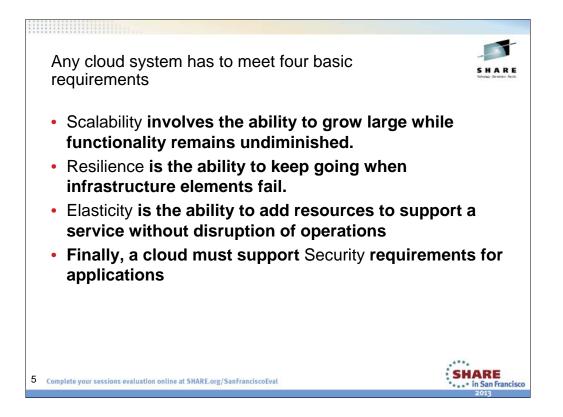
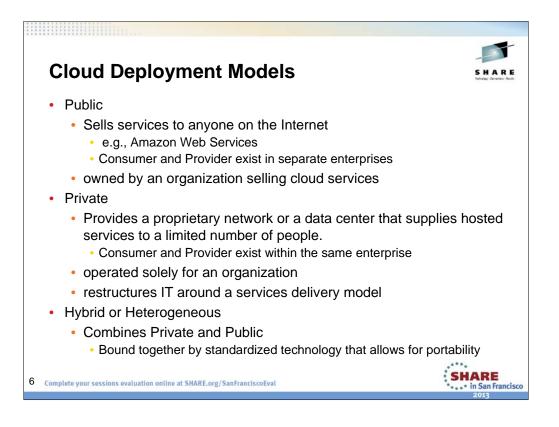


The term "Cloud" has a long—and growing—list of definitions. The National Institute of Standards and Technology (NIST) defines it as "a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources...that can be rapidly provisioned and released with minimal management effort or service provider interaction."1

The training organization Global Knowledge says cloud "can be broadly defined as several different methods to deliver information or services," and also describes it as "a new way of doing business that takes advantage of building efficiencies into the system that can scale out to provide services for multiple

companies."2

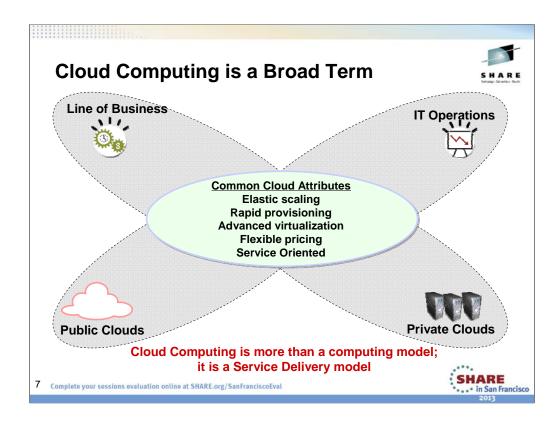




Public cloud. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Private cloud. The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.

Hybrid cloud. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).



Clouds mean different things to different people...

Regardless of roles and perspectives, the expectation for clouds include:

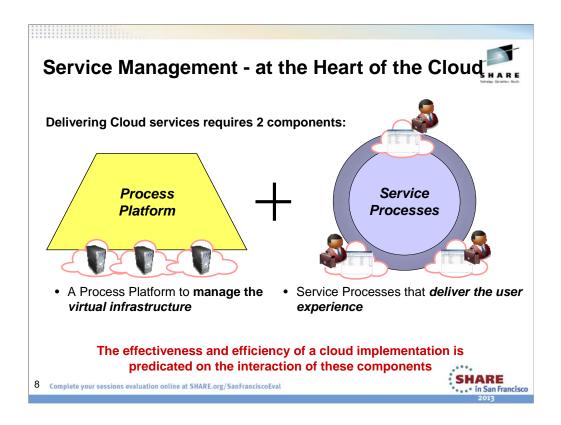
•A high quality of service delivered quickly and inexpensively and securely

•Support for business demands which include time to market, flexible costing and implementation transparency

•Support for IT Operations which demand scalability, security, lower cost and efficiency

Public clouds are concerned with delivering something simple but which can scale well –e.g. storage cloud.

Private clouds are about specific services (core and mission critical) being delivered securely and in a cost effective manner.

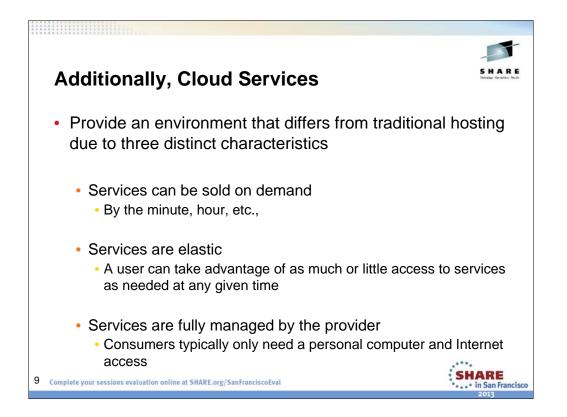


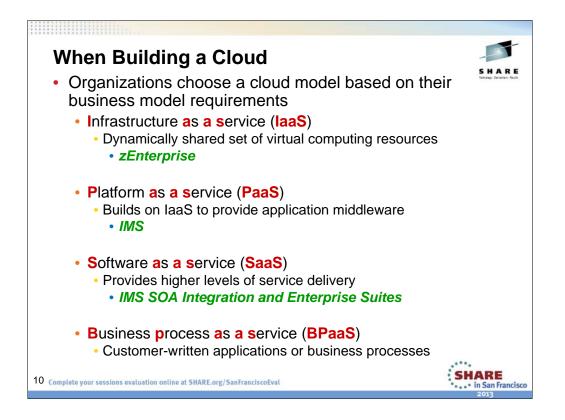
Process Platform:

- Discovery, monitoring, provisioning and security of virtual environments
- Optimization of resources, energy consumption and service delivery

•Service Processes

 Manage services as first-class objects across phases – definition, instantiation, deployment, management and termination



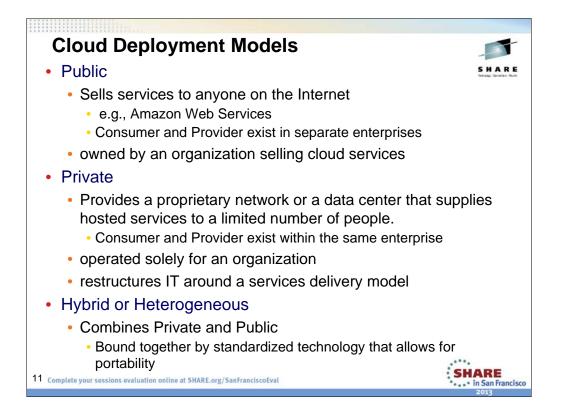


There are various cloud models that an enterprise can choose to create. The choice of which model depends on the enterprise's business needs.

•At the most basic level, infrastructure as a services(laaS) provides computing resources as a virtual set of resources that are shared and dynamically provisioned.

•Platform as a service (PaaS) builds on this model, providing application middleware (including databases and development tooling) as blank canvases ready for the deployment of a specific application.

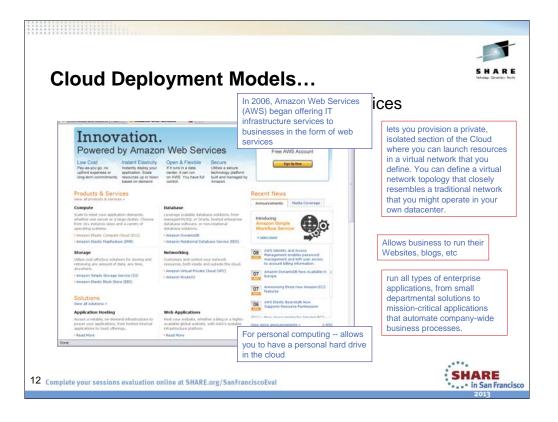
•Software as a service (SaaS) and business process as a service (BPaaS) enable higher levels of service delivery based on specific customer-written applications or business processes.

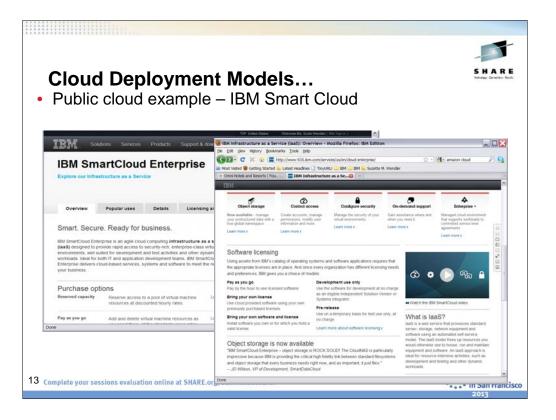


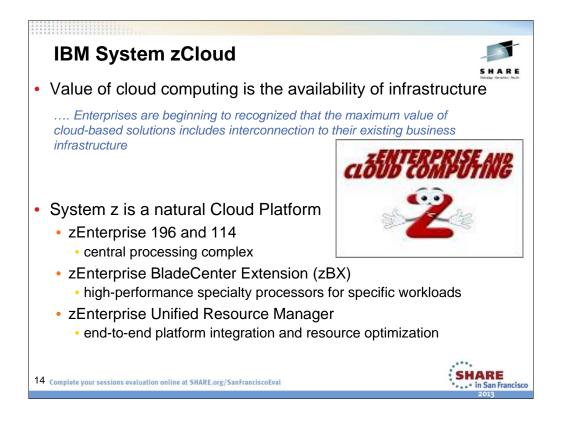
Public cloud. The cloud infrastructure is made available to the general public or a large industry group and is owned by an organization selling cloud services.

Private cloud. The cloud infrastructure is operated solely for an organization. It may be managed by the organization or a third party and may exist on premise or off premise.

Hybrid cloud. The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

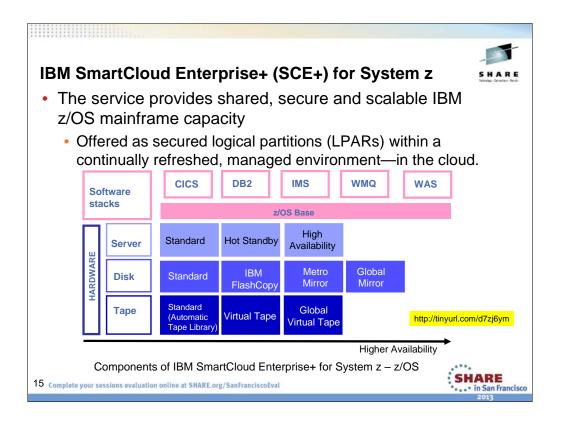


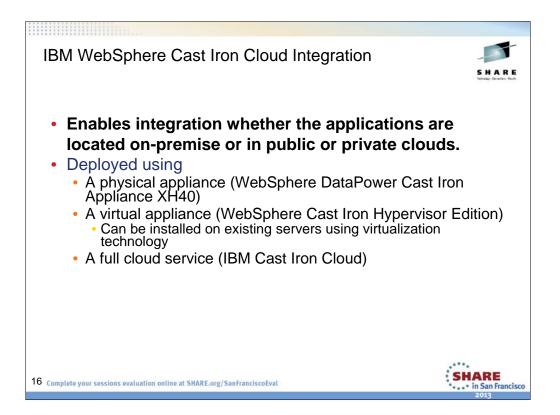




The value of cloud computing is the availability of technological infrastructure support that is always available, comprehensively useful, highly scalable and attractively priced. By this modern definition of cloud computing, System z has been an internalized cloud for decades and has found to have all the features to make it a natural Cloud Platform. The power of the mainframe's software and hardware virtualization, its shared resources architecture, its real and virtual memory controls, its in-memory internal communications structure, its specialty engines, its energy efficiency and its accommodation of Linux and Java natively are all cloud-relevant, and enable IBM System z machine to be a natural zCloud.

- The zEnterprise system enables clients to combine IBM's multiple processor architectures into one integrated pool of virtual resources, all managed under a single management framework. The zEnterprise system is comprised of three key elements:
- 1) The core server that comes in two fl avours (z196 with 80 cores and z114 with 10 cores),
- 2) The zBX which provides integrati on with IBM blade and accelerator technology, and
- 3) The Unified Resource Manager (zManager) which is firmware that integrates and manages the system as a unified set of virtual resources. The release of the core z196 represents a technology milestone for IBM, and the industry as a whole.
- The z platform boasts the fastest CPU in the world, running at 5.2GHz, and introduces some completely new capabilities, most notably the introduction of RAIM (Redundant Array of Independent Memory) which offers the same support for reliability for RAM as RAID does for disk storage.





Short for IBM WebSphere Cast Iron Cloud Integration, IBM Cast Iron is an offering from IBM that provides clients with a platform for integrating cloud-based applications from leading SaaS providers with on-premise applications from IBM and other companies. IBM acquired Cast Iron Systems in 2010 and debuted the WebSphere Cast Iron Cloud Integration offering shortly thereafter. IBM Cast Iron projects can be deployed using a physical appliance (WebSphere DataPower Cast Iron Appliance XH40), a virtual appliance (WebSphere Cast Iron Hypervisor Edition) or a full cloud service (IBM Cast Iron Cloud).

WebSphere DataPower Cast Iron Appliance XH40

Self-contained physical appliance that provides everything needed to connect cloud and on-premise applications.

WebSphere Cast Iron Live

Complete multi-tenant cloud service which consists of both web API services and cloud integration services.

WebSphere Cast Iron Hypervisor Edition

A virtual appliance that can be installed on your existing servers by using virtualization

WebSphere Cast Iron Express

Allows salesforce.com admins to rapidly configure, run and manage integration projects between Salesforce.com and other data sources.



...And one cannot overlook the new face of engagement. Last year there were 850K Android activations per day, over 372M iOS devices sold with 62M alone in Q4, 2011. By 2016, the estimate is that there will be 1 billion+ smart phones in market, 375 million+ tablets.

When it comes to transformational business approaches, enterprise mobile applications or "apps," therefore, represent a great leap forward. Rapid advances in mobile device connectivity and the evolution of new capabilities such as global positioning system (GPS) functionalities and cameras have compounded opportunities and helped drive explosive growth in the mobile marketplace. Businesses everywhere are now strategically employing enterprise mobile apps to support business objectives.

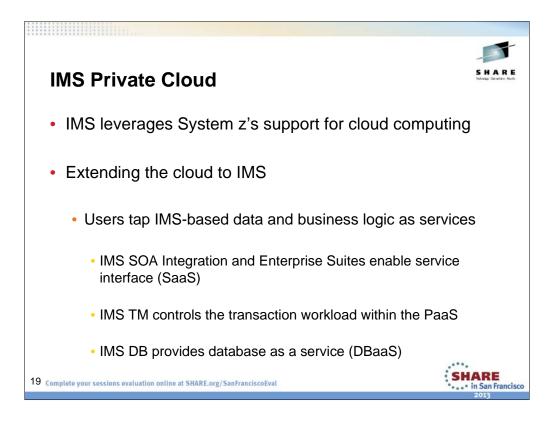
The IBM Worklight mobile application platform provides an open-standards-based approach to enterprise mobile application implementation that dramatically simplifies development for multiple devices and mobile operating systems. Not only does the platform provide tools and a console for easy deployment and maintenance of mobile apps, but it also supports enterprise integration with features that can enable security-rich device-server connectivity and security-rich storage on a device. By leveraging the Worklight platform, your organization can develop apps with enterprise connectivity for smartphones and tablets using one of several enterprise mobile application programming models.

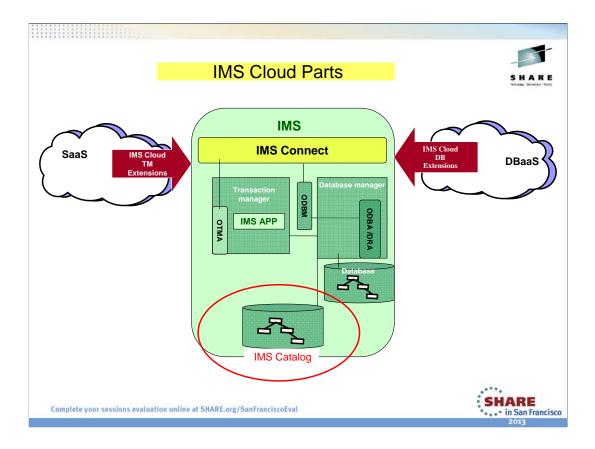


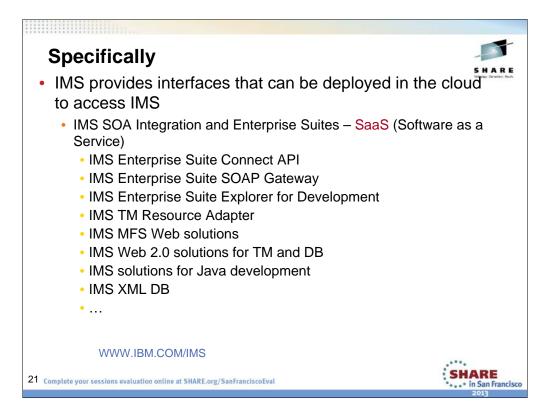
By leveraging the Worklight platform, your organization can develop apps with enterprise connectivity for smartphones and tablets using one of several enterprise mobile application programming models, including the hybrid programming model that supports cross-platform development using familiar web

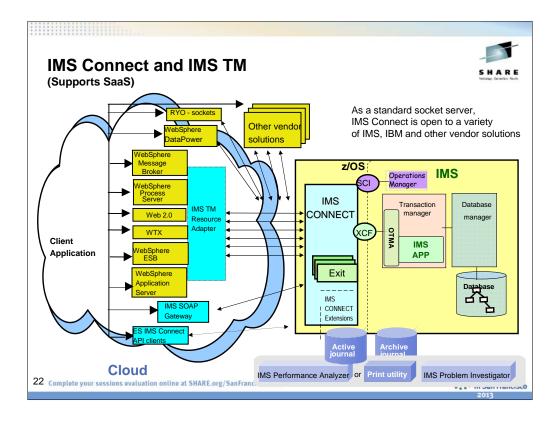
technologies, such as HTML5, Cascading Style Sheets (CSS) and JavaScript. Using the hybrid or web programming model allows your developers to write a mobile application once that is designed to run on a diverse device ecosystem that may include the iOS, Android, BlackBerry or Microsoft Windows operating

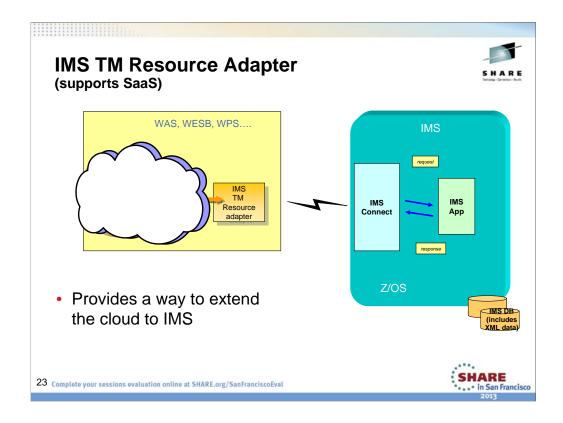
systems. You also have the option to add native customizations to the application to take advantage of the various features on a specific device.











IMS TM Resource Adapter is a J2EE Connector Architecture (J2C) resource adapter that supports IBM Java integrated development environments (IDEs) to create Java applications that can access IMS transactions. The IDEs import MFS, COBOL, PL/I and C source to generate the Java data bindings for the IMS transaction input/output message segments.

This diagram shows the dual functions of IMS TM Resource Adapter: to provide IDE support and to provide runtime environment support to access IMS applications using IMS Connect as the interface to IMS via OTMA. The following IDEs are supported by IMS TM Resource Adapter:

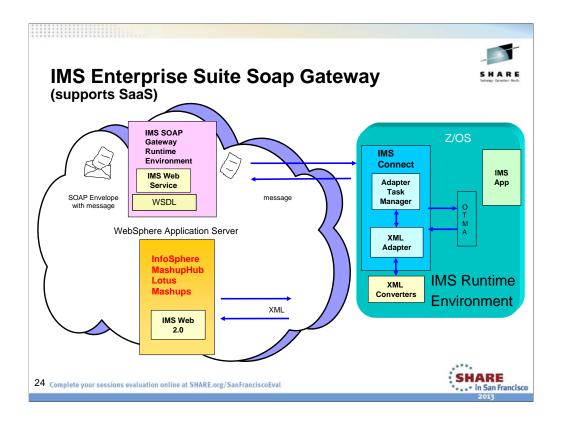
•Rational Application Developer (RAD)

•Rational Software Architect

•WebSphere Integration Developer (WID)

•Rational Developer for System z (RDz)

The IMS TM Resource Adapter also provides a plug-in for a J2EE application server by using the standardized system contracts. IMS TM Resource Adapter provides runtime support for WebSphere Application Server on distributed and z/OS platforms. A J2EE Application Server provides the runtime environment for application code. WebSphere Process Server (WPS) – extends an application server by providing a business process engine that executes business processes. WID is the development tool used to create standards based business processes that are deployed to WPS.



When IMS applications are enabled as web service providers, different types of client applications, such as Microsoft .NET, Java, and third-party applications, can submit SOAP requests into IMS to drive the business logic of the IMS applications.

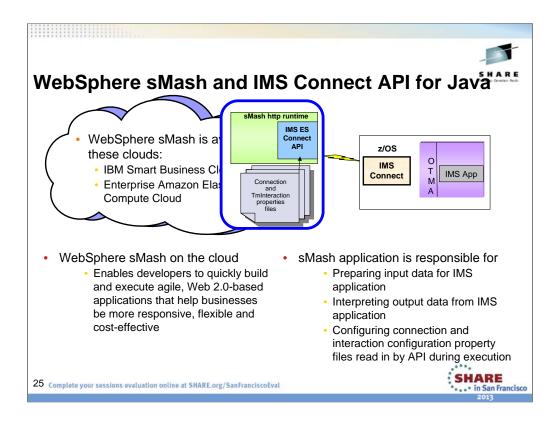
IMS Enterprise Suite Gateway is a web services solution that enables IMS applications to interoperate outside of the IMS environment through the SOAP protocol to provide and request services that are independent of platform, environment, application language, or programming model. It can be deployed on z/OS®, Linux on System z®, and Windows systems. Additionally, the IMS ES SOAP Gateway is compliant with the industry standards for web services, including SOAP/HTTP 1.1 and Web Services Description Language (WSDL) 1.1. It supports web services security (WS-Security), where user identity or security token is authenticated on a per-message basis. Authentication information can also be passed on a per-web service basis, where the user information is defined in the connection bundle for the web service.

•Easily customize IMS transactions without the need to modify the original application. Through a Web interface, you can specify which input parameters to expose to users, as well as the default parameter values to invoke the feed. Feeds can be further restructured and customized by using the operators and functions in InfoSphere MashupHub.

The IMS Web 2.0 solution is embedded in the InfoSphere[™] MashupHub and bundled in the IBM® Mashup Center. The IBM Mashup Center provides a Web interface that can unlock information buried in information silos. The IBM Mashup Center includes two components: InfoSphere MashupHub for creating, storing, transforming and remixing feeds, and for sharing with the Web 2.0 community; and Lotus® Mashups for rapidly assembling a variety of personal, enterprise, and Web content onto a Web page or into an application. IMS Web 2.0 allows you to:

[•]Create IMS feeds from IMS transactions that run on IMS with integrated IMS Connect. Both COBOL and PL/I applications are supported. An IMS feed is the connection information to the host IMS system using an XML-based file format (Atom or RSS).

[•]Utilize the tooling support in Rational Developer for System z Version 7.1.1 to generate the required XML converter driver and correlator. The XML converter driver is used by IMS Connect to transform the data between XML and bytes. The correlator is used by IMS Info 2.0 to map the request and response messages to the input and output data structures of the IMS application. IMS Version 10 customers are entitled to two complimentary, unwarranted copies of Rational Developer for System z Version 7.1.



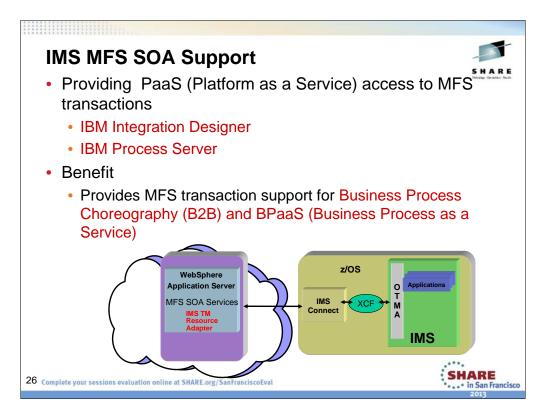
This slide is an example of an application that uses IMS Connect API for Java to access IMS transactions. The Connection and TmInteraction properties files provided with the Connect API for Java installation can be used to specify the connection parameters and the parameters that determine the type of interaction to perform with IMS Connect.

The IMS Web 2.0 Solution for WebSphere sMash allows for a more lightweight and cost-effective solution in generating IMS RESTful services from IMS Transactions for use in Web Oriented Architecture such as Web Mashups. WebSphere sMash is a runtime that allows for developing RESTful services Java (a popular object oriented computer language), PHP (a popular scripting language), and Groovy (a scripting language based on Java) and includes its own http stack for running these services. It is considered lightweight in the sense that it does not need an actual application server to deploy its services.

The IMS Web 2.0 Solution for WebSphere sMash shows how the IMS Connect API for Java (a free offering from IMS) can be used in all three languages (Groovy, Java and PHP) to connect to IMS and launch a transaction. Once the output of an IMS Transaction is returned to WebSphere sMash, the programmer can manipulate that information however they want prior to returning it to the end user. This allows for greater flexibility than the IMS Web 20 Solution for IBM Mashup Center which is all GUI/wizard based.

This Solution is provided as a document only that contains instructions and sample code. The reason for this is that both IMS Connect API and WebSphere sMash are available as separate deliverables. Developers can quickly build and execute Web 2.0-based applications that access IMS transactions

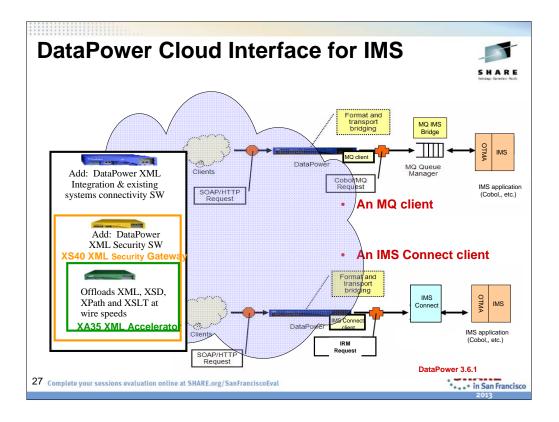
The Solution is available on the IMS Website: http://www-01.ibm.com/software/data/ims/mashup.html

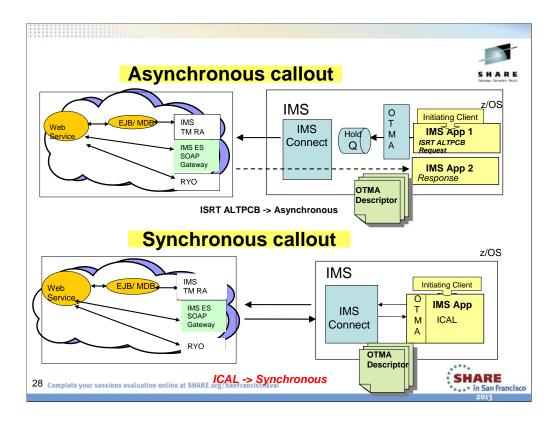


Message Format Service service-oriented architecture (MFS SOA) support enables the reuse of existing MFS-based IMS application program business logic as web services, web pages, or Enterprise JavaBeans (EJB) components, or integrate your MFS applications in business processes.

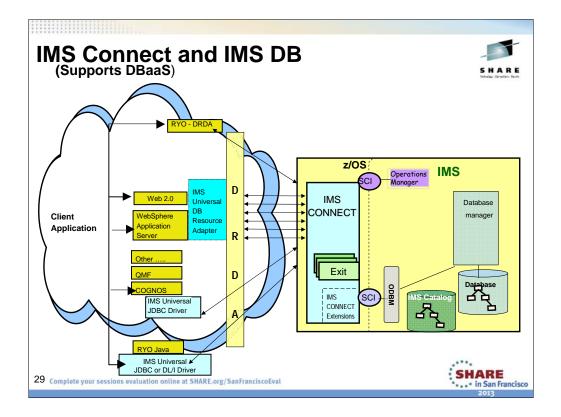
MFS SOA support uses the IMS TM resource adapter on both z/OS or distributed platforms to communicate with IMS Connect and IMS.

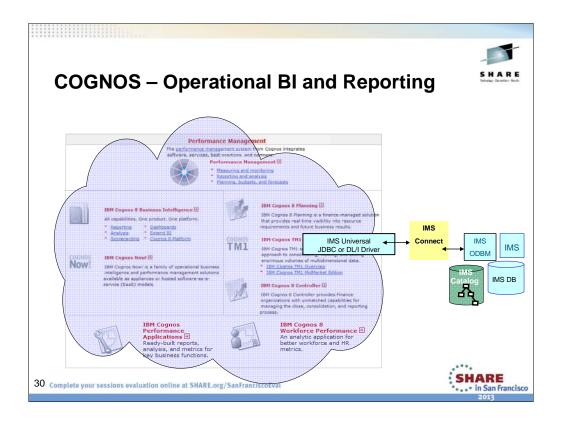
These APARS provides the ability to use the IBM Integration Designer to generate MFS-based services and integrate them into Business Process Execution Language (BPEL) processes that can be deployed to IBM Process Server.

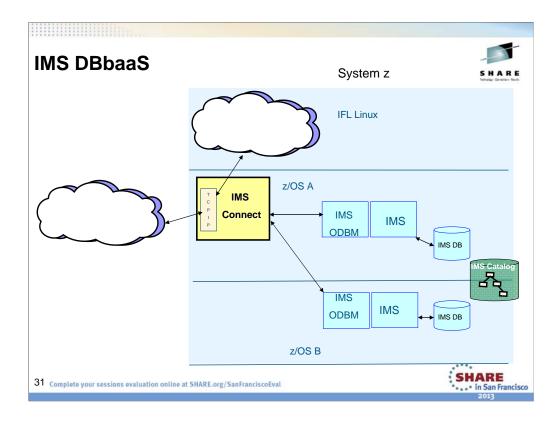


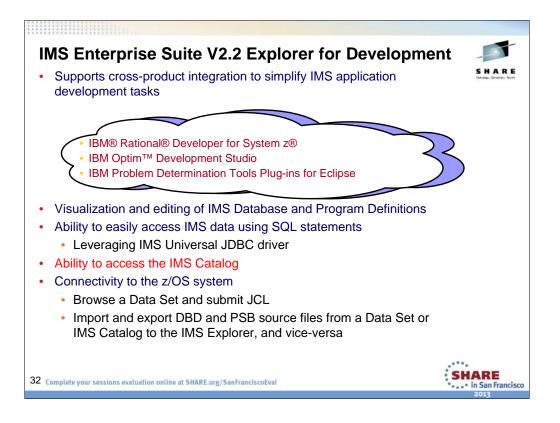


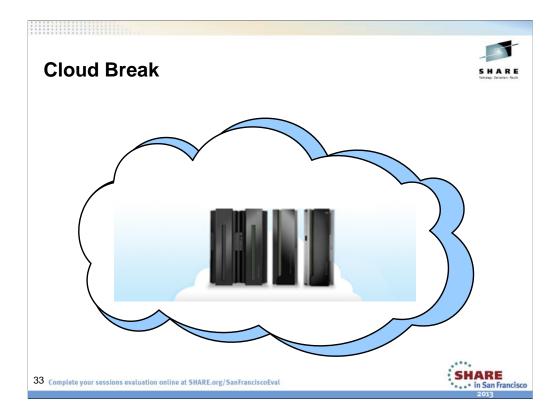
The structure of the library changed quite a lot with IMS Version 10. In Version 11 we will improve the content of information deliverables, but the overall library structure will remain largely the same, with the exception to the noted areas above.

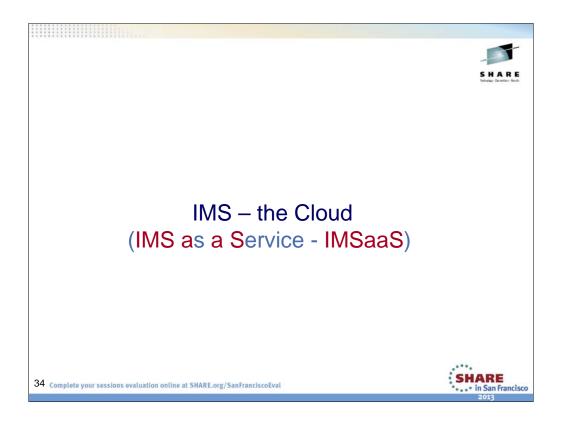


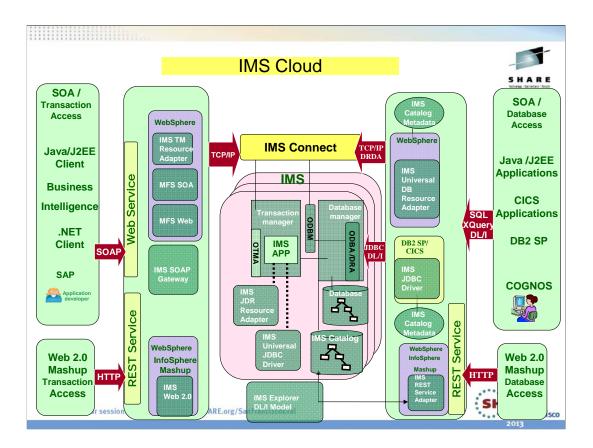


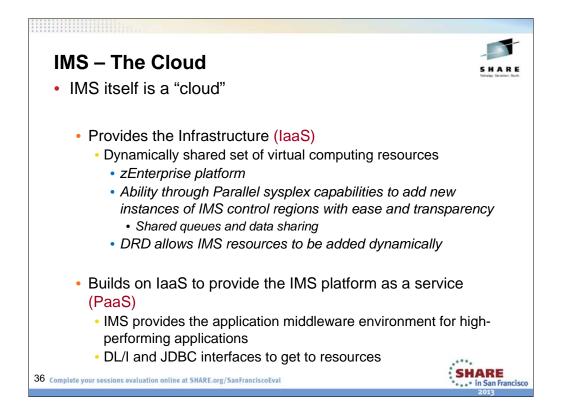


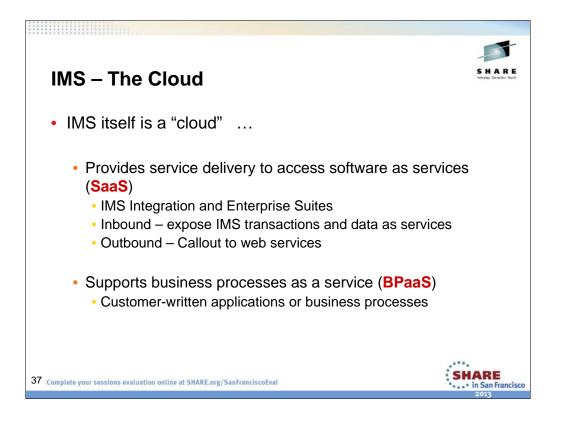


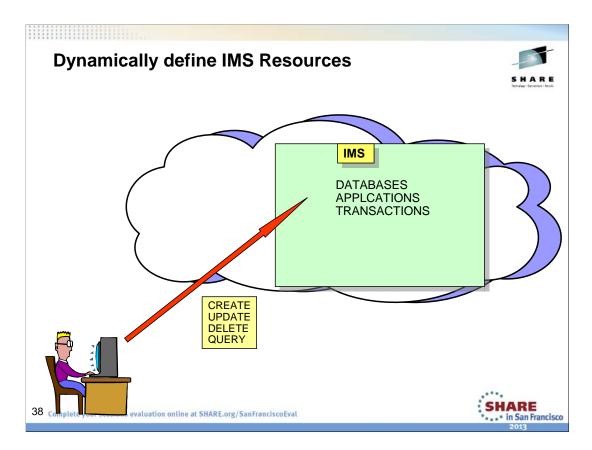




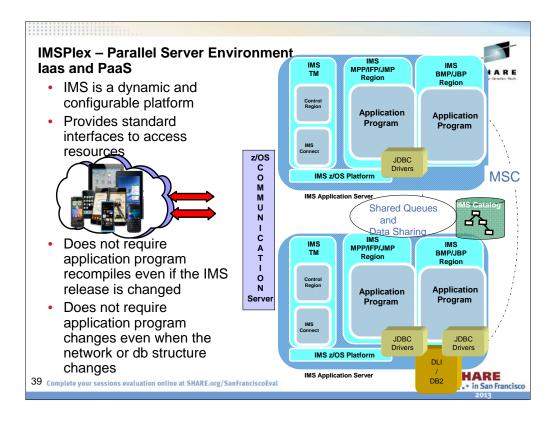


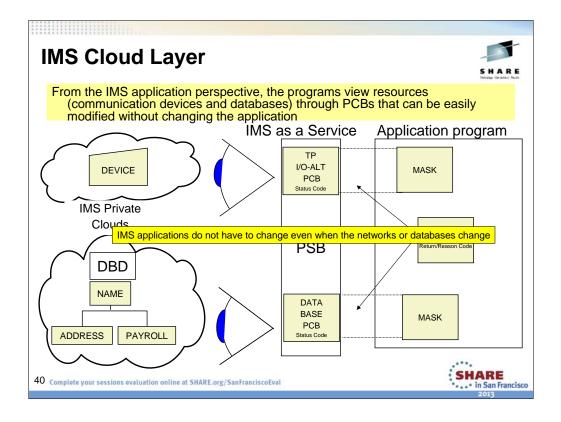


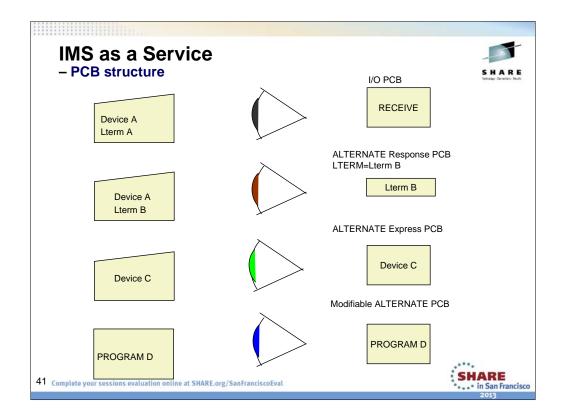




This chart shows the sysgen process being replaced by a user entering commands, and the cold start being replaced by an IMPORT function. EXPORT allows resources defined by the user dynamically to be saved in an RDDS for the next IMS cold start.







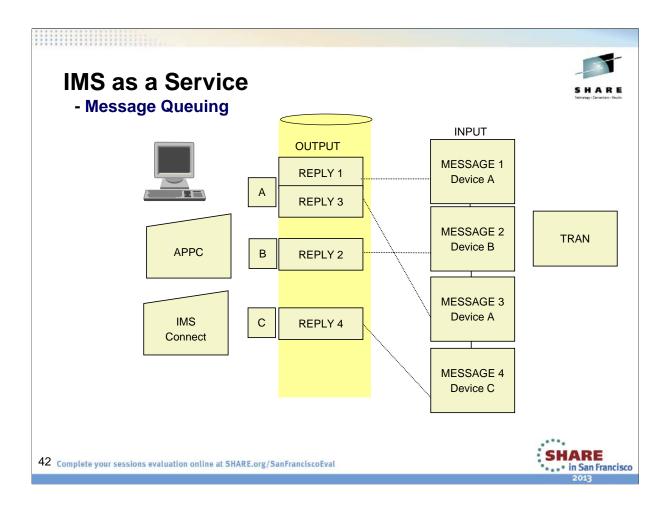
•To send a message to the logical terminal that sent the input message, the program uses an I/O PCB. IMS puts the name of the logical terminal that sent the message in the I/O PCB when the program receives the message. As a result, the program need not do anything to the I/O PCB before sending the Message.

•An *alternate PCB* is a data communication program communication block (ALTPCB) that you define to describe output message destinations other than the terminal that originated the input message.

•An alternate response PCB lets you send messages when exclusive, response, or conversational mode is in effect. You must specify SAMETRM=YES for response alternate PCBs used by conversational programs and programs operating with terminals in response mode.

•An express PCB is an alternate response PCB that allows your program to transmit the message to the destination terminal earlier than when you use a nonexpress PCB. EXPRESS messages can be sent to the destination terminal even though the program abends or issues a ROLL or ROLB call. For an express PCB, the message is available for transmission to the destination when IMS knows it has the complete message. The message is available when a PURG call is made using that PCB, or when the program requests the next input message.

•A modifiable alternate PCB feature allows for the dynamic modification of the destination name associated with this PCB.



IMS as a Service



- Supports multiple runtime environment
- Allows dependent regions to be added as needed for workload

	Application Programs							
		SUPPORTED BY IMS TM/DB CTL			STAND ALONE			
	IMS TM CONTROL REGION <i>(CTL)</i>		MESSAGE REGION	PROCESSING (BMP) BATCH Non-MESSAGE		DB BATCH REGION <i>(DLI)</i> TM BATCH REGION (<i>DB2</i>)		
			(MPP,IFP,JMP)					
								1
	FUNCTIONS			\ \			1	
	•QUEUING	•SCHEDULED BY	IMS	USER			USER	_
	•SCHEDULING	•ONLINE DB'S	YES	YES {	SOME		NO	
	 LOGGING 	•OS/VS FILES	NO	YES	>	RAMS (<u>YES</u>	_
	•I/O	•MSG Q	YES	YES [NTER-	<u>NO</u>	_
	- DATA BASE	•I/O PCB	YES	YES \	CHAN	GEABLE	OPTIONAL	
	- TERMINAL			/			ι	
43 Complete your sessions evaluation online at SHARE.org/SanFranciscoEval								RE an Francisco
2013								

A DB batch program can access full-function databases, DB2 for z/OS databases, GSAM databases, and z/OS files. A DB batch program cannot access DEDBs or MSDBs.

Using DB batch processing

Batch programs are typically longer-running programs than online programs. You use a batch program when you have a large number of database updates to do or a report to print. Because a batch program runs by itself—it does not compete with any other programs for resources like databases—it can run independently of the control region. If you use data sharing, DB batch programs and online programs can access full-function databases concurrently. Batch programs:

• Typically produce a large amount of output, such as reports.

•Are not executed by another program or user. They are usually scheduled at specific time intervals (for example, weekly) and are started with JCL.

•Produce output that is not needed right away. The turnaround time for batch output is not crucial, as it usually is for online programs.

A TM batch program acts like a DB batch program with the following differences:

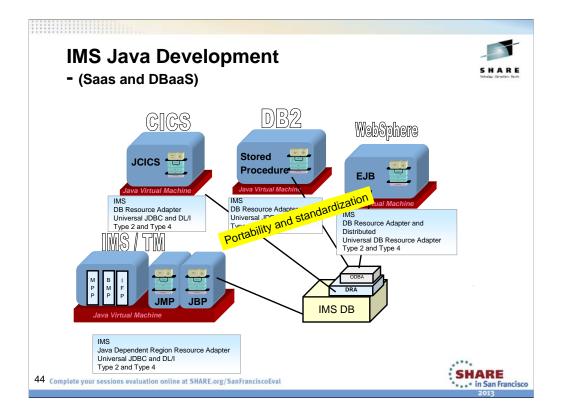
•It cannot access full-function databases, but it can access DB2 for z/OS databases, GSAM databases, and z/OS files.

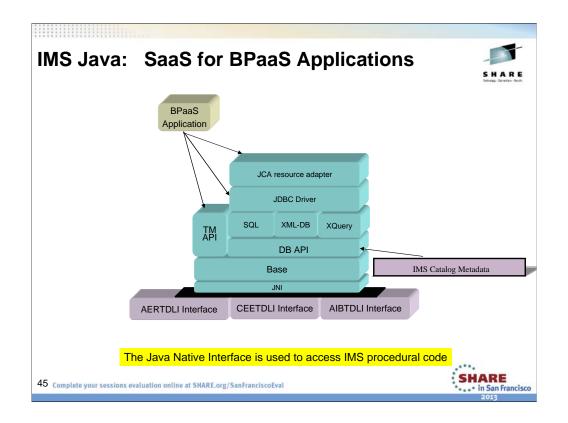
•To issue checkpoints for recovery, you need not specify CMPAT=YES in your program's PSB. (The CMPAT parameter is ignored in TM batch.) The I/O PCB is always the first PCB in the list.

•You cannot dynamically back out a database because IMS does not own the databases.

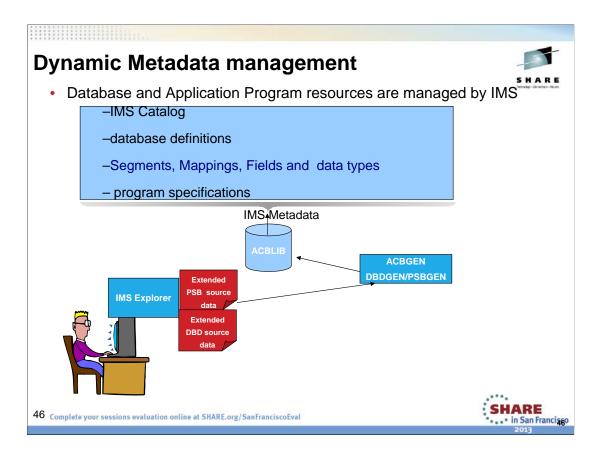
A Message Processing Program (MPP) is an online program that can access full-function databases, DEDBs, MSDBs, and DB2 for z/OS databases. Unlike

BMPs and batch programs, MPPs cannot access GSAM databases. MPPs can only run in DB/DC and DCCTL environments.





The Java Native Interface is used to access IMS procedural code



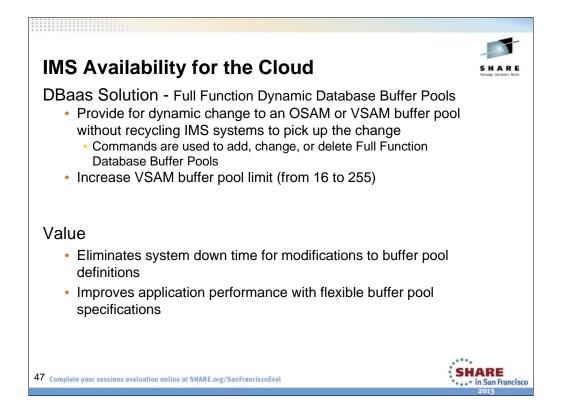
The IMS catalog contains information about IMS program resources, database resources that an IMS system controls, and relevant application metadata that an IMS application controls. This will include all program and database-related information defined to the IMS database system including databases, fields, segments, data types, etc., Changes made to any of these resources such as when you create, alter, or delete any IMS resource information will be reflected in the catalog. The only way to update the IMS catalog is through integrated IMS processes:

- Catalog Populate Utility
- PSBGEN
- DBDGEN
- ACBGEN
- Dynamic Database Change Database Alter database definitions, program specifications

Any changes to COPYBOOK or INCLUDE members will require DBDGEN and ACBGEN processes.

As part of IMS's growth strategy, the IMS catalog is a key component which will play an important role in:

- Simplification
- Integration



The Full Function Dynamic Buffer Pools enhancement allows you to dynamically manage your OSAM and VSAM buffer pools without having to take IMS down. You can now issue online UPDATE POOL commands to add new buffer pools, change or delete them. Imagine you experience an increase in your transaction volume. This could result in an increase in I/Os – however, if you detected this situation and dynamically increased the size of your buffer pools, less I/Os would be needed and processing overhead would be reduced.

Prior to IMS 12, you would have had to recycle your IMS system to pick up this change in buffer pool size. With IMS 12 the overall system availability is improved since you no longer need to take an outage.

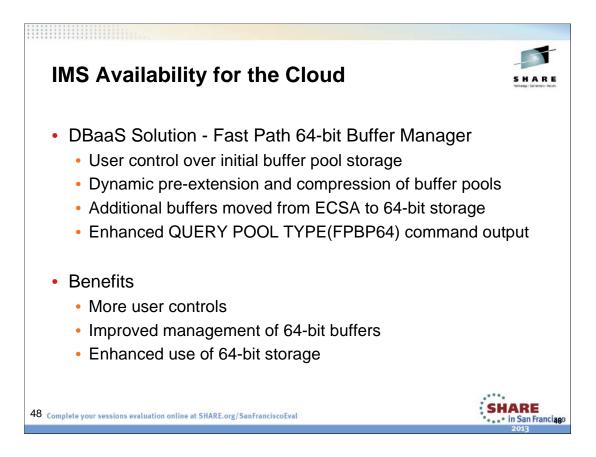
The dynamic changes are retained across an Emergency Restart because they are stored in the Restart Data Set. However, the changes are lost with a subsequent Cold or Warm Start. To make the changes permanent, it is necessary to change the DFSVSMxx proclib member, too. The UPDATE POOL command logs information in the x'22' log record for information purposes only.

The UPDATE POOL command itself is non-recoverable.

Additionally, the total number of possible VSAM shared buffer pool IDs has been increased from 16 to 255 – which provides constraint relief, allowing you to easily grow your applications and databases. More buffer pools also allows you to better segregate multiple databases across the buffer pools. Increase in VSAM shared buffer pools makes it possible for more databases and applications to grow.

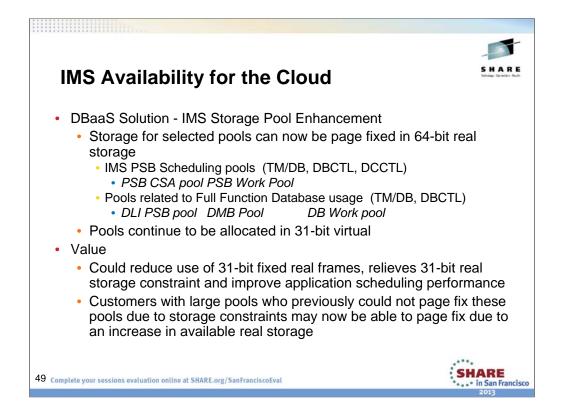
These enhancements improve availability and flexibility by allowing you to respond to your changing business needs right away, without having to take an outage or be constrained to a small amount of shared VSAM buffer pools.

This improves flexibility and system availability. This could eliminate systems downtime for buffer pool management.



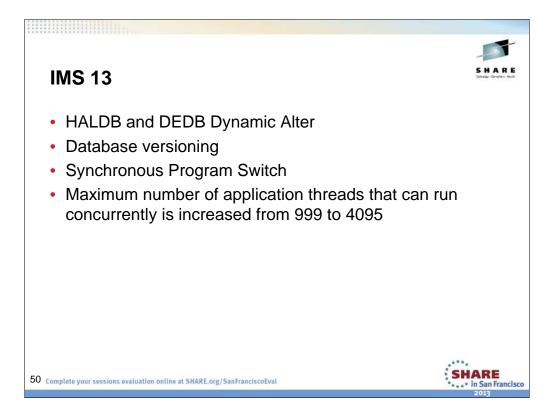
IMS 12 enhances the Fast Path 64-bit buffer manager which was introduced in IMS 11.

IMS 12 allows the user to specify the initial amount of 64-bit storage used for the buffer pool. Buffer pools are preexpanded, that is, expanded in anticipation of future needs. They are compressed when the use of a subpool drops. IMS 12 moves some buffers that were still in ECSA to 64-bit storage. Finally, IMS 12 enhances the QUERY POOL TYPE(FPBP64) command output.



Storage for the selected pools can now be backed by 64-bit real storage (when available), for the DB work pool (DBWP), DMB pool (DLDP), PSB CSA pool (DLMP), DLI PSB pool (DPSB), and the PSB work pool (PSBW). They remain allocated in 31-bit virtual storage. The movement to 64-bit real storage allows some users to fix these pools when they previously could not fix the pools due a shortage of 31-bit real storage. The pools that could see the greatest benefit by this are the CSA PSB pool and the DLI PSB pool since they tend to be the largest pools at most installations. This can provide a performance improvement for application scheduling.

Customers who page fix any of these pools and who have 64-bit real storage on their processor could see a reduction in 31-bit fixed real frames. Customers with large database pools who previously could not page fix them due to 31-bit real storage constraints may now be able to fix them, since the fixed pages will be backed in 64-bit real storage. This could improve the performance related to PSB pool castout processing. Today, if you have large PSBs that are infrequently used, they might get cast out after a period of time. If the pool is not page fixed, the pages for the large PSBs could be paged out. In order to do the cast out MVS would need to page in the storage only for IMS to throw it out. If the pool is page fixed, then the page out/page in operation would not be required for a PSB cast out process. For customers who would like to page fix their PSB pool to improve castout processing, this change would allow them to do that in 64-bit real.

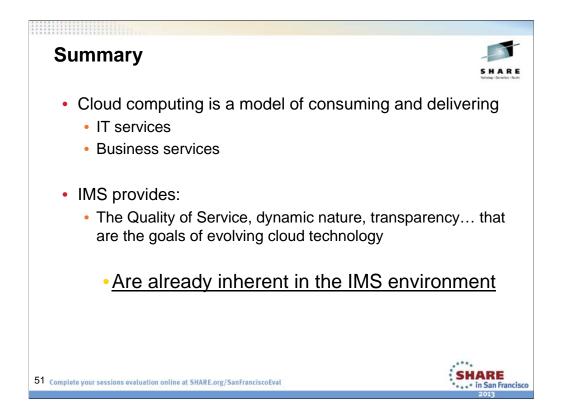


Full-function HALDB alter support enables an IMS HALDB structure to be changed without unload/reload of the database.

Fast path DEDB alter support enables DEDB area changes without unload/reload of the area.

Database versioning support enables programs to use different versions of the same physical database.

The maximum number of application threads that can run concurrently is increased from 999 to 4095. This supports newer system hardware with additional room for vertical growth, and enables IMS to add capacity without having to add more IMS images.



Cloud computing is a model of consuming and delivering IT and business services that enables users to get what they need, as they need it—from high-end, reliable database technology and business applications to IT infrastructure and platform services. Capabilities of self-provisioning and policydriven management provide for faster deployment and faster time to production, and they help reduce IT s aff involvement. Deploying IMS on z/OS in a cloud environment provides significant economies of scale, while allowing a broader audience greater flexibility in exploiting key functions and the wide capabilities of the System z information management portfolio.