

# 11

---

## *Cloud Service Providers*

---

---

### **Learning Objectives**

The main objective of this chapter is to provide an overview of different cloud service providers. After reading this chapter, you will

- Know about different companies that support cloud computing
- Understand open source/proprietary tools offered by the companies
- Know cloud services offered by the companies
- Understand the features and available architecture of different tools

---

### **Preamble**

This chapter provides an overview of cloud services offered by various companies. We begin with the introduction to cloud services. Subsequent sections talk about companies like Amazon, Microsoft, Google, EMC, Salesforce, and IBM that provide various tools and services in order to give cloud support. Each section briefly describes cloud features supported by these companies. It also gives an idea of tools and technologies adapted by companies in order to provide services to the users. In this chapter, we focus on giving readers brief information about various tools and technologies provided by different companies. After reading this chapter, the reader will be able to distinguish different services provided by various companies and make appropriate choices as per the requirement.

## **11.1 Introduction**

Cloud computing is one of the most popular buzzwords used these days. It is the upcoming technology provisioning resources to the consumers in the form of different services like software, infrastructure, platform, and security. Services are made available to users on demand via the Internet from a cloud computing provider's servers as opposed to being provided from a company's own on-premise servers. Cloud services are designed to provide easy, scalable access to applications, resources, and services and are fully managed by a cloud service provider. A cloud service can dynamically scale to meet the needs of its users, and because the service provider supplies the hardware and software necessary for the service, there is no need for a company to provision or deploy its own resources or allocate information technology (IT) staff to manage the service. Examples of cloud services include online data storage and backup solutions, web-based e-mail services, hosted office suites and document collaboration services, database processing, and managed technical support services.

Cloud services can be broadly classified into three types: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS). With growing technologies, many more services are emerging in this field, such as Security as a Service (SeaaS), Knowledge as a Service, and Data Analytics as a Service.

Many companies have come forward to adapt the cloud environment and ensure that the users as well as the companies benefit from this. Amazon, Microsoft, Google, Yahoo, EMC, Salesforce, Oracle, IBM, and many more companies provide various tools and services in order to give cloud support for their customers.

---

## **11.2 EMC**

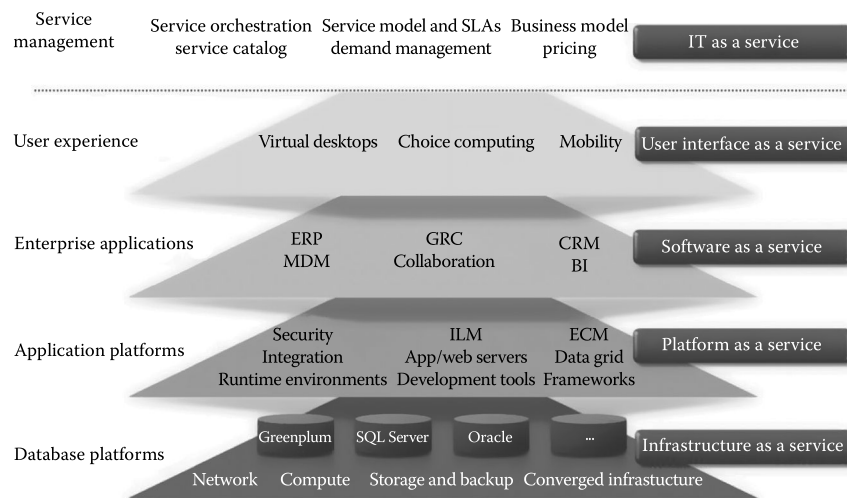
EMC is one of the leading global enterprises that require dynamic scalability and infrastructure agility to meet changing applications as well as business needs. EMC chose cloud computing as the ideal solution to reduce the complexity and optimize the infrastructure. Offering Information Technology as a Service (ITaaS) reduces the energy consumption through resource sharing.

### **11.2.1 EMC IT**

Virtualization is the main concept behind the success of EMC IT. By virtualizing the infrastructure, allocation of the resources on demand is possible. This also helps to increase efficiency and resource utilization.

EMC IT provides its business process units with IaaS, PaaS, and SaaS. Figure 11.1 gives an overview of the services offered by EMC, which are explained in the following:

1. IaaS offers EMC business units the ability to provision infrastructure components such as network, storage, computing, and operating systems individually or as integrated services.
2. PaaS provides the secure application and information frameworks on top of application server, web server, database, unstructured content management, and security components as a service to business units from which to develop solutions. EMC IT offers database platforms (Oracle Database as a Service, SQL Server as a Service, Greenplum as a Service) and application platforms (application development, Enterprise Content Management as a Service, Information Cycle Management as a Service, Security PaaS, Integration as a Service) for the purpose of development.
3. SaaS provides applications and tools in a service model for business enablement. EMC IT brought together several existing business solutions under the unified architecture named as Business Intelligence as a Service. It also offers Enterprise Resource Planning (ERP) and Customer Relationship Management (CRM) as a Service.
4. User Interface as a Service (UIaaS) provisions user and interface experience, rather than provisioning the actual device used.



**FIGURE 11.1** Cloud services by EMC. (Adapted from EMC IT’s journey to the private cloud, applications and the cloud experience, White Paper-EMC.)

### 11.2.2 Captiva Cloud Toolkit

EMC offers a tool called *Captiva Cloud Toolkit* to help in the development of softwares. EMC Captiva Cloud Toolkit is a Software Development Kit (SDK) comprised of modules that help web application developers to quickly add scanning and imaging functionality directly to their web-based business applications. It is ideal for document capture vendors, commercial software developers, and enterprises that want to create custom web-based applications that are fully scan enabled, complimenting their business solution offerings.

Using Captiva Cloud Toolkit, developers can quickly create a working scan-enabled web-based business application in as early as 1 week. As a result, time to market is shortened and development, testing, and support costs are greatly reduced. Also, the enterprise's return on investment is quickly achieved, and its ability to compete in an increasingly competitive distributed document capture market is accelerated.

There are a few modules that are commonly used in most of the process development. These are basic modules that import images from various sources like fax, e-mail, or scanner or from any repository. A few of these modules are as follows:

1. *Scan*: Scanning is importing activity of documents into Captiva from a scanner. Basically, scanning happens at page level to bring images page by page into Captiva. Scanning is the entry point to Captiva where one can import any kind of document like pdf, tiff, and jpg.
2. *MDW*: Multi Directory Watch is another entry point to Captiva. MDW can be pointed to any folder/repository from where Captiva could import documents directly. MDW is very useful if business is getting documents in the form of a soft copy, for example, as an attached file in an e-mail. MDW also acts as a scan module except it does not interlock with the scanner.
3. *IE*: Image enhancement is a kind of filter or repairing tool for images that are not clear. It enhances the image quality, so it could be processed easily through Captiva. One can configure IE as per business requirement and images being received. The functionalities of IE are deskew, noise removal, etc.
4. *Index*: Indexing is a data capturing activity in Captiva through which one can capture key data from various fields. For example, if bank form is being processed, the A/C no. and sort code could be the indexing field. Indexing could be added as per requirement of business. A validation field could be added to avoid unwanted data entry while indexing any document.

5. *Export*: Export is the exit point of Captiva where images/data are sent to various repositories like file, net, document, or data. The exported data are used for business requirements of various business divisions. For example, if we are capturing the A/C no. and sort code for a bank application, this could be mapped to any department where it is needed.
6. *Multi*: Multi is the last process in Captiva to delete batches that have gone through all modules and exported value successfully. Multi could be configured as per need of business. In the case when it is required to take a backup of batches, this module could be avoided.

The previously mentioned modules are very basic modules of Captiva for indexing and exporting. But for more flexibility and automation, dispatcher is used, which is more accurate to capture data.

---

## 11.3 Google

Google is one among the leading cloud providers that offer secure storage of user's data. It provides cloud platform, app engine, cloud print, cloud connect, and many more features that are scalable, reliable, as well as secure. Google offers many of these services for free or at a minimum cost making it user friendly.

### 11.3.1 Cloud Platform

Google Cloud Platform enables developers to build, test, and deploy applications on Google's highly scalable and reliable infrastructure. Google has one of the largest and most advanced networks across the globe. Software infrastructures such as MapReduce, BigTable, and Dremel are the innovations for industrial development.

Google Cloud Platform includes virtual machines, block storage, NoSQL datastore, and big data analytics. It provides a range of storage services that allow easy maintenance and quick access of user's data. The cloud platform offers a fully managed platform as well as flexible virtual machines allowing the user to choose as per the requirements. Google also provides easy integration of user's application within the cloud platform.

Applications hosted on the cloud platform can automatically scale up to handle the most demanding workloads and scale down when traffic subsides. The cloud platform is designed to scale like Google's own products, even when there is a huge traffic spike. Managed services such as App Engine or Cloud Datastore provide autoscaling that enables application to grow with the users. The user has to pay only for what he or she uses.

### 11.3.2 Cloud Storage

Google Cloud Storage is a RESTful online file storage web service for storing and accessing one's data on Google's infrastructure. Representational state transfer (REST) is an architectural style consisting of a coordinated set of architectural constraints applied to components, connectors, and data elements within a distributed system. The service combines the performance and scalability of Google's cloud with advanced security and sharing capabilities. Google Cloud Storage is safe and secure. Data are protected through redundant storage at multiple physical locations.

The following are the few tools for Google Cloud Storage:

- *Google Developers Console* is a web application where one can perform simple storage management tasks on the Google Cloud Storage system.
- *gsutil* is a Python application that lets the user access Google Cloud Storage from the command line.

### 11.3.3 Google Cloud Connect

Google Cloud Connect is a feature provided by Google Cloud by integrating cloud and the application programming interface (API) for Microsoft Office. After installing a plug-in for the Microsoft Office suite of programs, one can save files to the cloud. The cloud copy of the file becomes the master document that everyone uses. Google Cloud Connect assigns each file a unique URL that can be shared to let others view the document.

If changes are made to the document, those changes will show up for everyone else viewing it. When multiple people make changes to the same section of a document, Cloud Connect gives chance to the user to choose which set of changes to keep.

When the user uploads a document to Google Cloud Connect, the service inserts some metadata into the file. Metadata is information about other information. In this case, the metadata identifies the file so that changes will track across all copies. The back end is similar to the Google File System and relies on the Google Docs infrastructure. As the documents sync to the master file, Google Cloud Connect sends the updated data out to all downloaded copies of the document using the metadata to guide updates to the right files.

### 11.3.4 Google Cloud Print

Google Cloud Print is a service that extends the printer's function to any device that can connect to the Internet. To use Google Cloud Print, the user needs to have a free Google profile, an app, a program, or a website that incorporates the Google Cloud Print feature, a cloud-ready printer or printer connected to a computer logged on to the Internet.

When Google Cloud Print is used through an app or website, the print request goes through the Google servers. Google routes the request to the appropriate printer associated with the user's Google account. Assuming the respective printer is on and has an active Internet connection, paper, and ink, the print job should execute on the machine. The printer can be shared with other people for receiving documents through Google Cloud Print.

Because most printers are not cloud ready, most Google Cloud Print users will need to have a computer act as a liaison. Google Cloud Print is an extension built into the Google Chrome Browser, but it should be enabled explicitly. Once enabled, the service activates a small piece of code called a connector. The connector's job is to interface between the printer and the outside world. The connector uses the user's computer printer software to send commands to the printer.

If one has a cloud-ready printer, one can connect the printer to the Internet directly without the need for a dedicated computer. The cloud printer has to be registered with Google Cloud Print to take advantage of its capabilities.

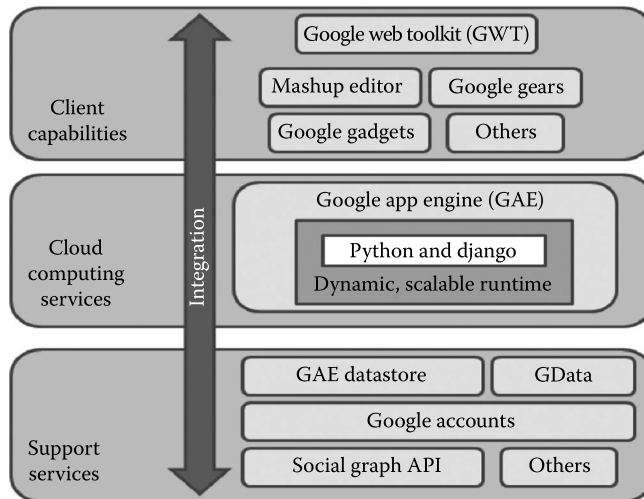
Because Google allows app and website developers to incorporate Google Cloud Print into their products as they see fit, there is no standard approach to executing a print job. Google Cloud Print depends on developers incorporating the feature into their products. Not every app or site will have Google Cloud Print built into it, which limits its functionality. Naturally, Google builds the service into its own products, but many people rely on services from multiple sources and may find Google Cloud Print does not have a wide enough adoption to meet all their needs.

### 11.3.5 Google App Engine

Google App Engine lets the user run web applications on Google's infrastructure. App Engine applications are easy to build, easy to maintain, and easy to scale as traffic and data storage needs grow. With App Engine, there are no servers to maintain: Just upload the application, and it is ready to serve users.

The app can be served from the user's own domain name (such as `http://www.example.com/`) using Google Apps. Otherwise, it can be served using a free name on the `appspot.com` domain. An application can be shared with the world or limit access to members of an organization. Figure 11.2 shows the different modules in Google App Engine. Integration of cloud computing services with support services and client capabilities is shown in the diagram.

Google App Engine supports apps written in several programming languages. With App Engine's Java runtime environment, one can build one's app using standard Java technologies, including the JVM, the Java servlets, and the Java programming language—or any other language. App Engine also features a Python runtime environment, which includes a fast Python interpreter and the Python standard library. App Engine also features a PHP runtime, with native support for Google Cloud SQL and Google Cloud Storage that works just like using a local MySQL instance and doing local file writes. Finally, App Engine

**FIGURE 11.2**

Google App Engine. (Adapted from <http://rdn-consulting.com/blog/tag/azure/>, accessed January 16, 2014).

provides a Go runtime environment that runs natively compiled Go code. These runtime environments are built to ensure that your application runs quickly, securely, and without interference from other apps on the system.

With App Engine also, the user has to only pay for what he or she uses. There are no setup costs and no recurring fees. The resources used by the application such as storage and bandwidth are measured in gigabyte and billed at competitive rates. One has to control the maximum amount of resources one's app can consume, so it always stays within one's budget.

App Engine costs nothing to get started. All applications can use up to 1 GB of storage and enough CPU and bandwidth to support an efficient app serving around five million page views a month, absolutely free. When billing is enabled for the application, free limits are raised, and one has to only pay for resources one uses above the free levels.

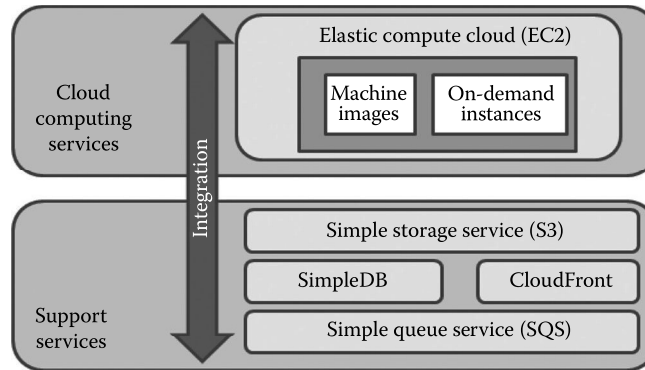
---

## 11.4 Amazon Web Services

Amazon Web Services (AWS) is a collection of remote computing services (also called web services) that together make up a cloud computing platform, offered over the Internet by Amazon.com. The most central and well known of these services are Amazon Elastic Compute Cloud (Amazon EC2), Amazon Simple Queue Service (Amazon SQS), and Amazon S3 as shown in Figure 11.3.

Amazon EC2 is a computing service, whereas Amazon SQS and Amazon S3 are support services. The service is advertised as providing a large



**FIGURE 11.3**

AWS. (Adapted from <http://rdn-consulting.com/blog/tag/azure/>, accessed January 16, 2014).

computing capacity (potentially many servers) much faster and cheaper than building a physical server farm. Amazon's data centers are located at Ashburn, Virginia, Dallas/Fort Worth, Los Angeles, Miami, Newark, New Jersey, Palo, Alto, California, Seattle, St. Louis, Amsterdam, Dublin, Frankfurt, London, Hong Kong, Singapore, Tokyo, etc.

#### 11.4.1 Amazon Elastic Compute Cloud

Amazon EC2 is an IaaS offered by AWS and is the leading provider of IaaS in the current market. Powered by a huge infrastructure that the company has built to run its retail business, Amazon EC2 provides a true virtual computing environment. By providing a variety of virtual machine or instance types, operating systems, and software packages to choose from, Amazon EC2 enables the user to instantiate virtual machines of his choice through a web service interface. The user can change the capacity and characteristics of the virtual machine by using the web service interfaces, hence named *elastic*. Computing capacity is provided in the form of virtual machines or server instances by booting Amazon Machine Images (AMI), which can be instantiated by the user. An AMI contains all the necessary information needed to create an instance. The primary Graphical User Interface (GUI) interface is the AWS Management Console (point and click) and a web service API that supports both Simple Object Access Protocol and Query Requests. The API provides programming libraries and resources for Java, PHP, Python, Ruby, Windows, and .Net. The infrastructure is virtualized by using Xen hypervisor, and different instance types are provided as follows:

- Standard instances—suitable for most applications
- Micro instances—suitable for low-throughput applications
- High-memory instances—suitable for high-throughput applications

- High-CPU instances—suitable for compute-intensive applications
- Cluster compute instances—suitable for high-performance computing (HPC) applications

The instances can be obtained on demand on an hourly basis, thus eliminating the need of forecasting computing needs earlier. Instances can be reserved earlier, and a discounted rate is charged for such instances. Users can also bid on unused Amazon EC2 computing capacity and obtain instances. Such instances are called as Spot Instances. Those bids that exceed the current Spot Price is provided with the instance, which allows the user to reduce costs. The Spot Price is varying and is decided by the company.

Instances can be placed in multiple locations, which are defined by regions and availability zones. Availability zones are distinct locations that are engineered to be insulated from failures in other availability zones and provide inexpensive, low-latency network connectivity to other availability zones in the same region. Thus, placing the instances in multiple locations enables fault tolerance and failover reliability. The Amazon EC2 instances can be monitored and controlled by the AWS Management Console and the web service API. However, AWS provides Amazon Cloud Watch, a web service that provides monitoring for AWS cloud resources, starting with Amazon EC2. It provides customers with visibility into resource utilization, operational performance, and overall demand patterns—including metrics such as CPU utilization, disk reads and writes, and network traffic.

Instances are authenticated using a signature-based protocol, which uses key pairs. Another important feature provided is the Amazon Virtual Private Cloud (Amazon VPC). The existing IT infrastructure can be connected to Amazon EC2 via a virtual private network (VPN). Isolated computing resources are provided in Amazon VPC, and the existing management capabilities such as security services, firewalls, and intrusion detection systems can be extended to isolated resources of Amazon EC2.

Elastic load balancing (ELB) enables the user to automatically distribute and balance the incoming application's traffic among the running instances based on metrics such as request count and request latency. Fault tolerance and automatic scaling can be performed by configuring the ELB as per the specific needs. ELB monitors the health of the instances running and routes traffic away from a failing instance.

An instance is stored as long as it is operational and is removed on termination. Persistent storage can be enabled by using either Elastic Block Storage (EBS) or Amazon Simple Storage Service (S3). EBS provides a highly reliable and secure storage, and the storage volumes can be used to boot an Amazon EC2 instance or be attached to an instance as a standard block device. Amazon S3 provides a highly durable storage infrastructure designed for mission-critical and primary data storage. Storage is based on units called objects whose size can vary from one byte to five gigabytes of data. These objects are stored in a bucket and retrieved via a unique, developer-assigned key.

It is accessible through a web service interface and provides authentication procedures to protect against unauthorized access.

#### **11.4.2 Amazon Simple Storage Service**

Amazon Simple Storage Service known as Amazon S3, is the storage for the Internet. It is designed to make web-scale computing easier for developers. Amazon S3 provides a simple web service interface that can be used to store and retrieve any amount of data, at any time, from anywhere on the web. It gives any developer access to the same highly scalable, reliable, secure, fast, inexpensive infrastructure that Amazon uses to run its own global network of websites. The service aims to maximize benefits of scale and to pass those benefits on to developers.

Along with its simplicity, it also takes care of other features like security, scalability, reliability, performance, and cost. Thus, Amazon S3 is a highly scalable, reliable, inexpensive, fast, and also easy to use service that meets design requirements and expectations.

Amazon S3 provides a highly durable and available store for a variety of content, ranging from web applications to media files. It allows users to offload storage where one can take advantage of scalability and pay-as-you-go pricing. For sharing content that is either easily reproduced or where one needs to store an original copy elsewhere, Amazon S3's Reduced Redundancy Storage (RRS) feature provides a compelling solution. It also provides a better solution in the case of storage for data analytics. Amazon S3 is an ideal solution for storing pharmaceutical data for analysis, financial data for computation, and images for resizing. Later this content can be sent to Amazon EC2 for computation, resizing, or other large-scale analytics without incurring any data transfer charges for moving the data between the services.

Amazon S3 offers a scalable, secure, and highly durable solution for backup and archiving critical data. For data of significant size, the AWS Import/Export feature can be used to move large amounts of data into and out of AWS with physical storage devices. This is ideal for moving large quantities of data for periodic backups, or quickly retrieving data for disaster recovery scenarios. Another feature offered by Amazon S3 is its Static Website Hosting, which is ideal for websites with static content, including html files, images, videos, and client-side scripts such as JavaScript.

#### **11.4.3 Amazon Simple Queue Service**

Another service of AWS is Amazon SQS. It is a fast, reliable, scalable, fully managed message queuing service. SQS makes it simple and cost effective to decouple the components of a cloud application. SQS can be used to transmit any volume of data, at any level of throughput, without losing messages or requiring other services to be always available.

Amazon SQS is a distributed queue system that enables web service applications to quickly and reliably queue messages that one component in the application generates to be consumed by another component. A queue is a temporary repository for messages that are waiting to be processed.

Amazon SQS offers various features like allowing multiple readers and writers at the same time, providing access control facilities, guaranteeing high availability of sending, and retrieving messages due to redundant infrastructure. It also gives provision for having variable length messages as well as configurable settings for each queue.

---

## **11.5 Microsoft**

Cloud computing provides a new way of looking at IT at Microsoft called Microsoft IT (MSIT). Cloud computing is now the preferred and default environment for new and migrated applications at Microsoft. MSIT has developed a methodology and a set of the best practices for analyzing their current application portfolio for possible candidates to migrate to cloud computing. This analysis enables MSIT to select the ideal cloud computing-based environment for each application. MSIT has captured these best practices and documented them for other Microsoft customers who wish to migrate their organizations to cloud computing.

### **11.5.1 Windows Azure**

Windows Azure Cloud Services (web and worker roles/PaaS) allow developers to easily deploy and manage application services. It delegates the management of underlying role instances and operating system to the Windows Azure platform.

The Migration Assessment Tool (MAT) for Windows Azure encapsulates all the information to be aware of before attempting the application migration to Windows Azure. Based on the response to a series of simple binary questions, the tool generates a report that outlines the amount of development effort involved to migrate the application, or the architecture considerations for a new application.

The Windows Azure Pricing Calculator analyzes an application's potential public cloud requirements against the cost of the application's existing infrastructure. This tool can help to compare current operational costs for an application, against what the operating costs would be on Windows Azure and SQL Azure.

Windows Azure Pack for Windows Server is a collection of Windows Azure technologies available to Microsoft customers at no additional cost for installation into their data center. It runs on top of Windows Server 2012

R2 and System Center 2012 R2 and, through the use of the Windows Azure technologies, it allows you to offer a rich, self-service, multitenant cloud, consistent with the public Windows Azure experience.

### 11.5.2 Microsoft Assessment and Planning Toolkit

The Microsoft Assessment and Planning Toolkit (MAP) is an agentless, automated, multiproduct planning and assessment tool for cloud migration. MAP provides detailed readiness assessment reports, executive proposals, and hardware and software information. It also provides recommendations to help organizations accelerate the application migration process for both private and public cloud planning assessments. MAP analyzes server utilization data for server virtualization and also server consolidation with Hyper-V.

### 11.5.3 SharePoint

Microsoft offers its own online collaboration tool called SharePoint. Microsoft SharePoint is a web application platform that comprises a multipurpose set of web technologies backed by a common technical infrastructure. By default, SharePoint has a Microsoft Office–like interface, and it is closely integrated with the Office suite. The web tools are designed to be usable by nontechnical users. SharePoint can be used to provide intranet portals, document and file management, collaboration, social networks, extranets, websites, enterprise search, and business intelligence. It also has system integration, process integration, and workflow automation capabilities. Unlike Google Cloud Connect, Microsoft SharePoint is not a free tool. But it has additional features that cannot be matched by Google or any other companies.

---

## 11.6 IBM

IBM is one among the players in the field of cloud computing offering various cloud services to the consumers. IBM cloud computing consists of cloud computing solutions for enterprises as offered by the global IT company IBM. All offerings are designed for business use, marketed under the name IBM SmartCloud. IBM cloud includes IaaS, SaaS, and PaaS offered through public, private, and hybrid cloud delivery models, in addition to the components that make up those clouds.

IBM offers an entry point to cloud computing whether a client is designing their own virtual private cloud, deploying cloud service, or consuming cloud workload applications. The IBM cloud framework begins with the physical

hardware of the cloud. IBM offers three hardware platforms for cloud computing, which offer built-in support for virtualization. The next layer of the IBM framework is virtualization. IBM offers IBM Websphere application infrastructure solutions that support programming models and open standards for virtualization.

The management layer of the IBM cloud framework includes IBM Tivoli middleware. Management tools provide capabilities to regulate images with automated provisioning and deprovisioning, monitor operations, and meter usage while tracking costs and allocating billing. The last layer of the framework provides integrated workload tools. Workloads for cloud computing are services or instances of code that can be executed to meet specific business needs. IBM offers tools for cloud-based collaboration, development and test, application development, analytics, business-to-business integration, and security.

### **11.6.1 Cloud Models**

IBM offers a spectrum of cloud delivery options ranging from solely private cloud to solely public cloud and numerous variations in between. IBM gives the option to build a customized cloud solution out of a combination of public cloud and private cloud elements. Companies that prefer to keep all data and processes behind their own firewall can choose a private cloud solution managed by their own IT staff. A company may also choose pay-as-you-go pricing that allows them to run lower-profile applications on a secure public cloud model. Hybrid cloud options allow for some processes to be hosted and managed by IBM, while others are kept on a private cloud or on a VPN or Virtual Local Area Network. IBM also offers planning and consultation throughout the deployment process. Cloud computing is the best choice for mobile software. IBM offers five different cloud provision models:

1. Private cloud, owned and operated by the customer
2. Private cloud, owned by the customer but operated by IBM (or another provider)
3. Private cloud, owned and operated by IBM (or another provider)
4. Virtual private cloud services, based on multitenant support for individual enterprises
5. Public cloud services, based on the provision of functions to individuals

The majority of cloud users choose a hybrid cloud model, with some workloads being served by internal systems, some from commercial cloud providers, and some from public cloud service providers.

For enterprise customers who perceive that the security risk of cloud computing adoption is too high, IBM specializes in secure private cloud offerings.

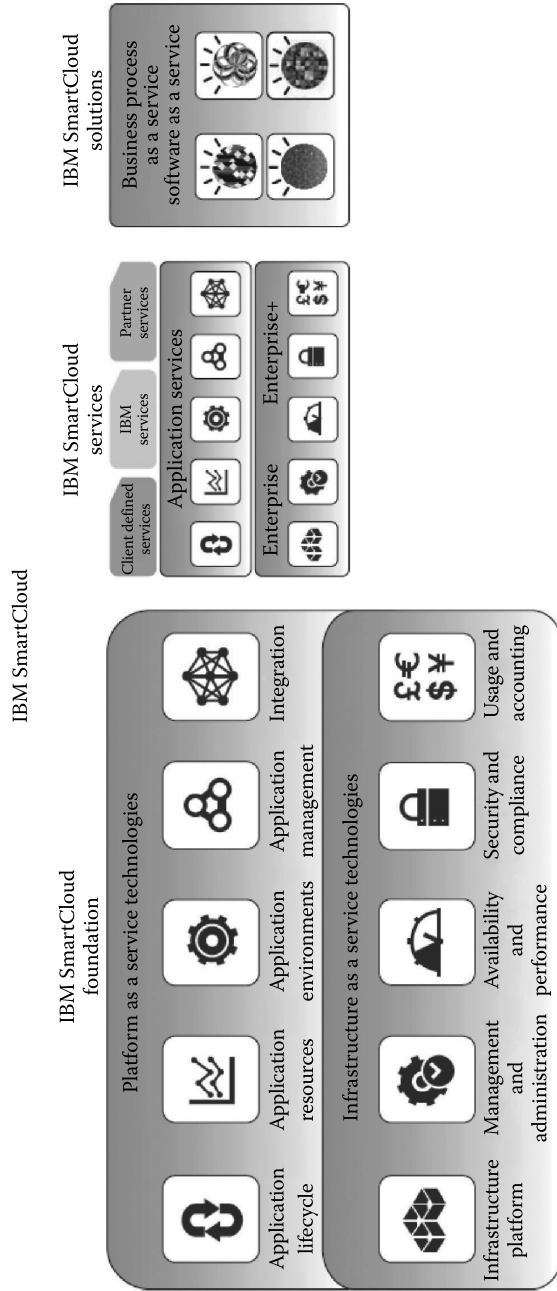
For building strictly private clouds, IBM offers IBM Workload Deployer and Cloudburst as ready-to-deploy, *cloud in a box*-style solutions. Cloudburst provides blade servers, middleware, and virtualization for an enterprise to build its own cloud-ready virtual machines. Workload Deployer connects an enterprise's existing servers to virtualization components and middleware in order to help deploy standardized virtual machines designed by IBM. For customers who prefer to perform their own integration of private clouds, IBM offers a choice of hardware and software building blocks, along with recommendations and reference architecture, leading the way to deployment. Clients may choose from IBM virtualization-enabled servers, middleware, and SaaS applications.

### 11.6.2 IBM SmartCloud

IBM SmartCloud is a branded ecosystem of cloud computing products and solutions from IBM. It includes IaaS, SaaS, and PaaS offered through public, private, and hybrid cloud delivery models. IBM places these offerings under three umbrellas: SmartCloud Foundation, SmartCloud Services, and SmartCloud Solutions. Figure 11.4 briefly explains the architecture of IBM SmartCloud.

SmartCloud Foundation consists of the infrastructure, hardware, provisioning, management, integration, and security that serve as the underpinnings of a private or hybrid cloud. Built using those foundational components, PaaS, IaaS, and backup services make up SmartCloud Services. Running on this cloud platform and infrastructure, SmartCloud Solutions consist of a number of collaboration, analytics, and marketing SaaS applications.

Along with IaaS, PaaS, and SaaS, IBM also offers Business Process as a Service (BPaaS). Infrastructure cloud services provide the consumer the provision of processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. In platform cloud services, a consumer can deploy consumer-created or consumer-acquired applications onto the cloud infrastructure created using programming languages and tools supported by the provider. Application cloud services allow consumers to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based e-mail). Business process cloud services are any business process (horizontal or vertical) delivered through the cloud service model (multitenant, self-service provisioning, elastic scaling, and usage metering or pricing) via the Internet with access via web-centric interfaces and exploiting web-oriented cloud architecture. The BPaaS provider is responsible for the related business functions.



**FIGURE 11.4** Architecture of IBM SmartCloud. (Adapted from Transitioning to IBM smart cloud notes, Smart Cloud White Paper-IBM.)



---

## 11.7 SAP Labs

SAP Labs makes enterprise software to manage business operations and customer relations. SAP is the leader in the market of enterprise applications in terms of software and software-related service. The company's best-known software products are its enterprise resource planning application systems and management (SAP ERP), its enterprise data warehouse product—SAP Business Warehouse (SAP BW), SAP Business Objects software, and most recently, Sybase mobile products and in-memory computing appliance SAP HANA. SAP is one of the largest software companies in the world.

### 11.7.1 SAP HANA Cloud Platform

SAP HANA Cloud Platform is an open-standard, Eclipse-based, modular PaaS. In SAP HANA Cloud Platform, applications are deployed via command-line tools to the cloud as web application archive (WAR) files or OSGi bundles. OSGi bundles are normal jar components with extra manifest headers. The applications run within the Java-based SAP HANA Cloud Platform runtime environment. It is powered by SAP HANA and can be maintained using web-based management tools.

The main features of SAP HANA Cloud Platform are as follows:

- Enterprise platform built for developers
- Native integration with SAP and non-SAP software
- In-memory persistence
- Secure data platform
- Lightweight, modular runtime container for applications

SAP HANA Cloud Platform lets the users quickly build and deploy business and consumer applications that deliver critical new functionality to meet emerging business needs. It also helps connect users with customers in more engaging experiences. It provides connectivity based on the cloud connectivity service. As a result, the platform streamlines the integration of new applications at the lowest possible total cost of ownership. Support for open programming standards provides a low barrier entry for developers. This makes them productive from the start in building enterprise applications that can integrate with any SAP or non-SAP solution. No new coding skills are required to work with SAP HANA.

### 11.7.2 Virtualization Services Provided by SAP

ERP virtualization increases a project's return on investment by maximizing hardware utilization. The business benefits of virtualization of ERP

applications are shorter development cycles, reduction in IT costs, improved availability, and energy saving. A joint service from SAP and VMware helps in transition to a more open and flexible private cloud platform based on proven virtualization technology.

---

## 11.8 Salesforce

Salesforce.com is a cloud computing and social enterprise SaaS provider based in San Francisco. Of its cloud platforms and applications, the company is best known for its Salesforce CRM product, which is composed of Sales Cloud, Service Cloud, Marketing Cloud, Force.com, Chatter, and Work.com. In addition to its products and platforms, Salesforce.com created AppExchange, a custom application building and sharing platform. The company also has consulting, deployment, and training services.

### 11.8.1 Sales Cloud

Sales Cloud refers to the *sales* module in Salesforce.com. It includes Leads, Accounts, Contacts, Contracts, Opportunities, Products, Pricebooks, Quotes, and Campaigns (limits apply). It includes features such as web-to-lead to support online lead capture, with autoresponse rules. It is designed to be a start-to-end setup for the entire sales process. Sales Cloud manages contact information and integrates social media and real-time customer collaboration through Chatter. The Sales Cloud gives a platform to connect with customers from complete, up-to-date account information to social insights, all in one place and available anytime, anywhere. Everything is automatically pushed in real time, from contact information to deal updates and discount approvals.

Salesforce.com created the Sales Cloud to be as easy to use as a consumer website like Amazon and built it in the cloud to eliminate the risk and expense associated with traditional software. With its open architecture and automatic updates, the Sales Cloud does away with the hidden costs and drawn-out implementations of traditional CRM software. By continuing to innovate and embrace technologies like mobile, collaboration, and social intelligence, the Sales Cloud has continued to pull ahead of the competition.

### 11.8.2 Service Cloud: Knowledge as a Service

*Service Cloud* refers to the *service* (as in *customer service*) module in Salesforce.com. It includes Accounts, Contacts, Cases, and Solutions. It also encompasses features such as the public knowledge base, web-to-case, call center, and self-service portal, as well as customer service automation. Service Cloud

includes a call center–like case tracking feature and a social networking plug-in for conversation and analytics.

The Service Cloud delivers the world’s first enterprise-grade knowledge base to run entirely on an advanced, multitenant cloud platform. That means one can get all the cloud computing benefits that Salesforce.com is known for delivering without expensive data centers or software. Just powerful knowledge management, without the hassle of on-premises software, is provided. Unlike stand-alone applications, this knowledge base is fully integrated with everything else. Service Cloud has to offer all the tools one needs to run the entire service operation. When the consumer’s knowledge base is a core part of CRM solution, knowledge as a process can be managed. One can continually create, review, deliver, analyze, and improve the knowledge. And, because it is delivered by the Service Cloud, user’s knowledge is available wherever other customers need it. Agents have the right answers at their fingertips to communicate over the phone, send out through an e-mail, or share via a chat client. The same knowledge base serves up answers to the service website is a part of company’s public site. If one wants to take advantage of social channels like Twitter or Facebook, one can easily share knowledge that is tapped into the wisdom of the crowd to capture new ideas or answers. All this is done securely.

The Service Cloud gives the tools that are needed to manage knowledge at enterprise scale. But it also delivers the same great ease of use that Salesforce.com is known for. That means user will benefit no matter what size or how complex the business is.

---

## 11.9 Rackspace

Rackspace Cloud, a part of Rackspace, is another player in the cloud computing market. Offering IaaS to clients, it has been used by a large number of enterprises. Rackspace Cloud offers three cloud computing solutions—Cloud Servers, Cloud Files, and Cloud Sites. Cloud Servers provide computational power on demand in minutes; Cloud Sites are for robust and scalable web hosting, and Cloud Files are for elastic online file storage and content delivery.

Cloud Servers is an implementation of IaaS where the computing capacity is provided as virtual machines that run in the Cloud Servers systems. The virtual machine instances are configured with different amounts of capacities. The instances come in different flavors and images. A flavor is an available hardware configuration for a server. Each flavor has a unique combination of disk space, memory capacity, and priority for CPU time. A varied set of instances are available for the user to choose from.

These virtual machines are instantiated using images. An image is a collection of files used to create or rebuild a server. A variety of prebuilt operating

system images are provided by Rackspace Cloud (64-bit Linux distributions—Ubuntu, Debian, Gentoo, CentOS, Fedora, Arch, and Red Hat Enterprise Linux) or Windows Images (Windows Server 2008 and Windows Server 2003). These images can be customized to the user's choice to create custom images.

The Cloud Servers systems are virtualized using the Xen Hypervisor for Linux and Xen Server for Windows. The virtual machines that are generated come in different sizes and measured based on the amount of physical memory reserved. Currently, the physical memory can vary from 256 MB to 15.5 GB. In the event of availability of extra CPU power, Rackspace Cloud claims to provide extra processing power to the running workloads, free of cost.

Backup schedules can be created to define when to create server images. This is a useful feature, which enables the user to continue work in the event of failures by using the backup images. Custom images are helpful in creating backup schedules. A type of images, referred to as *gold* server images, can be produced if the servers of that configuration are to be instantiated frequently.

Cloud Servers can be run through the Rackspace Cloud Control Panel (GUI) or programmatically via the Cloud Server API using a RESTful interface. The control panel provides billing and reporting functions and provides access to support materials including developer resources, a knowledge base, forums, and live chat. The Cloud Servers API was open sourced under the Creative Commons Attribution 3.0 license. Language bindings via high-level languages like C++, Java, Python, or Ruby that adhere to the Rackspace specification will be considered as Rackspace-approved bindings. The virtual machine instances are authenticated in the API by a token-based protocol that uses the HTTP x-Header. Private/public keys are used to ensure Secured Shell Access.

Cloud Servers scale automatically to balance load. This process is automated and initiated from either the Rackspace Cloud Control Panel or the Cloud Server API. The amount to scale is specified; the Cloud Server is momentarily taken offline; the RAM, disk space, and CPU allotment are adjusted; and the server is restarted. A Cloud Server can be made to act as a load balancer using simple readily available packages from any of the distribution repositories. Rackspace Cloud is working on beta version of the Cloud Load Balancing product, which provides a complete load balancing solution.

Cloud Servers are provided persistent storage through RAID10 disk storage; thus, data persistency is enabled leading to better functioning.

---

### 11.10 VMware

VMware, a leader in virtualization technology, has come up with enterprise cloud computing solutions. Having been a dominating player in the virtualization domain, VMware is currently providing a range of products for the development of private and public clouds and for leveraging the services

offered by both as a hybrid cloud, such as VMware vCloud Director, VMware vCloud Datacenter Services, VMware vSphere, and VMware vShield to name a few.

Private clouds enable the better usage and management of internal IT infrastructure than the traditional methods. Greater operational efficiency, secure, fault-tolerant, well-managed computing environments can be modeled and operated. VMware's private cloud offering provides greater standardization, rapid provisioning, and self-service for all applications and unparalleled cost savings by consolidating their physical infrastructures. VMware's modular technology enables the user to select from a variety of hardware, software, and certified service providers to result in efficient cloud computing. Thus, the family of products offered by VMware promotes compatibility and retains the choice of freedom for the users to obtain desired services.

Private clouds can be created by using the VMware vSphere and VMware vCloud Director. VMware vSphere is a robust virtualization platform used to transform IT infrastructures into virtual storage, compute, and network resources and provide them as a service within the organization. VMware vSphere provides services at both the infrastructure and application levels. At the infrastructure level, it provides options to perform efficient operation and management of the compute, storage, and network resources. At the application level, service-level controls are provided for the applications running on the underlying infrastructures, leading to available, secure, and scalable applications.

The VMware vCloud Director, coupled with VMware vSphere, is a software solution that enables enterprises to build secure, multitenant private clouds by pooling infrastructure resources into virtual datacenters and exposing them to users through web-based portals and programmatic interfaces as fully automated, catalog-based services. VMware vCloud Director abstracts the virtual computing environment from the underlying resources and provides a multitenant architecture that features isolated virtual resources, independent LDAP authentication, specific policy controls, and unique catalogs. VMware vShield technologies are used to provide security to these environments by using services like perimeter protection, port-level firewall, NAT and DHCP services, site-to-site VPN, network isolation, and web load balancing. The VMware vCloud Director allows users to catalog infrastructure and application services of the desired configurations and deploy and consume them as needed. Interactions with the virtual data centers or the catalogs are through a user-friendly web portal or the vCloud API. The vCloud API is an open, REST-based API that provides scripted access, complying with the open virtualization format (OVF). The API can be used along with VMware vCenter Orchestrator to automate and orchestrate operational processes like routine tasks, activities, and workflows.

Public and hybrid cloud solutions are provided by VMware by partnering with other companies, certified as service providers. VMware vCloud Datacenter Services and VMware vCloud Express offer efficient solutions for

utilizing IaaS either as a public cloud or a hybrid cloud. vCloud Datacenter Services provides a scalable environment, where internal resources are augmented with the external resources. vCloud Datacenter Services are built on the same technology and foundations as VMware vCloud Director and VMware vSphere to enable interoperability between cloud environments. Thus, the user is free to burst his private cloud into public cloud of his preferred service provider.

vCloud Express is an IaaS offering delivered by leading VMware service provider partners. It is a cobranded service that provides reliable, on-demand, pay-as-you-go infrastructure. The VMware vCloud Express providers are Virtacore vCloud Express, Hosting.com, Melbourne IT, and Terremark's vCloud Express. Instance types, load balancing, storage options, and pricing vary between service providers.

---

## 11.11 Manjrasoft

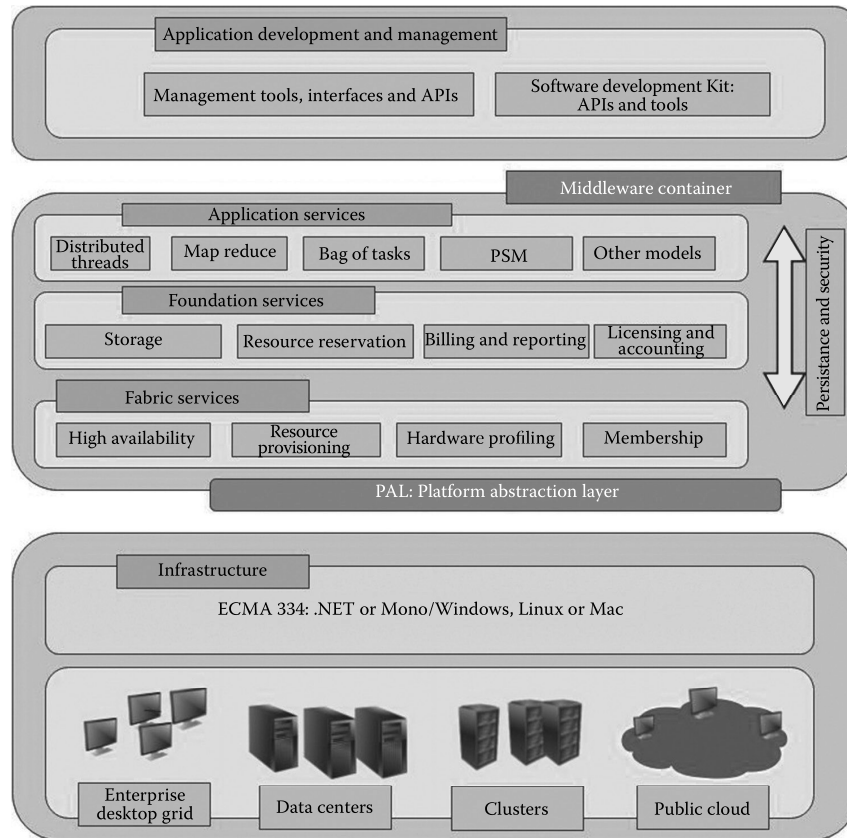
Manjrasoft is one of the nonmajor providers of cloud services. But it has come up with a platform called *Aneka* that provides a set of services that help the development of applications in an easier way. Manjrasoft develops market-oriented cloud computing platforms that allow one to build, accelerate, and manage the applications ultimately saving one's time and money, leading to enhanced business productivity and profit.

### 11.11.1 Aneka Platform

*Aneka* provides a set of services that make enterprise cloud construction and development of applications as easy as possible without sacrificing flexibility, scalability, reliability, and extensibility.

Figure 11.5 gives an overview of the Aneka platform. The key features supported by *Aneka* are as follows:

1. A configurable and flexible execution platform (container) enabling plug-gable services and security implementations. Multiple authentication/authorization mechanisms such as role-based security and Windows domain-based authentication are considered for this purpose.
2. Multiple persistence options including Relational Database Management System (RDBMS), Structured Query Language (SQL) Express, MySQL, and flat files.
3. Software development kit (SDK) supporting multiple programming models including object-oriented thread model, task model for legacy applications, and MapReduce model for data-intensive applications.
4. Custom tools such as Design Explorer for parameter sweep studies.



**FIGURE 11.5** Overview of the Aneka platform. (Adapted from [http://www.manjrasoft.com/aneka\\_architecture.html](http://www.manjrasoft.com/aneka_architecture.html), accessed January 8, 2013.)

5. Easy to use management tool for SLA and Quality of Service (QoS) negotiation and dynamic resource allocation.
6. Supports deployment of applications on private or public clouds in addition to their seamless integration.

*Aneka* allows servers and desktop PCs to be linked together to form a very powerful computing infrastructure. This allows companies to become energy efficient and save money without investing in a number of computers to run their complex applications.

Each *Aneka* node consists of a configurable container that includes information and indexing, scheduling, execution, and storage services. *Aneka* supports multiple programming models, security, persistence, and communications protocols.

## 11.12 Summary

In this chapter, we have discussed about various companies that support cloud computing by providing tools and technologies to adapt to the cloud environment. Each section briefly describes the cloud features supported in these companies. Few of the services like Google Docs and Google Cloud Print are free, whereas that of AWS, Microsoft, etc., are proprietary. Based on the specific requirements, the user has to make a trade-off between open source and closed source tools/services. An attempt has been made to list the tools/services offered by each company in Table 11.1.

Though there are a number of companies, we have chosen few companies that have progressed a lot in this area. Table 11.2 gives the information about few providers and the prices on a per hour basis along with the service

**TABLE 11.1**

Tools and Services Offered by Companies

Company Name	Tools/Services
EMC	Captiva Cloud toolkit
Google	Google App Engine, Google Docs, Google Cloud Connect Google Cloud Print
Amazon	Amazon EC2, Amazon S3, Amazon SQS
Microsoft	Microsoft Assessment and Planning Toolkit, Windows Azure Sharepoint
IBM	IBM Smart Cloud
Salesforce	Sales Cloud, Service Cloud
SAP LABS	SAP HANA Cloud
VMware	vCloud
Manjrasoft	Aneka Platform
Red Hat	OpenShift Enterprise, OpenShift Origin
Gigaspace	Cloudify

**TABLE 11.2**

Details of Cloud Service Providers

Provider Name	Service Model	Deployment Model	Server Operating System
Amazon Web Services	IaaS	Public	Windows, Linux
Google App Engine	PaaS	Public	Windows
Windows Azure	IaaS	Public	Windows, Linux
IBM Cloud	IaaS	Private, hybrid	Windows, Linux
Salesforce Platform	PaaS	Public	Windows, Linux
Rackspace	IaaS	Public, private, hybrid	Windows, Linux
SAP HANA Cloud	PaaS	Public	Linux



model and deployment model offered by that system. Further details regarding the companies can be seen in the references section.

---

### Review Points

- *MDW*: It is an entry point in the Captiva Cloud toolkit. This module helps to import documents from specified folders/repository (see Section 11.2.2).
- *REST*: It is an architectural style in a distributed system. Google Cloud Storage is a RESTful online storage web service (see Section 11.3.2).
- *AMI*: It is a special type of virtual appliance that is used to create a virtual machine in Amazon EC2 (see Section 11.4.1).
- *RRS*: It is a new storage solution by Amazon S3 that enables cost-effective storage of noncritical data by reducing the redundancy. This feature provides a better solution in the case of storage for data analytics in Amazon SQS (see Section 11.4.2).
- *MAT*: It is a tool provided by Microsoft Azure that addresses migration considerations including app server, database, integration, security, and instrumentation for different platforms (see Section 11.5.1).

---

### Review Questions

1. What do you mean by cloud service provider? Which are the major cloud service providers?
2. List the tools/services provided by Microsoft and explain them in brief.
3. What is Google Cloud Print? What are its advantages?
4. Explain SAP HANA Cloud in brief.
5. What are the services offered by EMC IT? Explain.
6. Explain the services provided by IBM SmartCloud.
7. What are the support services offered by Amazon Web Services? Explain.
8. What do you mean by *Knowledge as a Service*? Which company provides this service? Explain.
9. Explain the features of Aneka.
10. What is vCloud? Explain in brief.

---

## References

1. EMC IT's journey to the private cloud, applications and the cloud experience. White Paper-EMC.
2. <http://rdn-consulting.com/blog/tag/azure/>. Accessed January 16, 2014.
3. Transitioning to IBM SmartCloud notes. SmartCloud White Paper-IBM.
4. [http://www.manjrasoft.com/aneka\\_architecture.html](http://www.manjrasoft.com/aneka_architecture.html). Accessed January 8, 2014.

---

## Further Reading

Akamai solutions for cloud computing. White Paper-Akamai.

Cloud-based application evaluation, planning, and migration. Article, June 2012.

Dell vCloud featuring trend micro SecureCloud. White Paper-Dell.

EMC IT's journey to the cloud: A practitioner's guide. White Paper-EMC.

Getting cloud computing right. IBM Global Technology Services, White Paper. <http://googcloudlabs.appspot.com/>.

Service Cloud. <http://www.salesforce.com/in/service-cloud/overview/>.

Cloud Storage. <https://cloud.google.com/products/cloud-storage/>.

The Service Cloud—Knowledge as a service. White paper-Salesforce.com.

Varia, J. and S. Mathew. Overview of Amazon web services, 2012.

Vecchiola, C., X. Chu, and R. Buyya. Aneka: A software platform for .NET-based cloud computing. In Gentzsch, L., L. Grandinetti, and G. Joubert (eds.), *High Speed and Large Scale Scientific Computing*. IOS Press, Amsterdam, the Netherlands, 2009, pp. 267–295.

VMware vCloud hybrid services. White Paper-VMware.