops is highly dependent on the types of data remote users work with on the remote desktop.
- The performance of remote desktops is dependent upon the network connection bandwidth and latency between the consumer point of entry and the location of the desktop cloud.

- The client needs to consider the type of device their users use to access the remote desktop (desktop, notebook, or thin client device). Each device has different impacts on performance, security, and usability.
- SLAs might be difficult to define and enforce for a public cloud scenario. Clients with many application service requirements for performance and capacity might best be directed to the private or hybrid cloud scenario.

Scenario 3: Hybrid cloud deployment

An application development company is investigating new development tools and wants to evaluate those tools to see whether the new tools fit well in their existing toolset. They need quick turnaround, and do not want to buy servers and software licenses that might never be used in their actual development environment.

Solution strategy

The solution strategy is to contract with a development cloud service provider.

The development cloud is hosted by the service provider and includes a suite of images with development and test tools preinstalled and configured that the client can access.

The client also selects the virtual private network (VPN) option so that the image instances that they create have IP addresses, host names, and access security rules. These image instances allow the client to access the instances as though the instances were in their own data center.

Solution details

This solution strategy includes the following details:

- The consumer requests instances by using the self-service portal, then configures the development tools to integrate with their existing toolset.
- Total cost is based on metered usage of the development tool license, so the sales organization pays only for the time they use.
- If the integration is successful, the consumer can continue to use the development tool through the development cloud instance, or purchase their own license and install the tool in their own data center infrastructure.

Performance considerations

There are several performance considerations that must be addressed as this solution strategy is developed and implemented:

- The performance is highly dependent on the bandwidth and latency between the service provider data center and the consumer data center.
- Performance is impacted by the following factors:
  - Authentication access integration between the development cloud instances and the consumer toolset
  - Whether the development cloud instance must be authenticated by an enterprise authentication directory service in the consumer data center
- Some development tools require access to shared databases for code sharing. If these databases are remote from the development tool, this distance affects performance.

- If the consumer is considering moving the development tool from the development cloud into its standard toolset, they need to verify that the performance is adequate for daily usage by a larger user base than the evaluation user generated.
- SLAs need to consider the level of service between the cloud provider and the application integrator. The agreements also need to consider the level of service between the application integrator and the consumer of the cloud service.

For more information about the hybrid development cloud, see the IBM SmartCloud™ website:

http://www.ibm.com/ibm/cloud/smart_business

Solution considerations for the cloud

Many cloud computing solutions are available from different vendors. A common theme across those solutions is that the type of service that is offered highlights the products and services in which the cloud computing vendor specializes. The following cloud solutions are offered by different vendors:

- Oracle/Sun and Hewlett-Packard (HP) focus on hardware and services to enable cloud construction.
- Google offers Google Apps, a pure software-as-a-service solution.
- HP highlights its system management strengths as a software-as-a-service solution.
- IBM and Amazon offer a mixture of platform-as-a-service cloud solutions and complete turnkey public, private, or hybrid cloud computing packages.

The technical components, business focuses, and marketing directions that result from these cloud computing vendor decisions directly affect the performance implications of the underlying cloud computing solution and need to be analyzed and evaluated appropriately.

Cloud offerings from IBM

IBM SmartCloud Enterprise+ (SCE+) is a fully managed, security-rich, and production-ready cloud environment designed to ensure enterprise-class performance and availability. SCE+ offers complete governance, administration and management control along with service-level agreements (SLAs) to align your specific business and usage requirements. Multiple security and isolation options built into the virtual infrastructure and network keep this cloud separate from other cloud environments. [5]

In addition to SCE+, the following offerings are part of the IBM SmartBusiness cloud services overview.

CloudBurst

IBM CloudBurst® V1.2 is a prepackaged and self-contained service delivery platform that can be easily and quickly implemented in a data center environment. IBM CloudBurst V1.2 is positioned for enterprise clients who want to get started with a private cloud computing model. It is ideal for those clients who are looking for a solution to complement their existing IT infrastructure and assistance in efforts to realize a return on their cloud investment in a challenging economic environment.[6] Table 2 on page 21 lists the features of CloudBurst.

[5] For more information, see: http://www-935.ibm.com/services/us/en/managed-cloud-hosting/

[6] IBM United States Software Announcement 209-402, 2009

Table 2 CloudBurst features

Deployment model	Private cloud
Roles focus	Provider
Performance considerations	► Workload organization to support anticipated and unanticipated growth ► Consider the skill set of teams to support the CloudBurst platform ► Plan systems management capabilities and integration with other infrastructure ► There might be security integration implications

Smart Business Development and Test on the IBM Cloud

The IBM Smart Business Development and Test on the IBM Cloud is a dynamically provisioned and scaled runtime environment that provides everything needed to develop and test application code that highlights the software offerings of IBM. This environment includes tools to configure and manage the dynamic execution environment, an integrated development environment (IDE) that facilitates the direct use of the execution environment. The IDE also offers tools for building and testing that can use the execution environment. Table 3 lists the features of Smart Business Development and Test on the IBM Cloud.

Table 3 Smart Business Development and Test on the IBM Cloud features

Deployment model	Public or hybrid cloud
Roles focus	Provider
Performance considerations	► Performance is dependent on bandwidth and latency between the data centers ► Authentication between cloud and client tools affects performance ► Development tools might need access to shared databases ► Need client to verify performance for daily usage by developers

Smart Business Desktop Cloud

IBM Smart Business Desktop Cloud provides access to applications, information, and resources through thin clients or any other Internet-connected device. The Smart Business Desktop Cloud delivers a resilient, efficient, standards-based IT infrastructure for almost any traditional desktop application. This cloud employs portal, thin client, messaging, and security technologies delivered through a single, consistent framework. All that is needed is a thin client device or PC capable of running an Internet browser and Java™. Table 4 lists the features of Smart Business Desktop Cloud.

Table 4 Smart Business Desktop Cloud features

Deployment model	Public cloud
Roles focus	Consumer
Performance considerations	► Performance is dependent on the software on the remote desktop ► Performance is dependent on data users on the remote desktop ► Performance is dependent on the network connection bandwidth and latency ► The access device (PC or notebook) has different performance impacts

Tivoli Service Automation Manager

Tivoli Service Automation Manager provides the capability to request, deliver, and manage IT services. It is a strategic cross-IBM solution for the operational support systems that is necessary to help enterprise data centers benefit from cloud computing. Tivoli Service Automation Manager is designed to enable faster IT response and delivery capabilities, and help lower IT operational costs. Table 5 lists the Tivoli Service Automation Manager features.

Table 5 Tivoli Service Automation Manager features

Deployment model	Public or hybrid cloud
Roles focus	Developer and integrator
Performance considerations	► Building block for private clouds ► Workflows are a key component of performance ► Disk and networks must be optimized for workload performance

Cloud offerings from Amazon

Amazon Elastic Compute Cloud (Amazon EC2) is a web service that provides resizable compute capacity in the cloud. It is designed to make web-scale computing easier for developers. Amazon EC2 allows users to obtain and configure capacity with minimal friction. It provides complete control of resources that run on the proven computing environment of Amazon. Amazon EC2 reduces the time required to obtain and boot new server instances to minutes. Amazon EC2 allows users to pay only for the capacity that is used. Amazon EC2 provides developers the tools to build failure-resilient applications and isolate the applications from common failure scenarios.[7] Table 6 lists the features of Amazon EC2.

Table 6 Amazon EC2 features

Deployment model	Public cloud
Roles focus	Provider and integrator
Performance considerations	► Performance is dependent on bandwidth and latency between the data centers ► Authentication between cloud and consumer tools affects performance ► Development tools might need access to shared databases ► Consumer must verify performance for daily usage by developers

Cloud offerings from Oracle/Sun

Oracle/Sun combines open systems, software, and architectural expertise to build clouds and maximize their capabilities.

Responding to the demand for interoperability, Oracle/Sun is pursuing its vision of open, interoperable clouds in four key areas:

- Open standards-based software
- Open systems hardware for interoperability
- Microelectronics with multithreading, multicore computing to enable higher compute densities
- Professional services and systems integration to use the benefits of cloud computing[8]

7 http://aws.amazon.com/ec2/
8 http://www.oracle.com/us/products/consulting/cloud-computing-services/index.html

Table 7 lists the features of Oracle/Sun cloud.

Table 7 Oracle/Sun cloud features

Deployment model	Private cloud
Roles focus	Provider and integrator
Performance considerations	► Sun offerings are building blocks for private clouds ► Performance must be engineered into the cloud ► Open hardware and software increases flexibility, but might not increase performance

Cloud offerings from Google

Web-based messaging and collaboration applications from Google require no hardware or software and need minimal administration, thus creating time and cost savings for businesses. Clients can use mobile email, calendar, and information management access, with several options for mobile access to their information. Google guarantees that Google Apps is available at least 99.9% of the time, and each employee receives 25 GB for email storage so they can keep important messages and find them instantly with built-in Google search[9]. Table 8 lists the features of Google cloud.

Table 8 Google cloud features

Deployment model	Public cloud
Roles focus	Consumer
Performance considerations	► Performance is dependent on data the clients use in Google Apps ► Performance is dependent on the network connection bandwidth and latency ► The access device (PC or notebook) has different performance impacts

Cloud offerings from Hewlett-Packard

HP Cloud Assure consists of HP services and software, including HP Application Security Center, HP Performance Center, and HP Business Availability Center. HP also provides engineers to perform security scans, execute performance tests, and deploy availability monitoring. HP Cloud Assure helps clients validate security by scanning networks, operating systems, middleware layers, and web applications, and performing penetration testing. HP Cloud Assure also ensures that cloud services meet user bandwidth and connectivity requirements and satisfy service-level agreements and availability. Cloud-based applications are monitored to isolate potential problems, identify root causes with user environments and business processes, and analyze performance issues.[10] Table 9 on page 24 lists the features of Hewlett-Packard cloud.

[9] http://www.google.com/enterprise/apps/business/benefits.html

[10] *HP Unveils "Cloud Assure" to Drive Business Adoption of Cloud Services*: http://www.hp.com/hpinfo/newsroom/press/2009/090331xa.html

Table 9 Hewlett-Packard cloud features

Deployment model	Private cloud
Roles focus	Provider and integrator
Performance considerations	► HP offerings are building blocks for private clouds ► Performance needs to be engineered into the cloud ► HP performance and systems management tools are a strong part of offering

Conclusion

Cloud computing provides an efficient, scalable, and cost-effective way for organizations to deliver business and consumer IT services over the Internet. Various cloud computing models are available, providing both solid support for core business functions and the flexibility to deliver new services. Performance considerations must be addressed throughout the lifecycle of the cloud-based solutions if they are to deliver QoS levels that are typical of well-constructed enterprise applications. The responsibility for undertaking those performance management roles is shared by the consumer, provider, and integrator.

There are clear advantages to be gained today by using cloud technology for specific purposes and many organizations already are realizing the benefits. For example, solutions, such as Test Cloud enable test environments to be provisioned quickly reducing the time frame for deploying new applications into production. Those environments might support functional testing and even some level of non-functional testing. However, the cloud has yet to mature sufficiently to be the solution for all non-functional testing needs.

This paper provides an overview of cloud computing and its benefits, presented the different cloud computing perspectives, highlighted performance engineering interfaces with cloud solutions, and detailed performance considerations for the cloud. Usage scenarios were shown for clarification and elaboration of real-life cases in which performance considerations are part of the cloud solution strategy. Examples of offerings were provided to assist in the development of performing cloud systems.

Large enterprises that deploy their mission-critical applications over the cloud are still in the future. However, the current and emergent trends within our industry are fueling the demand for truly global business services that are agile and responsive to the needs of the marketplace and can be invoked from anywhere at anytime. The consequence is that the potential demand for successful services takes a quantum leap beyond the demands of more traditional commercial IT systems. This scenario has significant implications for how this new breed of services must be built and deployed. QoS is already becoming an increasingly important differentiator in the global market. The performance and scalability of the IT systems that underpin those business services are vital in enabling organizations to meet demands.

The challenge posed for Performance Engineering is substantial and the discipline needs to continue to innovate to meet that challenge. Industry watchers are predicting that Performance Engineering becomes even more critical to the success of the IT industry.

The team who wrote this paper

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Elisabeth Stahl is Chief Technical Strategist for the IBM Systems and Technology Group and has been working in systems performance for over 25 years. She holds a Bachelor of Arts degree in Mathematics from the University of Pennsylvania and an Master of Business Administration degree from NYU. Elisabeth is an Open Group Distinguished IT Specialist, an IEEE Senior Member, and member of the IBM Academy of Technology.

Lydia Duijvestijn is a senior IT Architect and Performance Engineer in the Global Business Services organization of IBM in the Netherlands. She is one of the co-leaders of the IBM worldwide Performance & Capacity Community of Practice. She led a large number of engagements with clients in the area of design for performance, performance testing, and performance troubleshooting. Two of the most recent performance-related international engagements were with the largest bank of Russia, based in Moscow, and a large bank in the Nordics, based in Stockholm. She has spoken at IBM conferences about subjects related to IT architecture and performance.

Avin Fernandes is a certified IT Technology Architect with over 25 years of experience. He leads the Performance Engineering Center Of Excellence for IBM Canada as part of the Performance and Testing Practice. He also is a core team member of the worldwide Performance & Capacity Community of Practice, providing leadership both within IBM and to external academia in support of student research and thesis development. As part of this program, Avin has led projects in SOA Performance Engineering and System Capabilities Assessments. He has various roles at IBM, including chief architect for the development and implementation of Air Canada's existing reservation system RESIII and more recently as the Test Manager for Rogers Enterprise Cloud PaaS solution.

Pamela K. Isom is a Global, Enterprise Principal Consultant at Dell Inc., where she leads very large cloud strategy and next generation data center engagements. Prior to Dell, Pamela was executive architect in IBM® Global Business Services® and a chief architect of Complex Cloud Integration and Enterprise Application Delivery in the Application Innovation Services, Interactive Solutions Practice. While at IBM, Pam was a member of the IBM Academy of Technology where she led smarter cities and cloud computing in highly regulated environment initiatives. She also managed the GBS/AIS patent board, having filed and received issuance of several patents with the U. S. Patent Attorney's office.

Dave Jewell is a Performance Engineer for the Systems Engineering, Architecture & Test organization in IBM Global Business Services. Since joining IBM in 1978, Dave has held a number of positions in application development, testing, performance analysis, and performance engineering. He also has been involved in performance work for clients in multiple industries over the past 20 years.

Martin Jowett is a Distinguished Engineer and Executive IT Architect specializing in the area of Performance Engineering and System Architecture and Design. Over the last 15 years, he has worked as Lead Performance Architect on many of the largest and most challenging of the IBM UK Systems Integration Programs. He has also worked as a performance consultant to many of the most significant IBM accounts, both within and outside the UK. He is currently Lead Performance Architect within Global Business Services IBM UK. He is a recognized leader across the worldwide IBM performance community, is currently co-lead of the worldwide Performance & Capacity Community of Practice, and global owner of the Architecting for Performance Class. He has played a leadership role in both the development and deployment of content related to Performance Engineering within the IBM Global

Services Method and accompanying education courses. He is a member of the IBM Academy of Technology.

Todd R. Stockslager is the Test Services Architect for the IBM Global Process Services Development team. Since joining IBM in 1997, he has provided performance testing and performance engineering support for several client-facing and internal IBM development efforts. In his current role, he implemented the performance test team, tooling, and processes for the IBM Global Process Services Testing Service. In this role, he has managed performance testing of the initial IBM Smart Cloud Enterprise offering and the IBM Smart Business Desktop Cloud offering. He is currently providing test services architecture support for infrastructure and performance testing to GPS Development clients in the Financial Services sector.

Acknowledgements

The authors would like to thank Pirooz Joodi, Damian Towler, Lewis Grizzle, Stefan Pappe and Ann Dowling for reading draft versions of this paper and providing many valuable comments.

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and client satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

http://www.ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks publications

- Find us on Facebook:

  http://www.facebook.com/IBMRedbooks

- Follow us on Twitter:

  http://twitter.com/ibmredbooks

- Look for us on LinkedIn:

  http://www.linkedin.com/groups?home=&gid=2130806

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

  https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm

- Stay current on recent Redbooks publications with RSS Feeds:

  http://www.redbooks.ibm.com/rss.html

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION “AS IS” WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

This document REDP-4875-00 was created or updated on April 3, 2013.

Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
  ibm.com/redbooks
- Send your comments in an email to:
  redbooks@us.ibm.com
- Mail your comments to:
  IBM Corporation, International Technical Support Organization
  Dept. HYTD Mail Station P099
  2455 South Road
  Poughkeepsie, NY 12601-5400 U.S.A.

IBM®

Redpaper™

Trademarks