Cloud Computing and the Value of zEnterprise

Redguides
for Business Leaders

Elsie Ramos
Kurt Acker
Robert Green
Sebastien Llaurency

- Deliver IT without boundaries and drive innovation
- Optimize your cloud infrastructure environment
- Plan for the integration and management capabilities of zEnterprise

Executive overview

Like many companies, your business probably has an IT environment that encompasses multiple, diverse information technology assets that require significant staff and budget to configure and manage. This complexity poses challenges to IT organizations as they try to keep pace with the increasing demand for more services.

You need to provide your users with rapid access to dynamically scalable IT resources. Unlike a client server environment, cloud computing has the elements of rapid scalability, on-demand self-service, ubiquitous network access, rapid elasticity, location independent resource pooling, and pay per use. These characteristics are key to achieve an optimal and effective cloud computing experience for the end user and to maximize the value to the business. Many of these elements are established strengths of the mainframe.

Cloud computing is a way to organize and streamline your IT environment to reduce the complexity and to realize IT without boundaries. Cloud computing facilitates better use of existing IT investments allowing you to do more with fewer resources. You can achieve a highly simplified and efficient IT environment, improve service delivery, and lower your costs through physical consolidation, virtualization, and better IT management. Cloud computing removes the barriers impeding the rapid delivery of new services and provides the opportunity to reinvent business processes and drive innovation.

In a typical cloud solution the application hosting options are usually limited to a single system architecture, resulting in system and workload inefficiencies that drive up costs and reduce operational flexibility. IBM® zEnterprise System enables you to build solutions optimized for the right platform. It provides the unique capability to create a heterogeneous cloud solution in one single managed environment, with the ability to integrate and manage workload optimized software solutions spanning multiple server architectures. Cloud computing on the zEnterprise System also provides service management capabilities that result in higher levels of staff productivity and enhances the qualities of services that are delivered to the consumer.

With a cloud computing solution on zEnterprise you can optimize your Capital Expense (CapEx) by being able to do more work with fewer hardware footprints then other systems and lower the overall Operating Expenses (OpEx) for your organization in a manner that can be measured and reported on. IBM System z® is a proven efficient platform based on virtualization that allows you to reduce the cost of running your cloud computing infrastructure.

Challenges of cloud computing

Executive focus on cloud computing is increasing because it promises to help organizations respond quickly to change, manage growth, unleash new and innovative ideas, gain competitive advantages, reduce risk, improve services, and empower end users, all while most importantly, reducing cost. More than 3,000 global CIOs responded in the IBM CIO 2011 study[1] and it showed that 60% of organizations view cloud computing as a way to grow their business and increase their competitiveness, as depicted in Figure 1.

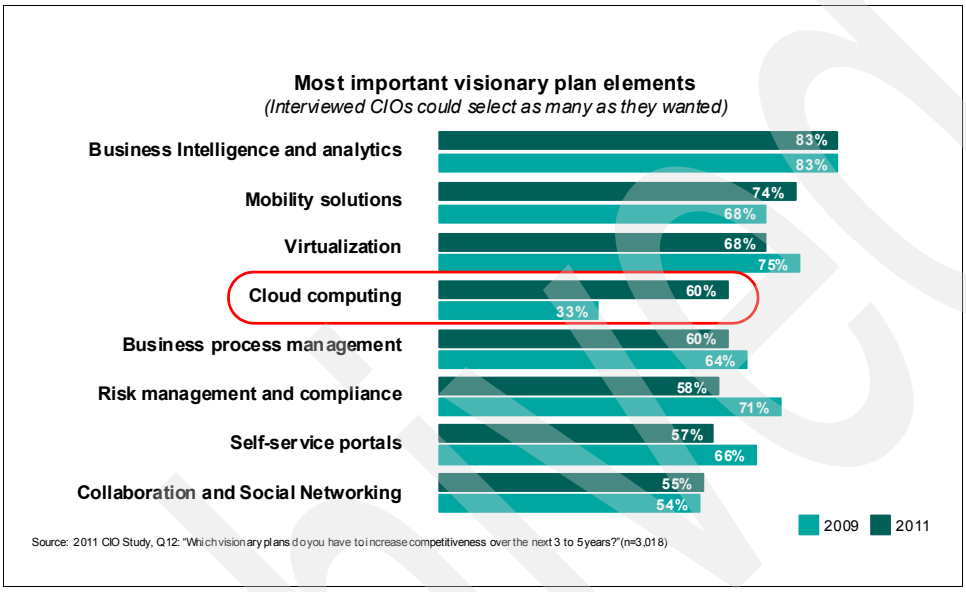

Figure 1 CIO visionary plan elements as measured in IBM 2011 CIO study

Cloud computing can transform how businesses use IT resources to:

- Drive down costs: Spend less on system maintenance and more on pursuing business opportunities
- Gain better control over costs: Provide more flexible options for workload hosting
- Unleash innovative ideas: Enable end-users to deliver new solutions with greater ease and less cost
- Gain competitive advantage: Provide better service and responsiveness

Implementing a cloud computing solution that delivers on these promises presents challenges that most organizations face today.

Prevalent issues with cloud environments

Many organization's cloud implementations are not meeting their expectations and they encounter many issues, such as security, availability, scalability, workload identification, support, standardization, and Quality of Service (QOS).

Security

Security is a top concern for IT management because it is a requirement to protect company information and to ensure system integrity. In the CDW 2011 Cloud Computing Tracking Poll[2],

[1] http://www-935.ibm.com/services/c-suite/cio/study.html
[2] http://newsroom.cdw.com/features/feature-05-26-11.html

41% of the respondents indicated that security is a major deterrent to implementing a cloud solution. There are specific security issues with public clouds, concerns that company data will reside on virtual servers that are on shared physical hardware, sometimes even shared with competing organizations. Another issue is whether private data can be accessed through the internet. However, in a private cloud model where all resources, including data, are behind the enterprise firewall, the respondent indicated that security is less of a concern.

Availability and reliability

Availability is another key concern. The infrastructure supporting the cloud must be reliable. Cloud service users expect the pools of resources to always be available for use. They also expect the availability of applications without any degradation of services.

Different hardware platforms have varying levels of acceptable reliability. Organizations are reluctant to move mission-critical applications from the mainframe to a cloud solution based on distributed platforms that do not provide the required reliability. As organizations move to cloud environments, there is also the concern with availability, which is generally addressed with redundant hardware and clustered software. But there is a cost to the organization to provision redundant servers to provide higher levels of availability.

Scalability

The cloud environment infrastructure must scale up easily to meet the processing requirements of the end users and scale out over time to meet the growth in cloud services. Organizations also want to understand the resource usage in their cloud environment since virtualization can drive resource utilization higher. This can be a challenge because each platform behaves differently when utilization is driven higher.

Elasticity

Elasticity is closely related to scalability. Cloud solutions must expand and contract with system usage. Organizations want the ability to dynamically provision additional capacity to their cloud service without having to tear down and rebuild. Likewise, when that extra capacity is no longer required, they want the ability to deprovision the excess capacity. These changes also must be metered so that the consumer can be charged according to what they provisioned at any given time.

Monitoring

Monitoring of a cloud environment presents a unique set of challenges. Organizations want to know how the environment is performing with respect to resource utilization. To be able to determine “hot spots”, if servers are consuming too many resources, and whether or not Service Level Agreements (SLAs) are being achieved is critical information. This problem is exasperated when the cloud infrastructure spans several platforms, operating systems, and physical locations. What is needed is the ability to control, manage, and view the environment based on the complete flow of any given workload from a single view.

Qualities of Service (QOS)

Traditionally, the expectation of QoS for the user is set by the architecture on which the workload is running, and it is defined as part of the services purchased. However, in a cloud environment, the underlying architecture is abstracted so that the consumer does not have an awareness of the platform or services that are used.

Organizations have come to expect a certain level of service and reliability from their IT infrastructure. Moving applications to a cloud environment does not change that expectation. Service Level Objectives (SLOs) and Service Level Agreements (SLAs) must still be met, regardless of where the compute resources are located or how they are being used.

Workload identification

Do not view cloud computing as a *one size fits all*. The choice on which platform to deploy applications is determined by the application requirements. An organization often needs several platforms with various operating systems and the cloud provider must be able to manage the different platforms.

Workload positioning is dependent on other applications and data. It is generally desirable to have the application “close to” the data being processed. If there are long network delays to retrieve data, the application performance suffers, although there might be plenty of excess processing capacity available to the application.

Seamless integration

Another factor to consider for a cloud environment is the integration of the workload and the combination of all the applications. Some platforms process different types of workloads more efficiently. The challenge in a cloud computing environment is to simplify and reduce the complexity while seamlessly integrating with the existing applications.

Application requirements, non-functional requirements and local factors all must be analyzed when deciding which workloads are ready to migrate to the cloud.

Charge back

Organizations must determine the cost per virtual server in a cloud environment to properly price the services for their consumers. A range of factors comprise the server cost, such as hardware, storage, network, hypervisor and operating system license and support, and middleware and cloud management software costs. Many products are licensed by the number of physical processors, or *cores* on which the software runs. The number of virtual servers deployed on a physical server, and the server utilization rate determines the software cost and the rate to charge the consumer for these services.

The number of virtual servers that can be deployed on a physical server and the server utilization rate determines the cost to the enterprise for the software and the rate to charge the consumer for these services.

Capacity planning

Organizations must forecast the cloud infrastructure capacity that is required to support their consumers and then plan for this capacity. For organizations that are not already taking advantage of consolidation and virtualization, this can be a difficult exercise. Likewise, an organization that is moving from a mix of platforms to a vendor’s single offering platform might not be able to gauge how much capacity is needed as applications change.

Support

Different vendor cloud solutions might not be able to manage and provide support outside of their own offering. This can result in vendor lock-in to a single platform/architecture solution, and resources already in use by the organization cannot be brought into the cloud environment. Since deploying a cloud solution requires a rip-and-replace strategy and results in assets that can no longer be used. The optimal cloud solution provides support across multiple disparate platforms and the ability to incorporate resources that are already in use.

Another area of support that must be addressed for the cloud is problem management, where groups of related incidents must be analyzed to determine the root cause. Reoccurring incidents can indicate an underlying problem that might require configuration changes or changes in capacity management functions. Also, all incidents must be reported through some type of event monitoring and management process. The incident and any associated action must be recorded for trending, and for a knowledge base search, should the same or

similar incident occur again. As cloud environments grow, automatic actions become a necessity so that the environment can become autonomous.

Release or change management also needs to be considered**.** Maintaining cloud virtual servers up-to-date on patches and changes can be a daunting task. A process must be planned to apply patches to the servers with minimal disruption. The install base for new virtual servers creation must be kept current on operating system and software fixes.

Asset lifecycle management must be taken into account. As assets reach the end of their useful life, there must be a plan on how to replace them in the environment. Organizations typically have a management process for doing technology refreshes of their IT resources. The implementation of cloud computing might require changes in that process.

Standardization

A critical factor in the success of a cloud implementation is the level of standardization that can be achieved for cloud resources. Typically an organization has multiple levels of operating systems and application software that they acquired over time. In a cloud environment, the objective is to provide a self-service portal for consumers to request provisioning of resources. This feature can be difficult to administer if the organization has to provide hundreds or thousands of server images for their consumers, and the consumers can quickly become overwhelmed to decide which services to request. The organization must be able to collapse the set of offerings down to a few distinct offerings that are clearly differentiated, for example by operating system and version. Middleware titles can also be offered as add-ons to the operating system but it must be provided as a manageable list. The management software that is used to administer the choices must be flexible and easy to use by the IT administration staff. It must be able to define titles and the associated options for a variety of platforms that are used within the enterprise.

System z Enterprise is uniquely positioned to address these challenges. Its robust heterogeneous virtualization environment makes it perfectly suited to host applications and workload across a variety of architectures. Its rich history of reliability and availability enable it to host cloud environments at levels of service that other platforms cannot achieve, creating in essence a “heterogeneous cloud”.

Solving cloud computing challenges

According to the U.S. National Institute of Standards and Technology (NIST) cloud computing is "*A model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.*"[1] A cloud can be public, private, both, or hybrid.

Cloud computing is a natural progression of the infrastructure transformation from standard managed services to cloud delivered services. Figure 2 illustrates the journey to an ideal cloud computing solution.

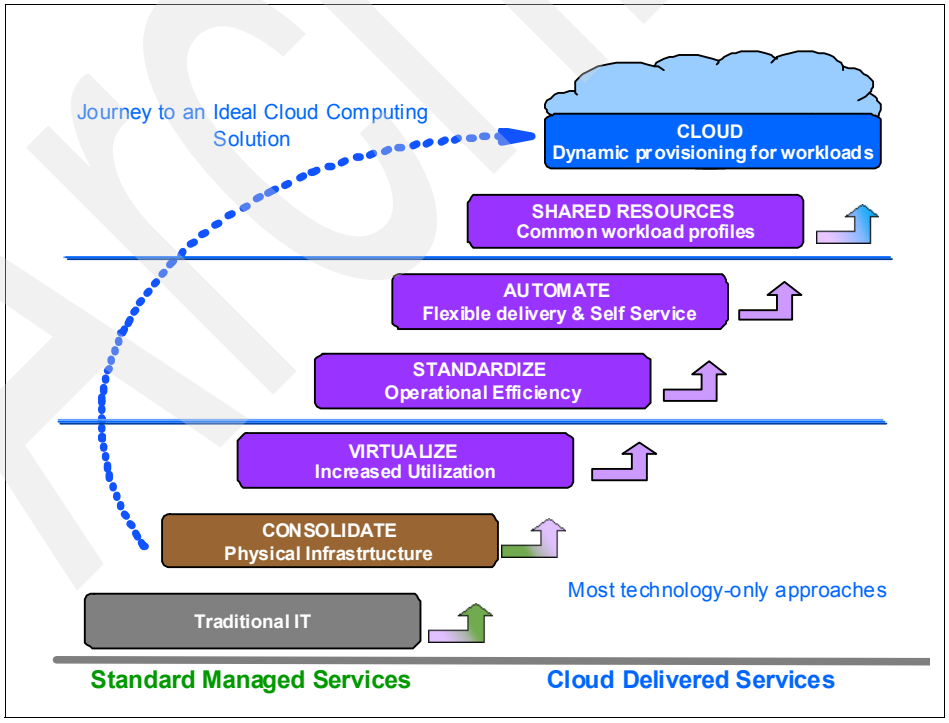

Figure 2 Considerations to get started with cloud computing

[1] http://csrc.nist.gov/groups/SNS/cloud-computing/index.html

This journey is an extension to the existing plans of most organizations to consolidate with virtualization to reduce expenses and increase asset utilization. The addition of standardization, automation, self service, and audit controls are a natural progression to these planning efforts. These characteristics are also key to achieving an optimal and effective cloud computing experience for the user with maximum value to the business.

Obtaining the ideal cloud computing solution begins with an assessment of the capabilities of your current environment versus the ones that compose an ideal cloud computing solution. This assessment requires a review of the key cloud characteristics, attributes, and models to identify your gaps to the cloud computing delivery model. It also helps to determine which elements to consider for your cloud computing solution and assess how to address the challenges to obtain the optimal return on investment that is expected from cloud computing solutions.

Use the following attributes and principles as a framework for your cloud computing solution. We discuss how using zEnterprise can mitigate the challenges previously described.

Cloud computing attributes

Cloud computing attributes are divided into two main families:

- Cloud services capabilities
- Cloud infrastructure capabilities

There is a dependency between these two families of attributes. The cloud-specific attributes inherit in a large part from the fundamentals ones, mainly in terms of non-functional requirements, such as scalability, security, resilience, availability, and so on.

These cloud attributes and characteristics must be incorporated as organizations progress on their journey to implement a cloud solution.

Cloud specific attributes

The NIST defined five essential characteristics or cloud specific attributes for a cloud computing solution:

- On demand self-service: The capability for the user to request a service without human interaction. The objective is to provide the choice and flexibility to the users for their service consumption without worrying about implementation, technical details, hardware licenses, or software licenses. The users can focus on requesting the appropriate service for their business activity.
- Resources pooling: The capability to structure physical and logical resources in pools. This provides you with the capability to dynamically affect your resources to a particular service. It is a source of cost reduction because it allows you to return the resources consumed by a service to their pool so it can be reused by another user. The objective of resource pooling is to optimize the utilization of your resources.
- Measured service: A key differentiator from the traditional IT environment. A cloud computing environment needs the capability to measure the utilization of the resources provided to a user. The objective of measuring resources is to charge back appropriately to the users of your services.
- Broad network access: This attribute has a direct impact on the utilization of the cloud service by the users. Your network must be responsive, secure, and highly available to

provide secure cloud services. The objective of broad network access is to deliver your services through multiple channels to a wide range of clients.

- Rapid elasticity: The capability of your cloud computing solution to scale rapidly on demand, to grow or reduce through provisioning mechanisms. The objective of rapid elasticity is to adapt your cloud computing environment to business demands and better align your IT resources to your business needs.

Cloud fundamental attributes

The cloud fundamental attributes are common in the traditional IT environment, and you might have already started to benefit from these. These attributes are fundamental to an efficient cloud computing solution and are the key elements for a cloud service:

- Standardization: For your cloud environment to be cost efficient you must select the services that you plan to offer and establish standards and rules for the creation and delivery of your services. This standardization must span your different IT organizations and teams. Having a standard process to manage a heterogeneous infrastructure is a key advantage because it provides the capability to react quicker to business changes.
- Consolidation: Initiate cost reduction by the consolidation of existing services to serve more than one project and avoid the purchase of new servers.This attribute can also generate environmental savings or remove constraints on your data center floor space.
- Virtualization: Is associated with consolidation and provides additional benefit in terms of IT flexibility. A virtualized environment and the ability to create virtual machines for applications' hosting is one major source of capital expenses savings.
- Automation: Reduces the service delivery time on two levels:
  - Technical automation: The automation of technical tasks and the grouping of tasks into a technical workflow that is automatically executed.
  - Process automation: The automation of the entire service delivery processes.

Cloud service models

Cloud computing can be delivered to consumers in several ways and there are various cloud models. Each model provides a level of service that is based on what the consumer needs from the environment, hardware, or infrastructure base giving the consumer the most flexibility with their environment.

Figure 3 on page 10 illustrates the four cloud services models. Across all cloud service models the management scope is determined by the provider. It is essentially the tasks that the operations staff of the provider takes on and not about the virtualization technology that is being used.

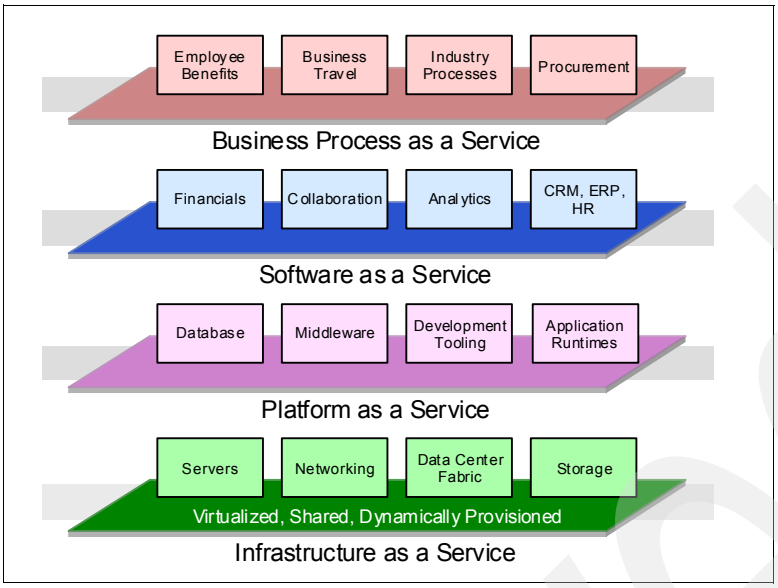

Figure 3 Cloud service models

The cloud service models are:

Infrastructure as a Service (IaaS): Cloud resources are provided in terms of underlying hardware and networking only. The consumer is responsible for installation and maintenance of operating systems, middleware, and applications, but does not manage or control the underlying cloud infrastructure. This offering gives the consumer the most flexibility and associated responsibility.

Platform as a Service (PaaS): The cloud infrastructure, operating system, and application hosting software are all provided. The consumer does not have any control or management capabilities of the underlying OS, but do control the applications. By abstracting the hardware, OS, and application hosting software, the consumer can focus their skills and efforts strictly on their applications.

Software as a Service (SaaS): Provides an end-to-end environment that includes the underlying infrastructure (processors, storage, networking), operating system, application hosting software, and applications. This service model allows applications to be made available to the user through a variety of devices through a thin client interface, such as a web browser.

Business Process as a Service (BPaaS): An additional service model that IBM defined. A cloud environment is provided for an entire business process, such as supply chain management or a health care application. All facets of the business process are included in the offering, and might include help desk support.

Summary: Use a single-management platform to support one or more of the service models: IaaS, PaaS, SaaS, and BPaaS.

Cloud architectural principles

To address the cloud computing challenges it is necessary to define the principles for your cloud computing solutions. These principles are dictated by the key business drivers, cost

reduction, speed-to-market, quality of service, better use of IT resources, and the monitoring and metering of services.

To benefit from the implementation of a cloud computing solution a reference architecture is key because cloud computing is not just an assembly of existing patterns used in the traditional IT environment. A reference architecture helps to address your entire environment across the lifecycle of your cloud computing services, from creation to delivery, taking into account the cloud consumers' requirements through the various business channels.

The IBM Cloud Computing Reference Architecture, shown in Figure 4, defines the fundamental architectural elements that constitute a cloud computing environment, and it was recently submitted to the Opengroup community[2]. This architecture aggregates experience across hundreds of cloud client engagements and implementations. It is required that all infrastructure components be managed from a single, central Common Cloud Management Platform with the ability to place instances of each cloud service on the corresponding infrastructure. The Governance aspect is the foundation to address the challenges of security, resilience, performance, and consumability of the cloud computing solutions.

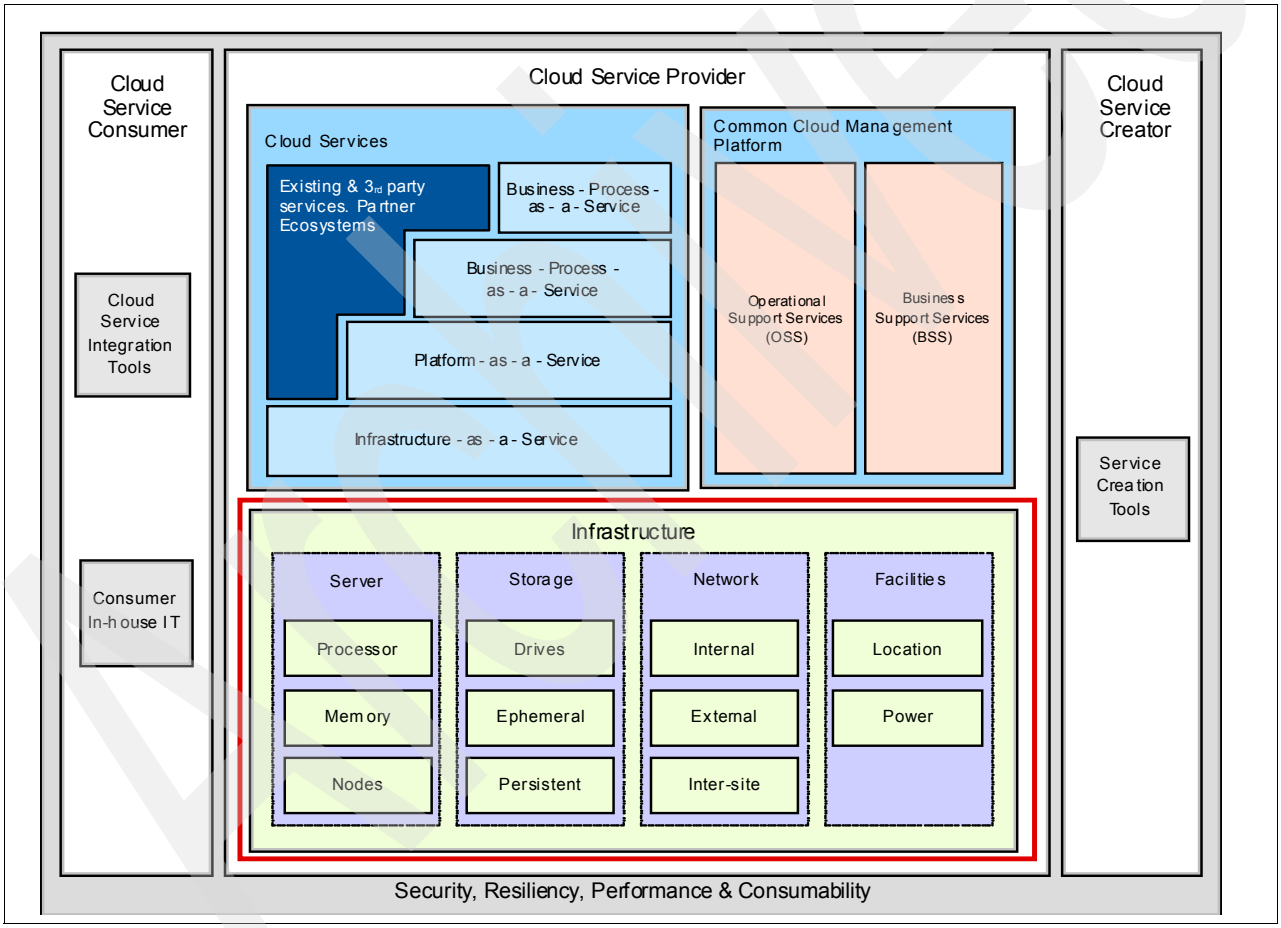

Figure 4 The IBM Cloud Computing Reference Architecture infrastructure details

The IBM Cloud Computing Reference Architecture describes four principles to consider solve cloud computing challenges:

- Design for Cloud-scale Efficiencies: Efficiency Principle
- Support Lean Service Management: Lightweight Principle

[2] https://www.opengroup.org/cloudcomputing/uploads/40/23840/CCRA.IBMSubmission.02282011.doc

- Identify and Leverage Commonalities: Economies-of-scale Principle
- Define and Manage generically along the Lifecycle of Cloud Services: Genericity Principle

Efficiency Principle

A complex environment has many interdependencies, high operational and maintenance costs, and consists of several manual activities and procedures for IT operations. A key objective of cloud computing is to change the operational inefficiencies in a traditional IT environment. A cloud management platform must be designed and implemented to manage all of these operational tasks; however, any changes to this environment or to provide a new IT system takes more time than is acceptable to businesses today. The implementation of a cloud solution following the efficiency principle can address these shortcomings and provide IT services with a high level of efficiency, flexibility, standardization, and automation.

Using zEnterprise you attain cost reduction and efficiency at all levels of your IT business services and benefit from the standardization of your operations with its unique integrated management solution.

The zEnterprise System demonstrates efficiency through the following elements:

- Virtualization is built into, not added on to, zEnterprise
- Easy consolidation of standardized images with multiple cloning options
- Integrated platform management across several servers architecture
- Electrical energy efficiency and space floor savings
- Dynamic resources allocation for optimized workload management
- Capability to scale and overcommit its resources
- High processor constant utilization rate

Summary: Using the efficiency principle you can drive down costs ($/ServiceInstanceHour) and time-to-response by orders of magnitude.

Lightweight Principle

Cloud computing promotes lightweight service management processes and technologies. As you introduce virtualization and automation in your IT environment, the need for a robust service management system increases. To serve many automated requests requires the ability to recover from failures and restore service or the operation. A single request can be easily handled manually. But you must be able to recover from all of the automated tasks of the different components, such as, network, storage, operating system, middleware, and applications.

You cannot simply map your existing processes to your cloud delivery model because those are often too long and require involvement from several support areas. You must select key parts of your existing processes and standardize them to be consumed during the cloud service delivery. An example is the approval process to handle any new cloud service requests. Standardize this process in a way to be consumed in an automated or semi-automated mode.

zEnterprise is the only system on the market that provides the same administrative console to your distributed and System z operational teams and the built-in standardization capability to manage your cloud computing environment. The cloud computing software stack that runs on this platform creates an enterprise service provided platform. You can benefit from zEnterprise for all management processes, such as incident and problem management, for your cloud self-service requests. The automated provisioning across various hypervisors on

both System z and non-System z, and the detailed accounting to chargeback services has run on IBM z/OS® and z/VM for years.

Summary: There is a need for the exploitation of a high degree of standardization in cloud environments to reduce management costs based on an Eliminate-Standardize-Optimize approach.

Economies-of-scale Principle

Cloud computing must be independent of specific product sets and utilize standard ones. When implementing a cloud solution you must identify the products that are back level and upgrade them. You also must identify all of the product functions to avoid supporting multiple products that have the same functions. The objective is to establish commonality in your cloud environment, which is the foundation for economies-of-scale.

As shown in Figure 5, as you share more of your infrastructure, it increases the utilization of your resources and reduces the duplication costs of maintaining different environments. As the levels of isolation decrease with the sharing of the infrastructure, you can than rely on a strong secure platform using a multi-tenant architecture. The relationship with the common "as-a-Service" model is also demonstrated.

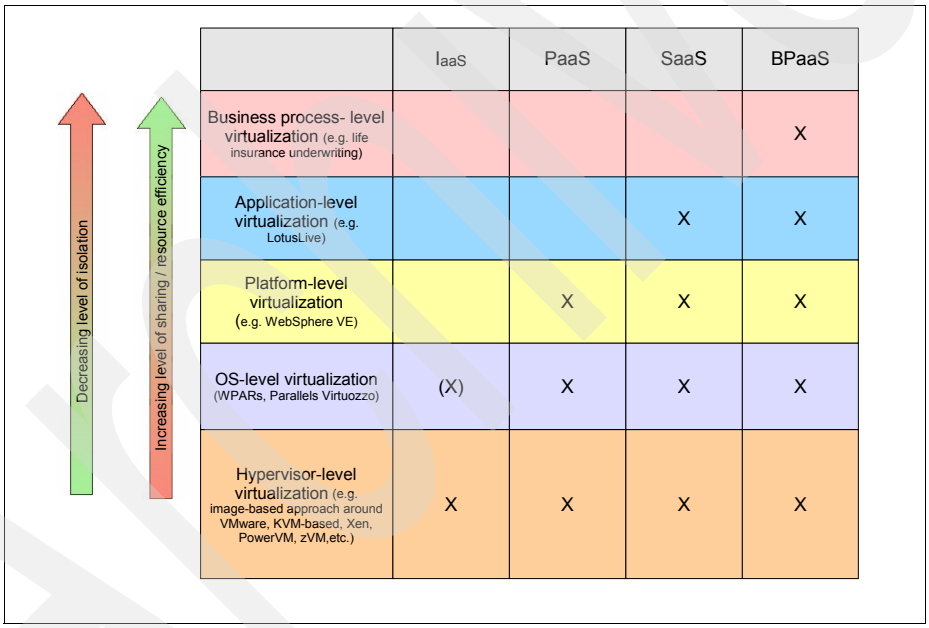

Figure 5 Increase your resources sharing for economies-of-scale

All commonalities are identified in the cloud service design. zEnterprise provides a high level of commonalities. On the System z platform your cloud services can benefit from IBM HiperSockets™ to communicate and exchange data between several partitions. Using this *in-memory* communication you can exchange data for your services at a high speed with minimal system and network overhead. Because you remain within the same server, there is less security exposures of your services and data. Since it eliminates the need for any physical cabling or external networking connection between the virtual connections, it reduces your costs as your cloud computing services continue to grow.

Summary: Use a maximum sharing of management components, infrastructure and managed platform across cloud services to reduce CapEx, OpEx, and time-to-market.

Genericity Principle

Cloud services must be generic across the different as-a-Service models and must provide a way to support the various cloud services using a common management platform. As you start to design your cloud computing service, consider the consumability of the service. The attributes drive the user experience and the quality of experience (QoE) of your cloud services. The QoE is critical for cloud computing services as opposed to traditional services. In cloud computing, the users consume the services through a self-service interface over a network connection using mobile computers, workstations, or mobile devices. Therefore poor response time and user experience can impact the ability of a business to compete and to grow.

A key market driver to consider is how your service can be used to do business easily and also provide a positive first-use experience. Another consideration is how to rapidly integrate with the existing environment and readily adapt to your customer requirements to simplify ownership and operation.

Consider your cloud service as a business workload, composed of several machines with different operating systems, middleware, and applications components. You must consider the entire service lifecyle to identify the challenges to address in your environment.

You do not need to create an offering for each configuration option for your cloud service. To be consistent in your service catalog offerings and provide a good user experience, you can use a generic definition of your service based on generic templates. At runtime, when the service is instantiated (delivered to the user), your cloud computing platform should take care of the appropriate options and implementations to use. You must be able to rely on a cloud computing service that is not just managing one kind of hypervisor or one kind of workload because your business is composed of mixed workloads.

The service instances must be managed during the service operation phase using a generic approach. Define the structure of your cloud service and plan for their management, which will be valid throughout their lifecycle. Make sure your service can run mixed workloads and optimize the deployments to servers on several architectures. This is a strong differentiator to enhance the quality of experience of the users and to improve the service quality. The management of your services instances must also support the evolution and the flexibility that is required by your business as requirements change.

zEnterprise uniquely positioned to address cloud challenges

The zEnterprise System offers all of these fundamental attributes and provides a solid foundation for a cloud environment. zEnterprise is a single-vendor solution that spans across multiple platforms and architectures, providing the best cloud solution available. It is a robust heterogeneous virtualized environment that makes it perfectly suited to host applications and workloads across a variety of architectures.

The capability of zEnterprise to run mixed workloads allows you to reduce the cost of running your IT infrastructure by optimizing the utilization of your resources. Throughout the day your business services and their workloads might not be at their peak simultaneously, and they can have various profiles.

zEnterprise provides the opportunity to optimize your cloud environment with the integration of the mainframe, IBM POWER®, and IBM System x®. The POWER servers with their RISC-based processors excel at compute intensive processing. The System z platform with its ability to off load IO processing is unsurpassed at processing input/output (IO) and mixed

types of workloads. System x is highly threaded and therefore suited to handle small discrete applications, such as web applications.

The management of this heterogeneous infrastructure occurs at a platform level from a single point of control with the IBM zEnterprise Unified Resource Manager.

In your cloud computing environment, you want the flexibility to grow or shrink the pool of resources on demand to help reduce the costs of running them. zEnterprise has an innovative and unique ability to handle platform and workload management across x86, Power blades, and System z architectures. Another key benefit of zEnterprise is the specialized application optimizers, for specific workloads, such as the IBM Smart Analytics Optimizer or the IBM WebSphere® DataPower Integration Appliance XI50 for zEnterprise (DataPower XI50z).

You have flexibility in the options to lower the costs of running your cloud services. As your cloud computing service catalog grows, you can control your operational costs. When properly implemented, you avoid the need and the costs to replace and integrate a whole set of distributed servers.

Cloud computing is not just about technology or automation; instead, it is about the integration and level of autonomicity between your service delivery capabilities. This autonomicity is first inherited from your underlying hardware and you must determine its reliability. zEnterprise can recover from memory errors, autonomously establish configurations, discover resources across several architectures, and automatically orchestrate the reallocation of resources. You can also decide to delegate these tasks to your application layer but this introduces complexity to your solutions and contains additional hidden costs.

The on demand capacity upgrade of the mainframe allows you to scale and adapt your infrastructure to the evolution of your cloud environment. By placing it at the core of your business applications, you have a solution to optimize your workloads while reducing your operational costs.

Summary: It is key that a cloud computing solution be based on three essential fundamentals: standardization, virtualization, and automation. zEnterprise is uniquely positioned to address these fundamentals and to address the challenges you face to provide a secure and reliable environment for your cloud services.

zEnterprise: The value of cloud in a box

The IBM zEnterprise System addresses the complexity and inefficiency of the multi-architecture data centers of today. As a heterogeneous compute system that can host many workloads that are integrated together, it provides the only multi-architecture cloud solution. zEnterprise provides the opportunity to gain the benefits of cloud computing today without the risk of uncertainty associated with most public cloud offerings.

zEnterprise made major strides in virtualization and scalability, and it can also deliver the proven qualities of service of the System z platform. Because it is a heterogeneous compute system, running a cloud on zEnterprise also extends the mainframe QoS to the distributed systems that are attached to it, providing greater levels of availability, reliability, and ease of hardware management.

zEnterprise provides multiple business values and benefits for cloud computing solutions. Figure 6 summarizes and describes the benefits of virtualization, security, availability, efficiency, and scalability. zEnterprise is able to deliver a heterogeneous cloud with these critical attributes all in one box.

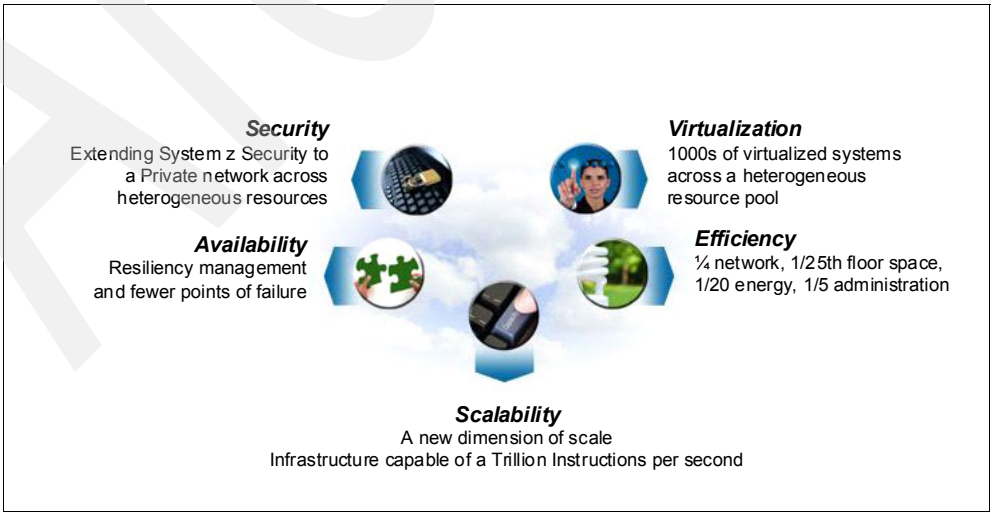

Figure 6 Value of zEnterprise for cloud computing

Virtualization is key

Virtualization is the first key step in building a cloud infrastructure. The fundamentals of virtualization have been part of the System z platform since its inception, and it is a leader in platform virtualization. This virtualization technology allows heterogeneous operating system instances to share the same system resources in a *shared everything* architecture while maintaining workload integrity.

The concept of a Virtual Machine (VM) defines the IBM ultimate virtualization hypervisor-based operating system IBM z/VM®. z/VM's ability to host large numbers of virtual machines makes it possible to achieve massive consolidation and same system growth with limited real resources.

Virtualization with zEnterprise allows the dynamic provisioning of resources, over allocation of real resources and the ability to easily consolidate standardized virtual images. zEnterprise can host more virtual machines with z/VM than any other platform, run thousands of virtual servers on a single IBM zEnterprise 196 (z196) or IBM zEnterprise 114 (z114), and can deploy virtual servers quickly and non-disruptively without the dependency to first acquire new hardware.

Higher security

Security remains a major concern for any cloud computing environment. For over 40 years the System z platform focused on providing security for applications and sensitive data in virtual environments. It is the most secure platform in the industry today, beginning with the hardware and integrated throughout the stack. System z has an Evaluation Assurance Level (ELA) of 5 for its logical partitioning (LPAR) technology, which remains the highest level of security certification in the industry. It also has an ELA rating of 4+ for the use of Linux under this environment.

The System z integrated hardware cryptographic coprocessors also provide additional levels of security. Linux virtual machines can take advantage of the same cryptographic coprocessors that z/OS uses.

The security of System z is extended to the high performance private network of zEnterprise, which helps to improve network security with lower latency. A private internal data network connects the virtual servers to the z/OS transaction and data services running in the same machine that support business critical data. This feature provides the opportunity to handle production workloads with more security than development or test workloads, and its security protects critical and sensitive business data in the cloud.

zEnterprise brings the value of reduced networking and simplified server control, resulting in tighter overall security for the complete cloud infrastructure. Its decades proven security stands out as a major reason for deploying a private cloud with zEnterprise.

Greater reliability and availability

Reliability and availability of resources are key factors when selecting a cloud delivery option. The System z platform has a long history of reliability and availability that enables it to host cloud environments at unmatched levels of service. System z has built-in hardware redundancy and fewer points of failure and is a proven reliable platform that provides high availability and achieves SLA's objectives. Memory and processors can be added and enabled dynamically on System z, allowing for continuous service and availability with no interruption. The System z platform five 9's of availability and Mean Time to Failure are

measured in decades. zEnterprise has improved and simplified cross-platform availability procedures.

With its upward migration compatibilities, System z elevates the process of recommissioning servers and reusing space in the data center. Another valuable System z feature is *push pull*, the server life cycle management capability to upgrade without major disruptions of service.

More efficient data center

The simplification of the infrastructure with zEnterprise results in greater efficiencies in the data center. zEnterprise reduces the energy, cooling, and floor space requirements, resulting in lower overall costs and greater operational efficiency.

With zEnterprise there are fewer parts or components to monitor. zEnterprise and IBM Tivoli® products can be used to manage the entire enterprise-wide IT environment, and the Unified Resource Manager can manage workloads across platforms.

Enterprise workloads are deployed on servers to work closely with business-critical data and to be located near where the data is housed. Cloud computing with zEnterprise supports this strong affinity to the z/OS environment that continues to be the backbone in many organizations. It allows moving the supporting distributed workloads closer to the data and to the processing.

Greater scalability

Another strength of the System z platform is its scalability, which meets changing business requirements and accommodates growth.Using System z you can dynamically add more capacity, increasing the elasticity of the environment. As your workload demands increases, vertical and horizontal growth become dynamic and automated, enabling your organization to focus on the business of software innovation rather than worrying about the acquisition and provisioning of new resources. With zEnterprise you can run multiple copies of z/VM on a single mainframe for enhanced scalability and failover capability for the virtual servers.

Increased productivity

zEnterprise provides increased productivity with easy technology refresh, rapid provisioning, and superior life-cycle management. Using the fast and easy technology refresh with zEnterprise you can react quickly to technology changes. The efficient, rapid provisioning of virtual and real system resources means you can quickly meet your business demands. Server life-cycle management allows for upgrades without disruption of service.

The increased productivity removes barriers to the rapid delivery of new services and provides you with the time to explore new technologies to drive innovation in your business.

Ease of use

The zEnterprise United Resource Manger provides integrated management with a single easy-to-use interface, which simplifies operations across all of the application environments. Another benefit of zEnterprise for cloud computing is that different workloads can run on disparate systems but behave and are protected like a single system. This also increases workload optimization because each application runs on the best-suited platform from an architectural point of view.

As priority resource control exists at guest, LPAR, or system level, maintaining a defined SLA for production or critical workloads for each defined environment is simplified because everything can run from within the same system. Problem determination is also simplified.

Development and test environments can be deployed on the same zEnterprise system with little or no impact to production workloads that run on the same box. After a release or project passes acceptance testing parameters they can be pushed right into production.

Cost savings

Many organizations are focused on reducing expenses. Implementing a private cloud on zEnterprise can dramatically reduce the total cost of ownership (TOC), specifically for IT and cloud operating expenses. zEnterprise delivers by offering the most consolidation per platform possible today, which results in significant cost savings.

Most cloud computing models include a pay-as-you-go or grow ability feature and requires the automation of the processes that define this ability. zEnterprise On/Off Capacity on Demand processors can be turned on temporarily to meet business peak demands. Payment for this additional capacity is only for the time period requiring the extra resources.

Cost savings are also obtained with the ability to run multiple instances of a software product per processor, reducing the per core charges for most workloads. The benefits with these economies of scale increases because the pool of available resources and systems are managed from a single place, the Unified Resource Manager. This can reduce the number of people required to maintain the infrastructure which helps lower the total administration costs.

The zEnterprise heterogeneous environment across platforms is optimized for workloads in a private cloud model and can result in a dramatic reduction in costs compared to like services offered from other cloud computing models. As shown in Figure 7, there is a 72% lower overall TCO when deploying different types of workloads on a private cloud on zEnterprise.

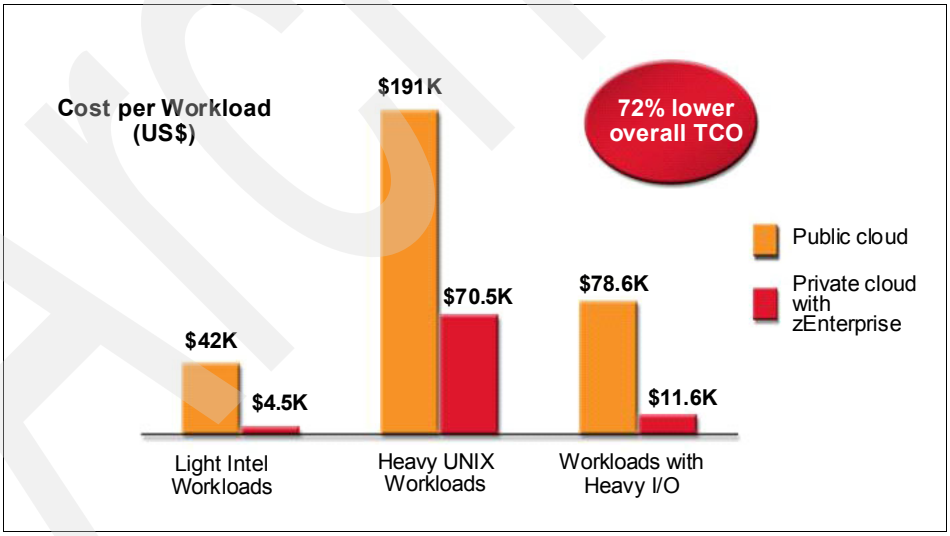

Figure 7 Private cloud savings with zEnterprise

Consolidating servers onto the mainframe can result in big savings. The Clabby Analytics research report explains *How to Save Over Million Dollars Using an IBM System z as a Linux Cloud Server*[1].

[1] https://www14.software.ibm.com/webapp/iwm/web/signup.do?source=stg-web&S_PKG=us-en-sz-ar-clabbysavemillion&S_CMP=web-ibm-sz-_-ws-zcloud

The zEnterprise end-to-end integrated architecture

The zEnterprise is a heterogeneous system that can host many workloads that are integrated together and managed as one single entity. The IBM zEnterprise System includes the zEnterprise CPCs (z196 and z114), the zEnterprise IBM BladeCenter® Extension (zBX), and the Unified Resource Manager. A zEnterprise can deploy and manage workloads across mainframe and distributed technologies with a single management interface.

System z environment and distributed systems

In 1964, IBM pioneered data processing with the introduction of the S/360, the first mainframe computer to be used for general business. In 1967 IBM extended the ability of the S/360 with the introduction of the first hypervisor, CP-40/CMS, later to become z/VM. This technology introduced the capability to run *virtual machines* through the use of virtualized hardware on the S/360.

In the 1980s, the mainframe was an island of computing with no connection to distributed systems. Computers were large, expensive, and required a dedicated team of specialized professionals for maintenance. Logical partitions (LPARs) were created introducing the capability for the System z platform to run workloads that were completely isolated and to run different operating systems in the same box.

In the 1990s, the mainframe was mainly considered as a back-office system and some distributed servers started running application servers to access the transitional subsystems, such as IBM IMS™ or IBM CICS®. The standardization of the platform was improved with the web enablement of traditional sub-systems and the era of e-business.

In the 2000s, as the IT industry transformed towards internet-based services, the System z environment continued its adoption of emerging standard technologies, such as Java (“write once, run everywhere”) and XML (“create once, distribute everywhere”). This adoption of standards and the associated birth of Linux on System z initiated the openness of the mainframe and it started to be considered as a part of the distributed application

environment. During this period, a set of best practices for the technical management and the automation of System z and distributed servers environments was established.

Many IT organizations still have two separate operational management processes, one for distributed systems and another for System z. Operational teams are organized in silo mode, not talking to each other, but they still must follow similar maintenance and support processes to deliver IT services and guarantee their quality of service.

An economic crisis at the end of the 2000s forced companies to accelerate their focus on cost reduction. The arrival of cloud computing during this period and its ongoing development changed the rules of IT service delivery. System z, already recognized as a key platform for critical consolidation projects, once again innovated to support IT organizations that wanted to reduce their costs.

With the introduction of the IBM zEnterprise System, this first *system of systems* enabled operational and capital expense reduction. It avoided the previous duplication of processes by standardizing and unifying the management processes across multiple server architectures. It follows two simple principles:

- The more you virtualize your IT infrastructure, the lower your capital expenditures
- The more you standardize your infrastructure, the lower your operating expenses

Figure 8 illustrates the positioning of System z and distributed systems. With the introduction of the zEnterprise System, the distributed systems are now part of the mainframe and they are managed as a single entity thus reducing operational management costs. You now have the capability to use standard processes to deliver your IT services across several server architectures while having the benefits of workload optimization. The zEnterprise provides a unique foundation for the new way to deliver IT services.

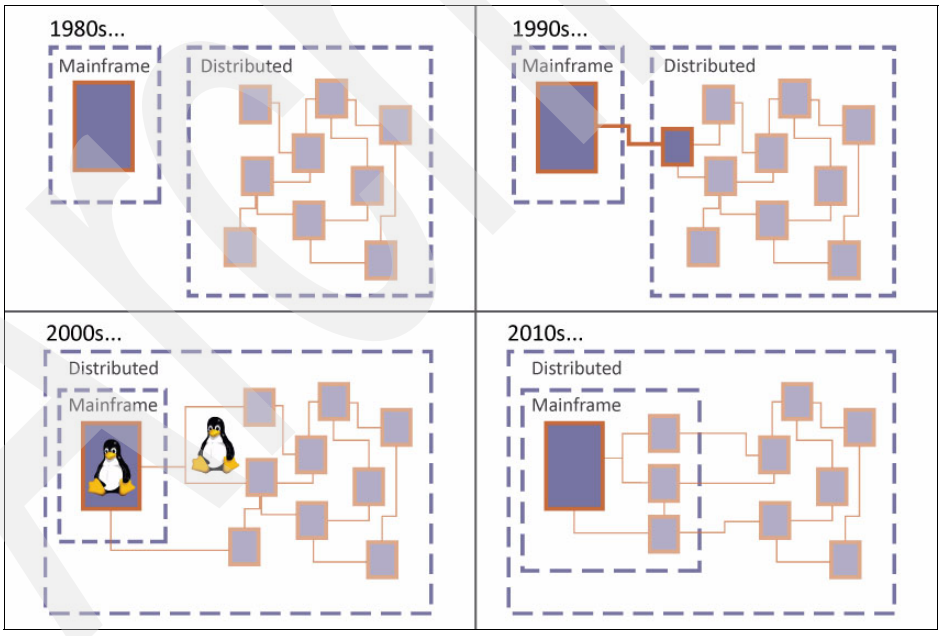

Figure 8 Positioning of System z and distributed systems

zEnterprise hardware

The zEnterprise is a *system of systems,* a heterogeneous compute system that can host many workloads integrated together and provide a virtualization environment in one system. The zEnterprise goes beyond traditional mainframe applications with the ability to run applications that are based on POWER and System x architectures. The zEnterprise can run z/OS, IBM z/VSE™, and z/TPF for traditional mainframe applications. The z/VM as a hypervisor to run thousands of virtual servers, z/Linux, both Red Hat Enterprise Linux and SuSE Enterprise Linux, to consolidate and virtualize Linux servers, IBM AIX® with IBM PowerVM™ on the POWER7® blades, and Linux (and in the future Windows[1]) on System x blades.

The zEnterprise System includes the zEnterprise CPCs (z196 or z114), the zEnterprise BladeCenter Extension (zBX) with its integrated optimizers, and the Unified Resource Manager.

The foundation of the z196 CPC is the 5.2 GHz quad-core processor. Each core can be assigned to a particular process, giving the enterprise the ability to assign capacity to workloads on a per-core basis. The partitioning capability of the z196 permits these cores to be shared among different processes, with processing weights assigned to allow access by mission critical workloads sharing a core with a non-production workload. The z114 employs the same technologies as the z196 and is suitable for a variety of workloads for mid sized businesses, which provides companies with an extreme level of virtualization and flexibility of resources.

The zBX is an optional blade center attachment to the zEnterprise CPCs, which houses IBM POWER7® blades or System x blades for general-purpose work and optimizers for specific workloads, such as the IBM Smart Analytics Optimizer or the IBM WebSphere® DataPower® Integration Appliance.

The zBX is directly connected to the z196 system with a 10 GbE private data network eliminating network hops. The management function is provided on a separate 1GbE network, therefore management functions and data are isolated from each other. The zBX is ideal for applications that require a POWER or System x environment and need to be tightly coupled with data residing on the mainframe.

The heterogeneous virtualization of the zEnterprise allows System z virtualization and resource management to be extended to other platforms. The zEnterprise CPCs (z196 and z114), and the zBX combine to form a highly-optimized platform for hosting workloads that benefit from using multiple architectures.

Hypervisor overview

A hypervisor uses a layer of code in the firmware to achieve dynamic resource sharing, providing flexibility in how virtual resources are defined and managed. It is the primary technology for system virtualization. IBM has worked on this type of technology since the 1960s with the evolution of z/VM. This type of methodology made LPARs possible.

The zEnterprise Unified Resource Manager runs on the Hardware Maintenance Console (HMC) and it works with the various hypervisor technologies, as shown in Figure 9 on page 24. The hardware virtualization of the zEnterprise resources is achieved through the various hypervisors—IBM Processor Resource/System Manager (PR/SM™), z/VM,

[1] All statements regarding IBM future direction and intent are subject to change or withdrawal without notice and represents goals and objectives only

PowerVM, and System x blades integrated hypervisor using Kernel based Virtual Machines (KVM):

- PR/SM supports the definition of logical partitions or LPARs in the zEnterprise CPCs so that many logical partitions can share the same physical resources.
- z/VM provides the capability to run operating systems, such as Linux on System z, z/OS, and others as *guests* of z/VM so that many virtual machines can share the same physical resources.
- PowerVM Enterprise Edition offers virtualization capabilities for AIX. It allows the creation of logical partitions on the blades and enables the sharing of resources between multiple operating systems.
- The Unified Resource Manager's integrated hypervisor (based on KVM) provides a virtualization solution for Linux on x86 hardware.

The links into the various hypervisors and indirect ties to z/OS, as required, create a complete heterogeneous environment. It is this hypervisor technology that makes consolidation with virtualization possible, which is a key step in building a cloud infrastructure.

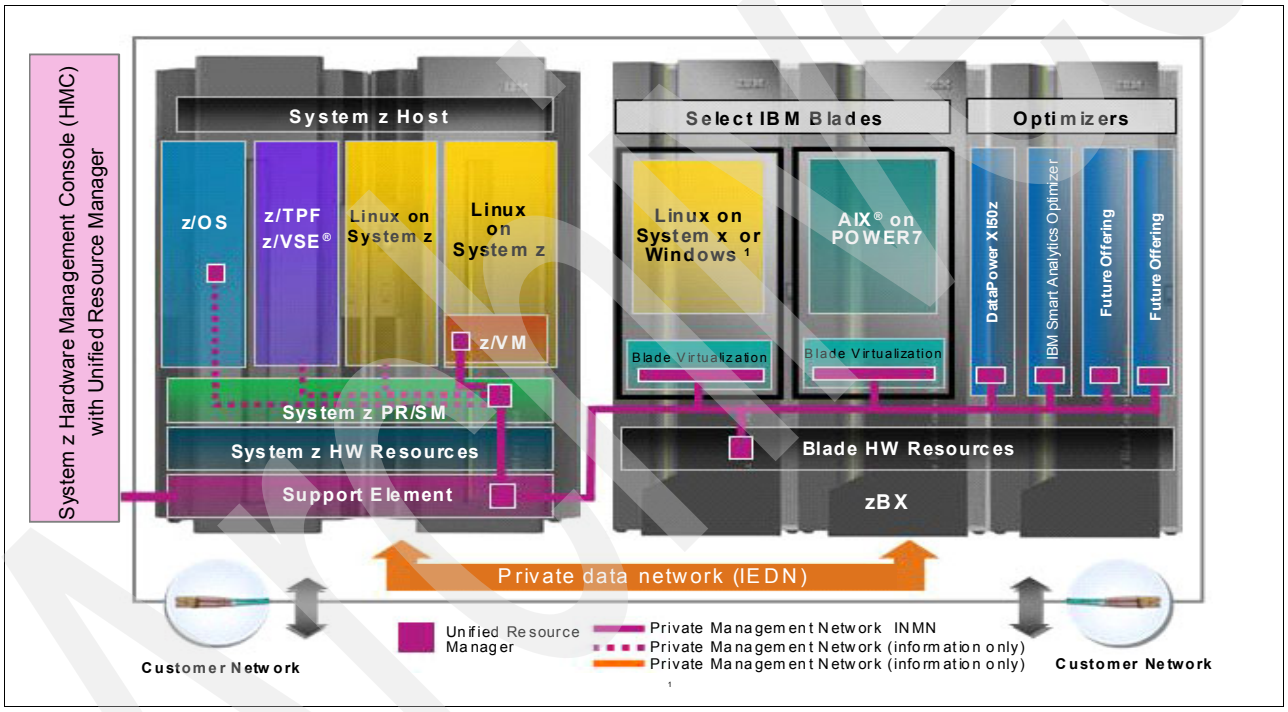

Figure 9 Inside IBM zEnterprise System

The ability of zEnterprise to manage all the real or virtual machines working under the control of the different hypervisors delivers the greatest value for cloud computing solutions. Although each hypervisor has its own functional control of the resources that reside within it, the ability to manage the workloads running across all of them is critical and the Unified Resource Manager can manage workloads across the different platforms. The workload management functions of the Unified Resource Manager for z/VM or POWER7 blades provides the ability to ensure that SLAs are met even during times of expected high demand on system resources.

Integrated Software Stack

A cloud computing environment not only relies on robust hardware or hypervisors, it must have the functions to ease the deployment and the management of your cloud services and for the platform management itself. These integrated service management functions provide the capabilities for centralizing management of the operating systems, middleware, storage, networks, and other resources that are needed to deliver business services.

Unified Resource Manager

The zEnterprise Unified Resource Manager, Figure 9 on page 24, can be used for platform management and for virtualization management functions.

Unified Resource Manager consists of six management areas:

- Operations: Operational controls for various management functions
- Virtual server: Directed and dynamic virtual server provisioning across hypervisors from a single point of control
- Hypervisor: Management of hypervisors and support for application deployment
- Energy: Energy monitoring and management capabilities
- Network: Create and manage virtual networks
- Performance: Management of CPU resources across virtual servers hosted on same hypervisor to achieve workload performance policy objectives

This Unified Resource Manager has some APIs for integration with the software described in this section. These APIs enable the unique end-to-end integrated capabilities of the zEnteprise solution that are particularly important in a cloud computing environment.

Cloud computing integrated software stack

Figure 10 on page 26 illustrates the System Director, the Provisioning Manager, and the Service Automation Manager in the Cloud computing integrated software stack and demonstrates the zEnterprise capabilities for cloud computing “management from” and “management to” environment.

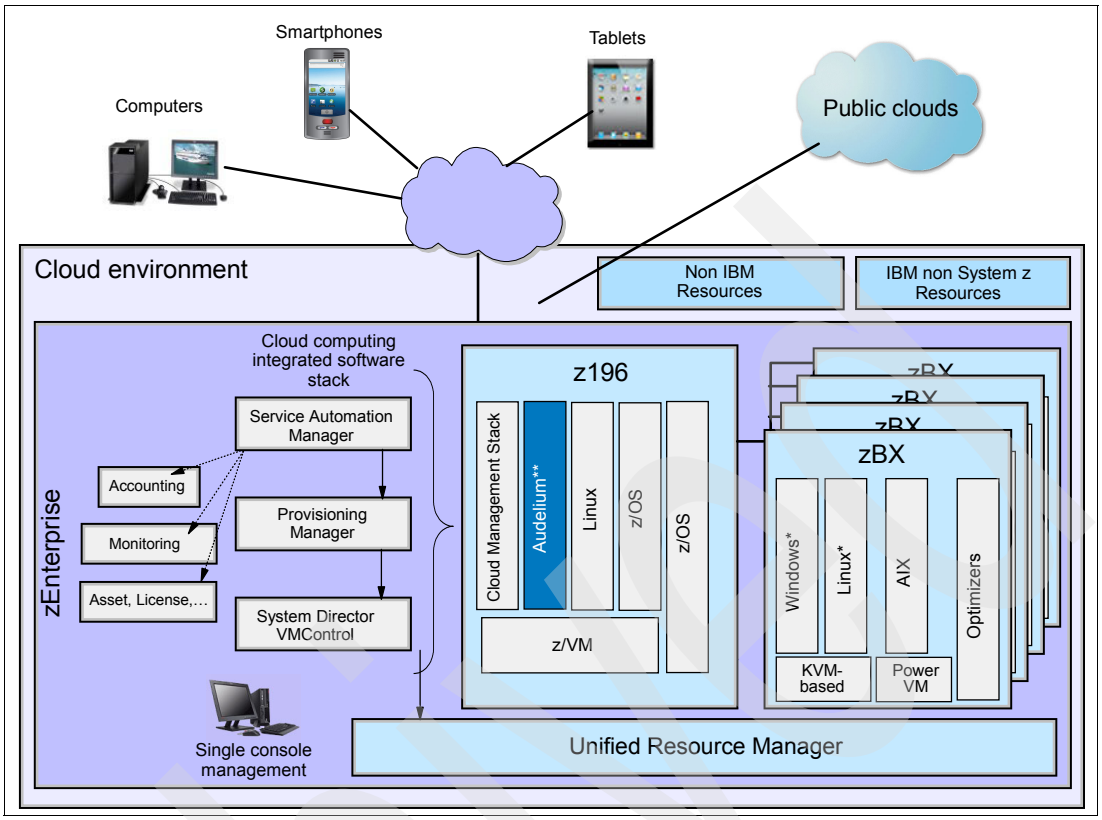

Figure 10 Overview of zEnterprise cloud computing capabilities

Provisioning management

The Provisioning Manager component provides a solution for the automation of middleware, applications, physical switches, storage, and so on, for the technical deployments across your datacenter. It can be placed directly on top of your current hypervisors management suite or on top of the previous virtualization management suite for example, IBM System Director VMControl. It provides a wide range of automation capabilities and comes with a code to support the automation of data center physical and logical assets (OS, software, patches, and so on). You can implement cloud computing automated deployments and benefit from the integration of your entire software stack down to the Unified Resource Manager, a key component of the zEnterprise.

Service Automation management

In the cloud computing software stack, the provisioning management layer is integrated with a service automation management suite. This solution adds the service management layer you must automate, not only the technical deployments but also the service delivery processes. It has a self-service interface to present your cloud services catalog to your end users. It also has a service request management component to keep track of your requests and to integrate with incident and problems managements activities.

Virtualization management

In your cloud environment, you might need to provide virtualization management for IBM and non-IBM resources. To extend the management of virtual machines and their operating

system images across several hypervisors, for example, z/VM, PowerVM, ESX (VMware), Hyper-v (Microsoft), Xen. A virtualization management suite, such as IBM System Director VMControl, can complement the Unified Resource Manager to use the same console for zEnteprise and non-zEnterprise hypervisors.

Summary: Using the integrated service management stack you can have end-to-end efficient management of your cloud computing environment. You can establish a bridge between the cloud services for your end users in the service catalog and the technical layer in charge of automating the tasks onto your managed infrastructure. This link between the business services and the technical services improves your agility and speeds up your time-to-market, as this integration is provided on the IBM zEnterprise System.

How others started the journey

Building a smarter planet is the IBM vision on how intelligence is being infused into the systems, processes and interconnected technologies that make the world work. We must think and act in new ways to make our systems more efficient, productive, and responsive.

Organizations are under pressure to make more intelligent choices. How did others start this journey? How did they make their choice for an innovative infrastructure to address the challenges of today and the opportunities of tomorrow?

The cloud computing journey with the IBM zEnterprise System allows you to:

- Improve service: High availability and quality of existing services and meeting customer expectations for real-time and dynamic access to innovative new services.
- Reduce cost: Containing cost and achieving breakthrough productivity gains through virtualization, optimization, energy stewardship, and flexible sourcing.
- Reduce risk: Managing the risk requirements of today, with security, resiliency and compliance challenges and getting in front of new risks of a more connected and collaborative world.

Customer cloud implementations

The trend is clear, the System z platform is being used to deliver and manage cloud computing services in various businesses and industries, such as banking, insurance, transportation, utilities, telcos, media, and gaming. The matching between the cloud computing attributes and the IBM zEnterprise System provided an innovative and competitive cloud solution.

In this section, we highlight some customer cases from various industries—Nationwide Insurance, Transzap, Inc., and Universita di Bari. You can get more information about these cases at:

http://www-03.ibm.com/systems/z/success/index.html

Flexible and evolutive platform for Nationwide Insurance

Nationwide Insurance wanted to increase its agility in delivering IT Services to address the pain points of the increasing back log of resource requests, the high amount of time spent by the management resolving resource issues, and the problem of server proliferation; therefore, Nationwide Insurance deployed System z mainframes running Linux.

To address these pain points, Nationwide Insurance now runs hundreds of virtual Linux servers and offers a broad set of tested-and-certified Linux images as part of their catalogue of services:

- WebSphere Application Server
- WebSphere Portal Server
- Apache Server
- DB2 LUW Server

Virtual server images are deployed in minutes after a request is submitted using pre-approved templates and a *shared root* file structure made possible by running Linux on z/VM. Services are monitored and managed with complete back up capabilities to run in either of their two Class 4 Data Centers. Services are terminated and resources are reclaimed when user contracts expire.

Nationwide Insurance successfully migrated their Linux on System z cloud infrastructure to new generations of mainframe technology with incredible speed. In 45 minutes the Nationwide team moved their System z Linux infrastructure onto their new System z machines, upgrading hundreds of virtual servers with no application breakage. Several of these applications moved multiple times from Linux-on-z to AIX and back to Linux-on-z to satisfy various business requirements.

A secure and scalable environment for Transzap

As a small business with tens of billions of dollars in client transactions flowing through their systems each year, Transzap, Inc. needed an economical, reliable platform to provide clients with high availability, while enabling the capacity to accommodate growth within their Software as a Service (SaaS) business model. The company is positioned as a leader in SaaS provider of ePayable, digital data and spends analysis solutions. It supports 44,000+ users in 4,200 companies. The company processes $80 billion in transactional detail and meets enterprise service level agreements with 24/7 operations.

System z provides a trusted foundation for secure operations from service provider to consumer. Transzap decided to consolidate on an IBM System z platform to provide the stability and scalability needed to accommodate triple digit volume growth, enabling them to focus on the business of software innovation. Transzap migrated to System z and virtualized its critical applications on Linux on System z, a platform that supports Transzap's dynamic Java and Oracle environments.

An enabler of innovative services for Universita di Bari

The Universita di Bari is strongly committed to developing cloud-based solutions for communities and businesses in southern Italy, as shown in Figure 11 on page 31. The University needed a platform to facilitate cost-effective, flexible application development. The objectives were to provide cloud services for small and medium business located in the south of Italy. To provide services on demand, manage the charge back and billing of services, and develop a new solution for local markets.

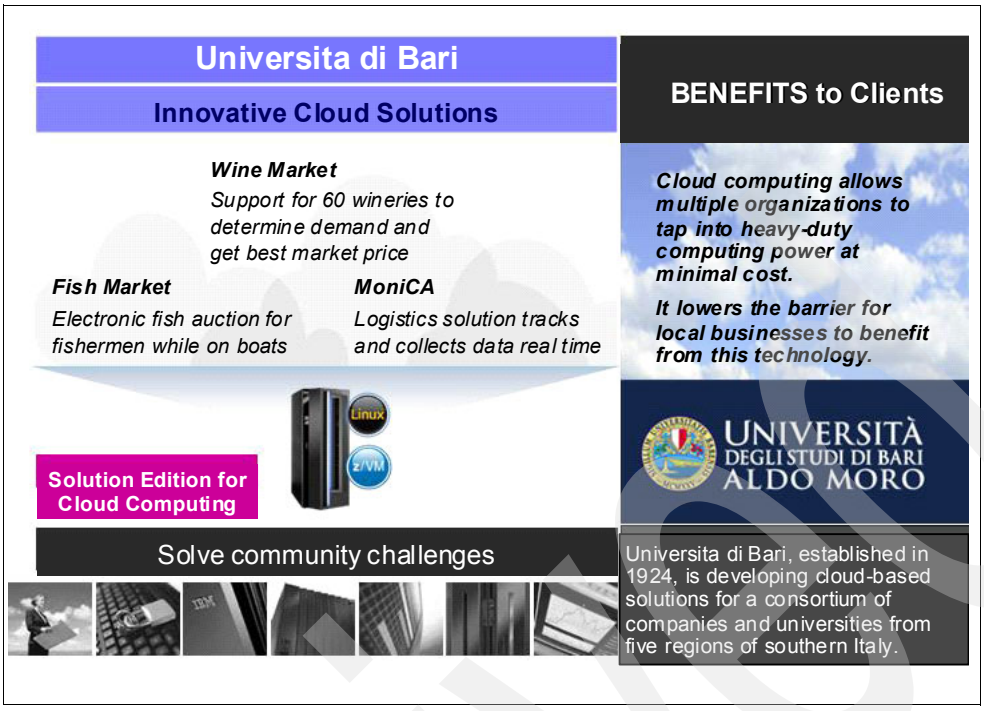

Figure 11 Universita di Bari cloud implementation

The University leveraged the IBM System z Solution Edition for Cloud Computing — a virtualized infrastructure that uses IBM System z, IBM System Storage®, SUSE Linux Enterprise Server for IBM System z and IBM Tivoli Service Automation Manager to enable intelligent management of Linux virtual machines. The System z cloud enabled the development of innovative applications for the local fishing, wine-making, logistics industries, and for the University itself.

IBM cloud transformation cases

In this section, we will review a couple of IBM internal cloud computing implementations:

- Smart Analytics Cloud
- IBM Smarter Banking™ Cloud Mobile Payment

Smart Analytics Cloud

The proliferation of business intelligence (BI) and performance management tools and deployments around the enterprise often results in multiple tools and (if not more) overall deployments. Departmental BI solutions are common place in large enterprise accounts.

It is difficult to support and comply with corporate and regulatory standards when no two BI deployments in the enterprise look exactly alike. There are also rising costs for delivering BI/analytics/performance management within the enterprise because each instance requires ongoing operations costs plus initial start-up costs, which include hardware, set-up costs, software, and so on.

Blue Insight is the IBM BI and analytics private cloud deployment that serves as the template for the Smart Analytics Cloud for System z, as presented in Figure 12 on page 32. IBM decided to adopt a corporate wide BI strategy created on cloud computing that is supported

(and attractive) at the highest levels in the company. It enables a core BI service that is standardized across the lines of business without removing the LOB (line of business) control to own and manage their own end-to-end BI solution. This centralized service strategy for BI supports simpler enforcement of corporate and regulatory standards, including security and boarding processes. It provides a BI report usage-based pricing model, which enables departmental budgets to stretch further and cover more users.

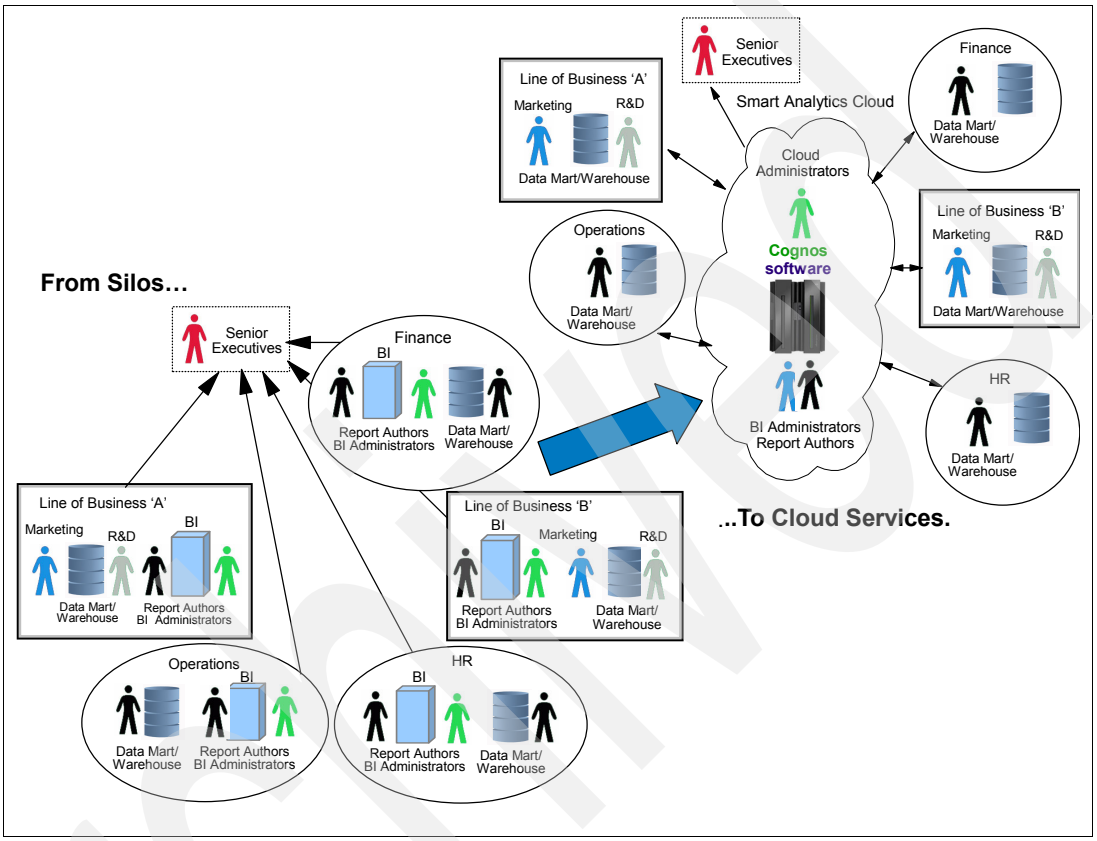

Figure 12 Smart Analytics Cloud offering

Having the solution hosted on System z, reduced overall costs, increased operational efficiency and provided a centralized service strategy based on economies of scale. You can get more information about this solution at:

http://www.ibm.com/systems/z/solutions/cloud/smart.html

Smarter Banking Cloud Mobile Payment

The IBM Systems and Technology Group Product and Solutions Support Center (PSSC) in Montpellier implemented a cloud computing solution to resolve challenges with big data and improve existing services for smarter banking mobile payments.

The following challenges were addressed:

- Build a catalog of mobile services (expose several applications, such as payments alerts, contactless payments, loyalty transactions, and authentication/biometric).
- Sustain a load/traffic increase on the mobile banking platform by providing new resources on demand.

- Create integrated services with core banking and payments channels.
- Extend existing systems management (Tivoli Suite and Unified Resource Manager) and HA/DR services to Mobile Banking Services.

Mobile Banking consists of multiple services offered across a range of devices. The Mobile Banking Services within a cloud are shown in Figure 13. By providing a common consumer (user) and provider (Admin) portal, subscriptions and services are offered to clients from a catalog that is role based. For example, Retail Banking Private clients are offered services, such as payments, transaction management, alerts, or loyalty schemes. Corporate clients are offered business dashboards/reports or ledger services, and Financial Market clients are offered trading information.

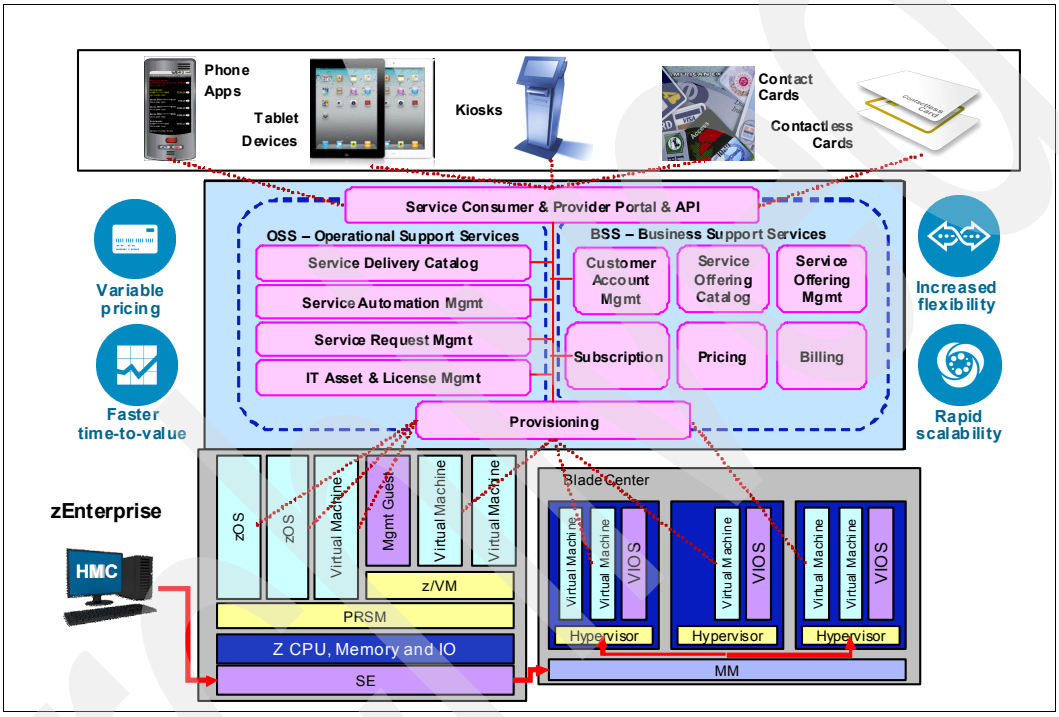

Figure 13 zEnterprise at the core of the Smarter Banking Cloud Platform

Service already running and others that are planned:

- Payments Alerts: WebSphere Business Events evoke user intervention using a Smartphone application. CICS, DB2, and VSAM transactions are captured as ATOM feeds and consumed through a Mashup process into a message that is processed by a mobile Android application.
- Contactless Payments: Use a mobile phone for low-value purchases and customer service transactions. This solution is based on contactless (NFC-communication) technology. A web service is exchanged between a device and terminal. This XML-based transaction is processed by a payment application accessing the core banking systems through WebSphere Application Server for z/OS.
- Multi-factor mobile authentication: Use mobile phones and terminals to enhance authentication (PIN + biometric identification). This involves capturing voice prints through a terminal and converting to web services for processing against a Linux on Intel application before integrating with the core-banking systems.

A benefit is the ability to build a repository of reporting services to cater to the varied roles and requirements across the retail banking sector. The optimization of resources by exploiting a common analytics infrastructure that can access information resources across the

environment. These services are integrated with core banking and payments transactions. The automation, provision, and management of the services are managed within the cloud. These services and the transactions themselves can be provisioned dynamically upon a number of zEnterprise hybrid locations, including z/OS (for example, WebSphere Business Events), Linux on z (Cloud Management with TSAM), POWER, and Linux or Wintel[1] on zBX. All these services are taking advantage of the superior Quality of Services, such as scalability, security, reliability, availability, and serviceability of the zEnterprise System.

System z cloud computing offerings

IBM has several System z cloud computing offerings, as shown in Figure 14, to help companies consolidate workloads on a large scale and provide superior qualities of service.

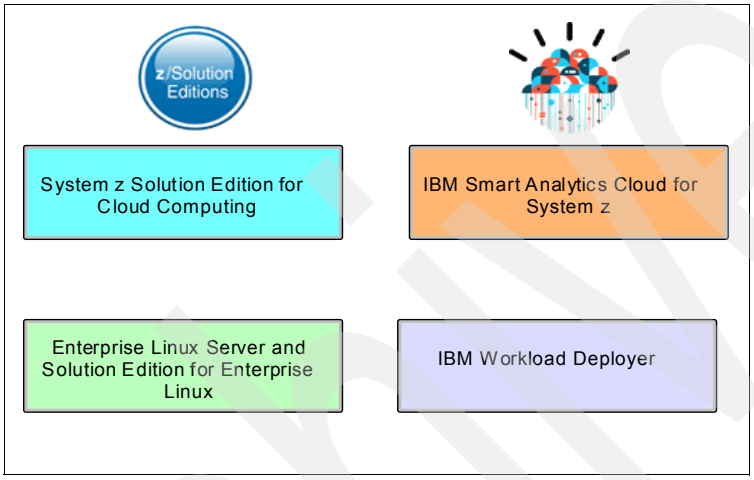

Figure 14 IBM System z Cloud Computing offerings

The key difference of the IBM System z cloud computing offerings are:

- A workload-optimized approach with systems design, management, and deployment. It considers the various workload needs and offers special configured and optimized systems for individual workloads. These offerings accommodate the fact that your infrastructure is always heterogeneous.
- Flexible sourcing options enable you to choose between an appliance (hardware, software and services) or use your own infrastructure to build up a private cloud.

Further details about the System z Solution Editions are available at:

http://www-03.ibm.com/systems/z/solutions/editions/

Enterprise Linux Server and Solution Edition for Enterprise Linux

Enterprise Linux Server and Solution Edition for Enterprise Linux is a system offering that provides a basic level of cloud infrastructure support that well suited for deploying a development or test cloud. It enables clients to add a similar solution stack of hardware and software to an existing System z machine and obtain the same business advantages of consolidating workloads on System z using Linux. This can also be a good option for existing mainframe clients if they plan to deploy a large virtual Linux server environment.

[1] All statements regarding IBM future direction and intent are subject to change or withdrawal without notice, and represents goals and objectives only.

System z Solution Edition for Cloud Computing

The Solution Edition for Cloud Computing is an extension to the Solution Edition for Enterprise Linux / Enterprise Linux Server and delivers a competitively priced offering providing a services automation and management framework for the adoption of cloud computing workloads. This offering can be used as the foundation to provide both a “managed from environment” and a “managed to environment” on System z, while also being extendable to manage other IBM platforms. It is a a cloud computing foundation solution that can be customized by the client for a wide range of cloud workloads.

IBM Smart Analytics Cloud for System z

IBM Smart Analytics Cloud for System z a cloud computing solution for the delivery of business intelligence and analytics optimized for the large enterprise client. The Smart Analytics Cloud creates a private cloud business intelligence solution built on the System z platform providing greater efficiency at less cost and resources.

IBM Workload Deployer

IBM Workload Deployer V3.0 (previously known as IBM WebSphere Cloudburst Appliance) is an appliance that dispenses and manages IBM middleware to virtualized servers and private cloud computing environments. It can be used for the rapid provisioning of WebSphere workloads onto System z.

Audelium STG Lab service offering

Audelium stands for Automated Deployment of Linux and z/VM systems. It is an environment running in a z/VM LPAR, which provides rapid and safe provisioning functions for Linux or z/VM systems. This offering consists of a set of IBM services, providing the installation and the configuration of the Audelium package on an existing z/VM or z/OS LPARs. This offering also includes onsite skills transfer to the customer including mainframe and open systems teams. Using this service, the customer gets an initial environment to deploy a set of automated infrastructure services.

The Audelium package provides functions to provision Linux and z/VM systems in the local z/VM hypervisor or in a remote location, such as another z/VM or a native LPAR. It also provides capabilities to move systems between an Audelium LPAR and other managed z/VM hosts.

This Audelium Lab service offering is provided on-site and provides you with the following benefits:

- Multi-tenant environment capabilities
- Simplified administration
- Highly virtualized environment
- Optimal usage of HW resources by pooling of resources
- Simplified network using internal network at high speed
- Improved security, less actions directly on the system, users can use Lightweight interface
- Improved service delivery through automated deployments
- Benefits of collocation with existing mainframe workloads
- Allows an incremental cloud computing infrastructure setup on System z
- Integrates with IBM service management software
- Optimize your workloads with consolidation under zVM

Further information is available in the section *Other Key Services for IBM System z* at:

http://www-03.ibm.com/systems/services/labservices/platforms/labservices_z.html#show-hide

Summary

For companies of today cloud computing is rapidly becoming the choice to deliver IT without boundaries. Any enterprise can now implement cloud computing to reinvent the way they do business and to drive innovation. Cloud spending is expected to grow at a compound average rate of 26% from 2009 to 2015, compared to just 3% for traditional on-premises IT spending. As with any technology implementation, there are challenges as organizations adopt cloud computing. This Redguide publication highlights the challenges an organization must be prepared to face and explains how the IBM zEnterprise System is uniquely able to address these challenges while maintaining your required traditional IT infrastructure.

There are common cloud principles that must be considered when addressing the challenges of implementing cloud solutions. The zEnterprise through the nature of its virtualization and management capabilities meets those cloud design principles. The System z platform has a long rich heritage of supporting multiple workloads running simultaneously and has been on the forefront of virtualization for decades. Those strengths evolved the zEnterprise into a powerful, scalable system that can support thousands of virtual servers through its ability to virtualize all resources.

The zEnterprise system provides an enterprise class heterogeneous platform spanning the mainframe, UNIX, and x86 technologies. At the underlying hardware layer, the System z platform is designed to run in a partitioned mode. All processors, networking, and IO resources can be shared across similar partitions. In each partition, System z supports multiple operating system types. The nature of its heterogeneous environment positions the zEnterprise to host all types of workloads under a single management interface.

The value of the zEnterprise for cloud computing is based on its ability to address the challenges that arise in implementing cloud environments. With the System z platform QoS and reliability it allows you to focus more on your enterprise and less on your IT infrastructure.

The zEnterprise, a *system of systems* with its built-in management capabilities, is perfectly positioned to meet a company's cloud computing requirements. With its heterogeneous environment, virtualization, integrated hardware, hypervisors, and management software, together with the traditional reliability, availability, scalability, and security, of the System z platform, it really is a cloud in a box.

Other resources for more information

For more information about transforming your business with cloud computing on IBM System z, refer to:

http://www-03.ibm.com/systems/z/solutions/cloud/
http://w3.itso.ibm.com/abstracts/redp4711.html?Open

For more details about the IBM Smart Cloud and cloud services and solutions refer to:
http://www.ibm.com/cloud-computing/us/en/.

The team who wrote this guide

This guide was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO).

Elsie Ramos is a Project Leader at the International Technical Support Organization, Poughkeepsie Center. She has over 28 years of experience in IT supporting various platforms, including System z servers.

Kurt Acker is a certified IT Specialist who currently works in technical sales for the implementation of new workloads on System z. He started working for IBM in 1999 helping with the design and implementation of web-based services for the payroll department before becoming part of the z/VM service, operation, test, and performance teams.

Robert Green is a Cloud Solution Architect in the US. He has over 20 years of experience in the IT field, 15 of those with IBM. He started at IBM in the OS/390® development laboratory and then moved to technical sales for the mainframe. He now architects cloud solutions across all IBM platforms. His focus continues to be in virtualization and consolidation on the mainframe.

Sebastien Llaurency is an Integration Architect. He joined IBM in 1999 and started to work in the IBM Montpellier PSSC Customer Center with customers on projects using leading edge technology. He currently works at the New Technology Center, PSSC in Montpellier, France. He participated in the creation of the Audelium STG Lab Services offering, a solution that extends the provisioning capabilities on System z in a cloud computing perspective. He is now leading virtualization and IT Optimization for the zEnterprise team and is working on cloud computing architectures with several customers from different industries.

Thanks to the following people for their contributions to this project:

Ella Buslovich
International Technical Support Organization, Poughkeepsie Center

Douglas Lin
Hank Meetze
IBM United States

Stay connected to IBM Redbooks

- 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 IBM 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 give 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.

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-4763-00, was created or updated on September 9, 2011.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at
http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

AIX®	OS/390®	System Storage®
BladeCenter®	PowerVM™	System x®
CICS®	POWER®	System z®
DataPower®	PR/SM™	Tivoli®
DB2®	Redbooks®	WebSphere®
HiperSockets™	Redguide™	z/OS®
IBM®	Redbooks (logo) ®	z/VM®
IMS™	Smarter Banking™	z/VSE™

The following terms are trademarks of other companies:

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Microsoft, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java, and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.