

Understand the New Enterprise Cloud System

August 4, 2014

Kershaw Mehta
kershaw@us.ibm.com



Introduction to IBM Enterprise Cloud System

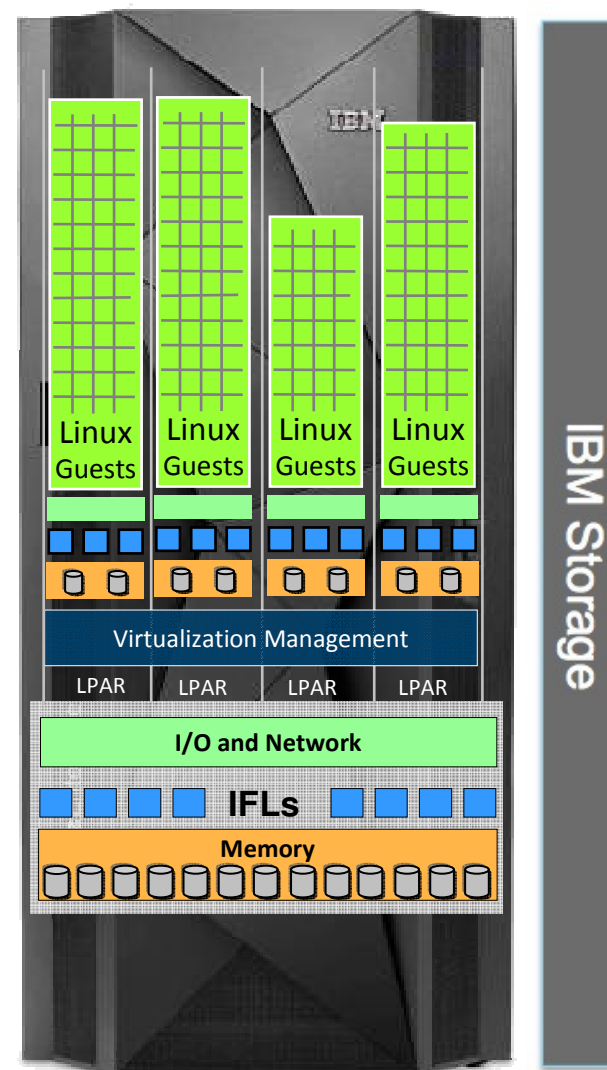
Converged Infrastructure-as-a-Service solution, providing a highly available, secure, cloud platform based on System z technology

- **Pre-configured and integrated system**

- Includes Processor, Disk, Hypervisor, Cloud Management Software and Services
- Pre-installed cloud management software that leverages open source such as OpenStack and Linux to deliver orchestration, provisioning and monitoring
- Integration performed at IBM's Customized Solution Center and onsite by STG Lab-Based Services

- **Flexible configurations**

- No fixed sizes - flexibility on hardware configurations allow customers to choose the right amount of resources for their workload
- Sample configurations will be provided to Sales Team as guidance and comparison



Composition of Enterprise Cloud System

Server



zEnterprise BC12

or



zEnterprise EC12

Storage



Storwize V7000

or



DS8870

Software

- z/VM with following priced features:
 - Directory Maintenance
 - Resource Access Control Facility
 - Performance Toolkit
 - Single System Image
- IBM Wave for z/VM
- Cloud Management Suite for System z
- Operations Manager for z/VM
- Backup and Restore Manager for z/VM
- RHEL and/or SLES Linux for System z

Services

- Integration Services
 - Performed by WW Customized Solutions Center in Poughkeepsie, NY
 - Will integrate server and storage devices and pre-install software prior to shipment to the customer
- On-Site Personalization Services
 - Performed by STG Lab Based Services to complete SW installation and personalize Enterprise Cloud System for the customer

Composition Notes for Enterprise Cloud System

- All components of Enterprise Cloud System are generally available today
 - A customer can put together the same solution by themselves
- The “secret sauce” of Enterprise Cloud System is the **integration**.

Integration means:

- Pre-loaded software – The software is installed and configured at IBM prior to delivery to customer
- Pre-tested Server & Storage Configurations – The server and storage are put together at IBM and pre-tested *together* prior to shipment to ensure no missing parts
- Single Order – The IBM salesperson orders all parts of solution as part of the same order
- Single Simplified Contract – Instead of having separate contracts for server, storage, software and services, a single contract is provided to the customer to sign
- Single Invoice – A single invoice is provided to the customer that covers all four components
- Bottom Line Price – A single bottom line price is provided to the customer for full solution
- Coordinated Planning & Delivery – All components of the solution are in one shipment and arrive at same time at the customer

High Level Overview of Hardware for Enterprise Cloud System

Server

- Choice between zBC12 and zEC12 is based on:
 - Future growth potential
 - Price point
 - Footprint limitations
- zEnterprise server is an IFL only machine
 - No general purpose CPs – all CPs are assigned to be IFLs
 - No zIIPs, no zAAPs, no ICFs



zBC12



zEC12

Storage

- For simplicity, Storwize V7000 will be for SCSI implementations, while DS8870 will be ECKD implementations
 - SCSI disks will be implemented as EDEV (Emulated Devices) under z/VM
- Choice between DS8870 and Storwize V7000 based on:
 - Future growth potential
 - Price point
 - Enterprise Class Qualities of Service
 - Footprint limitations
 - Need for SSI/LGR
- Storage unit will be control unit or base cabinet only
 - Expansion units will be handled outside of this order
 - DS8870 will be directly attached to zEC12 or zBC12
 - V7000 will be attached to SAN switch (customer supplied)



Storwize V7000



DS8870

High Level Overview of Software for Enterprise Cloud System

▪ **z/VM V6.3** with following priced features:

- Directory Maintenance (Dirmaint) feature
- Resource Access Control Facility (RACF) feature
- Performance Toolkit for VM feature
- Single System Image (SSI) feature

Industry leading hypervisor providing virtualization and efficiency at scale and OpenStack enablement

▪ **IBM Wave for z/VM V1.1**

A simple, intuitive graphical tool providing hypervisor management and automation for a z/VM environment, supporting Linux virtual servers.

▪ **Cloud Management Suite for System z V1.1**

which includes:

- OMEGAMON XE on z/VM and Linux
- Tivoli Storage Manager Extended Edition
- SmartCloud Orchestrator

A suite of products providing automated provisioning using OpenStack technology (via SmartCloud Orchestrator), Cloud Monitoring (via OMEGAMON XE on z/VM and Linux) and Cloud Backup and Recovery (via Tivoli Storage Manager Extended Edition)

▪ **Operations Manager for z/VM V1.5**

Supports automated operational monitoring and management of z/VM virtual machines and Linux guests.

▪ **Backup and Restore Manager for z/VM V1.2**

Enables the back up and restore of files and data on z/VM systems and Linux guest images.

Functional Mapping of Software Products

Function	Product	Task Description	User Interface
Hypervisor Management	IBM Wave for z/VM	<ul style="list-style-type: none"> • Discovers z/VM resources and the relationships among them • View and manipulate the state of resources like network setup and storage configuration • Simplify and automate the process of performing specific functions across multiple virtual machines and z/VM systems 	Java-based client running on Windows workstation
Operational Monitoring and Automation	Operations Manager for z/VM	<ul style="list-style-type: none"> • Monitoring of console messages for z/VM infrastructure service machines and Linux guests • Monitoring “state” information for infrastructure service machines and Linux guests • Monitoring spool and page space on the z/VM system 	3270 Interface
Resource Monitoring	OMEGAMON XE on z/VM and Linux - part of Cloud Management Suite for System z	<ul style="list-style-type: none"> • Provides real-time monitoring • Monitoring of individual Linux guest environments • Ability to set service level thresholds and generate events when exceeded • Historical view of monitoring data 	Browser-based

Functional Mapping of Software Products *(cont)*

Function	Product	Task Description	User Interface
Data Management	Backup and Restore Manager	<ul style="list-style-type: none"> • Backup selected files to disk • Backup and restore Linux images • Provides exit points that supports optional data compression and data encryption 	3270 Interface
	Tivoli Storage Manager Extended Edition - part of Cloud Management Suite for System z	<ul style="list-style-type: none"> • Provides file level backup of Linux data with automated scheduling • Provide policy-based backup, archive, recovery and space management capabilities • Allows for multiple versions of files to be archived 	Browser-based CLI for TSM Server
Cloud Management	SmartCloud Orchestrator - part of Cloud Management Suite for System z	<ul style="list-style-type: none"> • Provision / deprovision a virtual system instance • Stop and start a virtual system instance • Provides OpenStack APIs for user extendibility • Provides end-to-end automation and coordination across many data center domains 	Browser-based
Security Management	Resource Access Control Facility (RACF) feature of z/VM	<ul style="list-style-type: none"> • Provides fine-grained access controls of VM resources used by users and guests • Provides auditing capabilities of "VM events" by ensuring security policy enforcement • Separates the disciplines of security administrator, auditor and operations staff 	3270 Interface

Inclusion of Linux with Enterprise Cloud System

- Given that many of the infrastructure service machines from IBM Wave, Tivoli Storage Manager and OMEGAMON XE run in a Linux environment, Enterprise Cloud System will deliver a distribution of Linux on System z from either Red Hat and/or SUSE.
- Moreover, SmartCloud Orchestrator specifically requires a RHEL (Red Hat Enterprise Linux) distribution to operate in both the z/VM environment and on the x86 server.
- An Enterprise Cloud System customer will first need to acquire a license for RHEL and/or SLES (SUSE Linux Enterprise Server) prior to the delivery of the Enterprise Cloud System solution from IBM.
 - Legally speaking, this then allows IBM to deliver Linux for System z with Enterprise Cloud System on behalf of the customer or as an agent of the customer.
 - Also, additional legal license agreements have been acquired from both Red Hat and SUSE to allow IBM to use our distribution of Linux for delivery with Enterprise Cloud System.
- Levels of Linux distributed with Enterprise Cloud System:
 - SLES 11 SP3 plus patches
 - A variety of RHEL 6.x
 - Primarily RHEL 6.5 plus patches
 - RHEL 6.3 (SCO sample master image)
 - RHEL 6.4 (z Region server)

Overall Design Principles for Enterprise Cloud System

1. We have implemented IBM “best practices” in Enterprise Cloud System to help close skills gap and to provide a sound foundation for the customer to build on
2. Since we do not have fixed configurations (no T-shirt sizes) and allow for flexibility in hardware configurations for Enterprise Cloud System, we have designed the software stack to also be flexible so it can operate in a variety of different hardware configurations
3. In addition, the Enterprise Cloud System has been designed to scale when the customer decides to grow their environment in the future
4. Enterprise Cloud System is a disk only solution. Tape drives are not included in the solution. If the customer wants to add tape, this can be done after customer has taken delivery of Enterprise Cloud System.
5. The software stack for Enterprise Cloud System is fixed. No additional software can be added; no software products can be deleted. Additional software can be installed by the customer after delivery of the Enterprise Cloud System solution.
 - Exception – When the solution includes a DS8870, the Single System Image feature of z/VM will be included. When the solution includes a Storwize V7000, the Single System Image Feature of z/VM will not be included.

Overall Design Principles for Enterprise Cloud System (cont)

6. Enterprise Cloud System has been design to fit into a variety of different network implementations
7. Since Enterprise Cloud System can be delivered with either a SCSI or ECKD device, we have maintained similar implementations and only deviated from this principle where it made sense (ie. as a result of exploiting Single System Image functionality)
8. Perform as much customization as possible in the WW Customized Solution Center prior to shipment to the customer, in order to minimize the on-site activity
9. Any customization that is dependent on customer environment (ie. IP addresses) will be performed at the customer site. This activity is called *personalization*
10. We have selected shortwave technology for both FICON/FCP and OSA connections, since we believe this technology will satisfy our customer set.
 - Shortwave has a distance limit of up to 100 meters for copper cables and 300 meters for fiber optic cables
 - If longwave is required for FICON and OSA connects, additional cards can be purchased afterwards (MES)

Overall Design Principles for Enterprise Cloud System *(cont)*

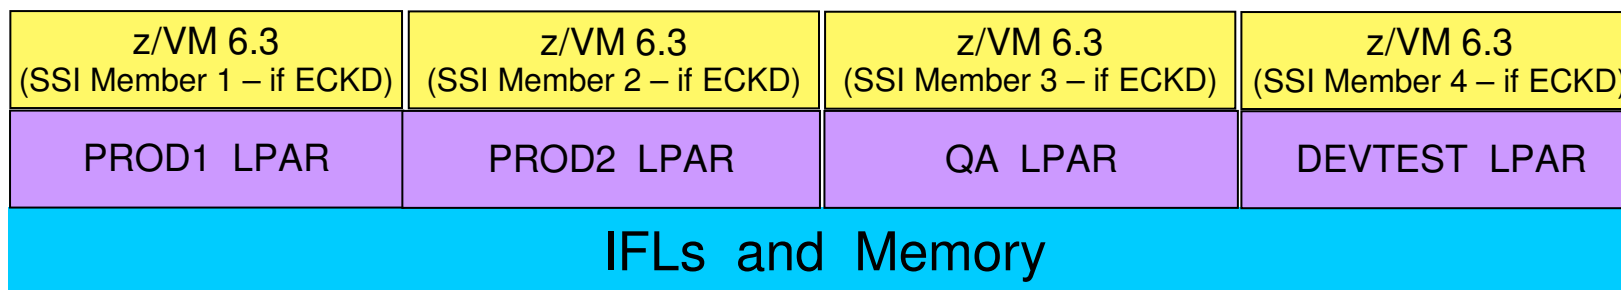
11. All aspects of the configurations were designed in accordance to best practices. As such some minimum hardware requirements have been established, which should be followed when ordering Enterprise Cloud System

12. The I/O configuration has been constructed to accommodate future growth. Each I/O definition file will contain definitions for SCSI devices, ECKD devices, OSA connections, consoles, and CTC connections.

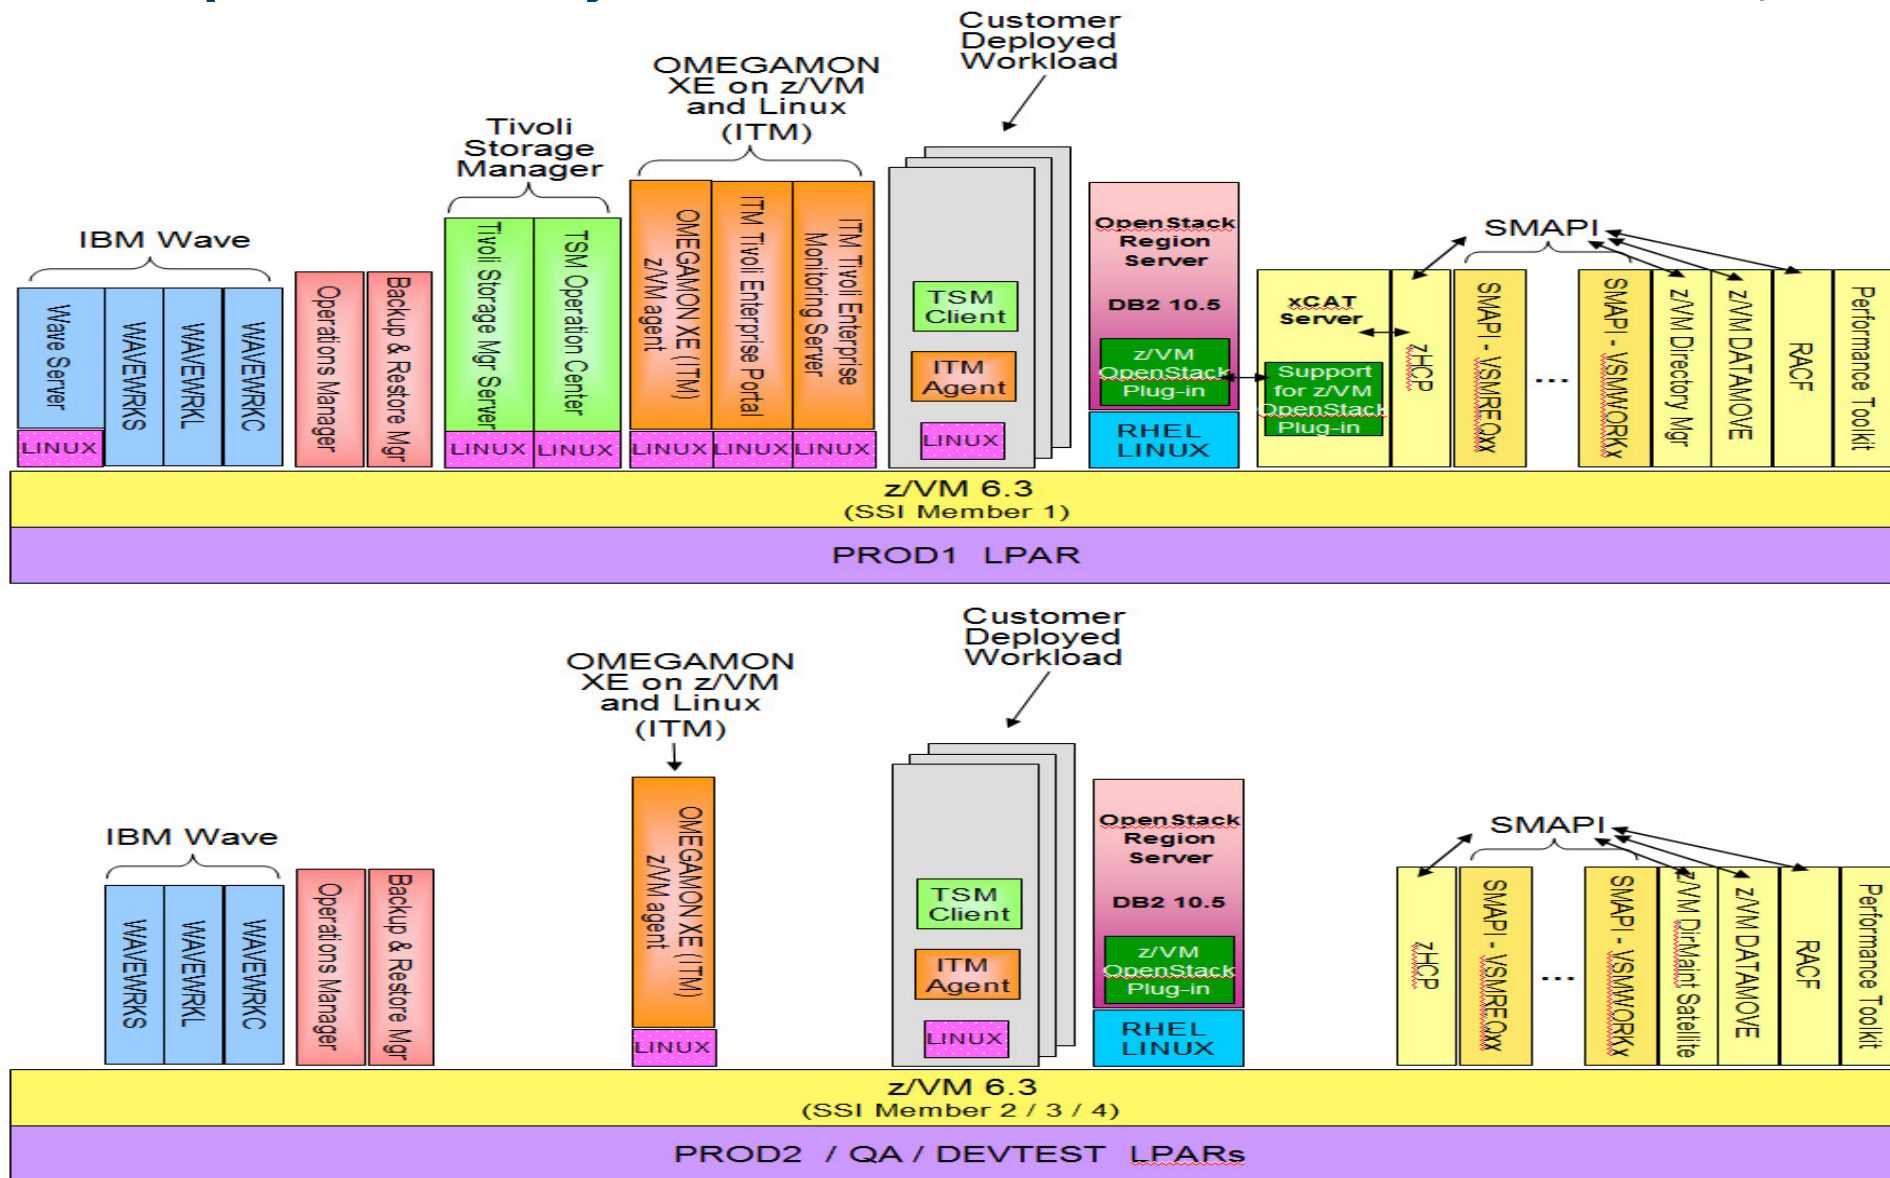
Implementation Architecture

Hardware Virtualization for Enterprise Cloud System

- All IFLs will be shared across all LPARs
 - No dedicated IFLs
- All Memory will be split evenly between LPARs
- Enterprise Cloud System will establish four LPARs
 - Two for Production
 - One for QA
 - One for Dev/Test
- Each LPAR will have its own page and spool volumes
- Customer can choose to add or delete additional LPARs as they see fit.
- Each LPAR will be running z/VM 6.3 and Single System Image will be established when Enterprise Cloud System configuration contains DS8870.



Enterprise Cloud System with Customer Workload - ECKD Setup



Differences Between ECKD and SCSI Setup in Enterprise Cloud System

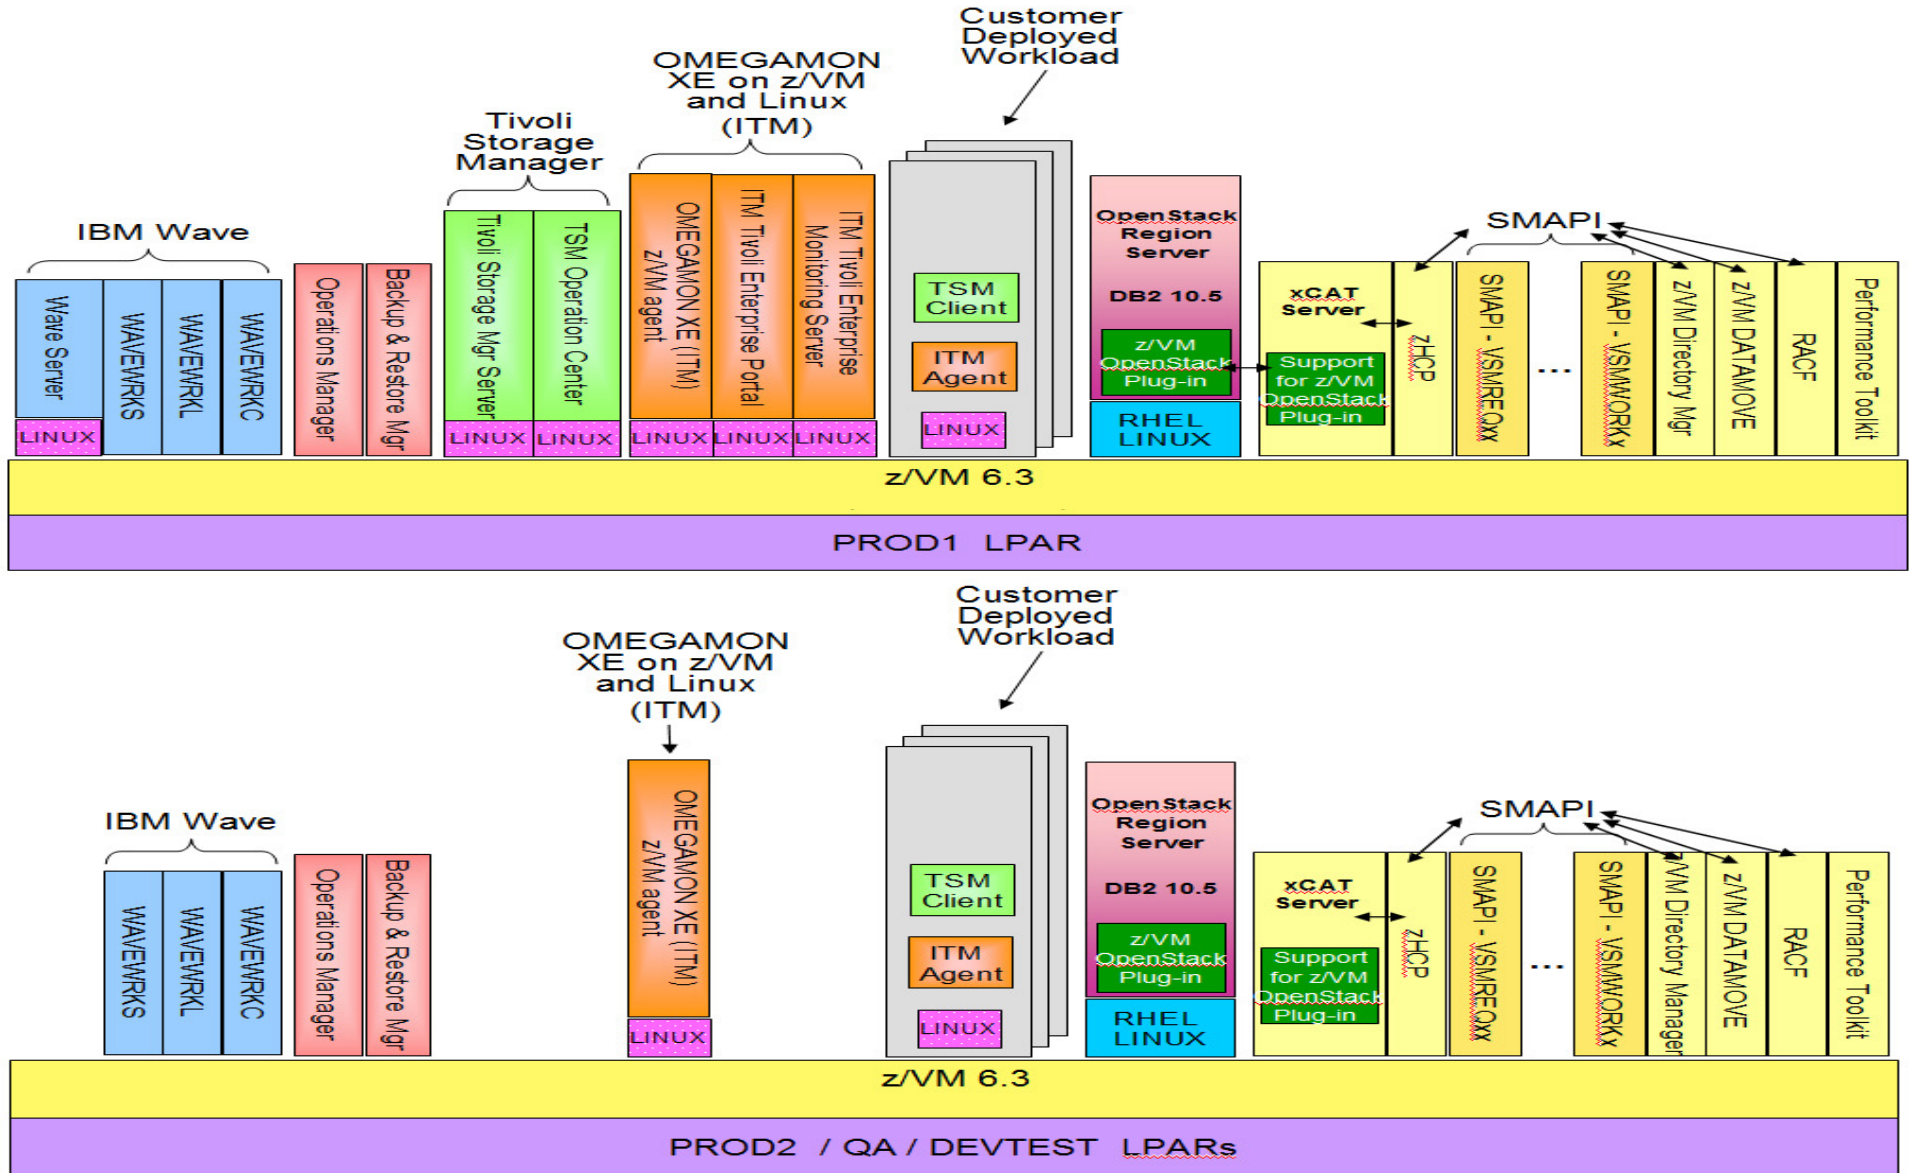
ECKD

1. Single xCAT server running on one LPAR with multiple zHCP servers running on all LPARs
2. RACF database is shared between LPARs
3. User Directory is shared between LPARs
4. Directory Manager (DirMaint) running on one LPAR and DirMaint satellite servers running on others
5. Channel-to-Channel (CTC) connections used for communicating between LPARs

SCSI

1. One xCAT server and one zHCP server running on each LPAR
2. Each LPAR has their own RACF database
3. Each LPAR has their own User Directory
4. Directory Manager (DirMaint) running on all LPARs
5. No Channel-to-Channel (CTC) connections – no communication between LPARs

Enterprise Cloud System with Customer Workload - SCSI Setup



Storage Implementation Details

Specific SCSI-related Configuration Details - Storwize V7000

Attachment

- Since a Storwize V7000 device cannot be directly attached to a zEnterprise server, the Storwize V7000 storage device will need to be attached via two customer provided SAN switches (two for redundancy)
 - SAN switches were not made part of the Enterprise Cloud System infrastructure, because many customers standardize on a specific switch vendor across their enterprise
 - Will require 8 shortwave ports – four in each SAN switch (two incoming, two outgoing) for Enterprise Cloud System deployment

- For Enterprise Cloud System solutions delivered with Storwize V7000, a minimum of four FICON cards (8 ports) will be required
 - Each channel on the FICON Express 8S SX cards will be in FCP mode with NPIV enabled
 - Shortwave only (distance of up to 100 meters for copper cables and 300 meters for fiber optic cables)
 - One port from each card will connect to one of the SAN switches
 - Each SAN switch will then have two connections (total of four connections) to Storwize V7000 device
 - 4 spare FCP ports will remain

Specific SCSI-related Configuration Details *(cont)*

Rack

- Since a Storwize V7000 control enclosure is 2U device, including a full-rack (that is 42U) or 3/4 rack (that is 36U) with Enterprise Cloud System would seem as overkill.
 - Hence, we are requiring the customer to provide space in their own industry-standard rack for the Storwize V7000 device
 - Ordering a rack for the Storwize V7000, must be handled as a separate order outside of the Enterprise Cloud System

Format

- As part of Enterprise Cloud System, RAID 5 (distributed parity) + Spare technique will be employed
- All disks in the Storwize V7000 will be configured as 20GB LUNs – even those for customer use
- Each LUN will be configured in z/VM as an EDEV (emulated FBA device) with four paths

Capacity

- For Enterprise Cloud System, a minimum of 8 drives will be required
- The software configuration in Enterprise Cloud System will utilize 3 TB (assuming 4 LPARS)

Specific SCSI-related Configuration Details *(cont)*

Other

- 100 FCP devices per FCP CHIPID, 18 FCP CHIPID defined in I/O configuration
- No additional features in Storwize V7000 will be configured
 - No thin provisioning, no mirroring, no compression, no Flashcopy – These features can be turned on afterwards by customer

Models Supported

- 2076-124 – Supported for first six orders
- 2076-524 – Available in June 2014.
 - Will be used in Enterprise Cloud System in September 2014.

Drives Supported

- 600 GB 10K rpm 6Gb SFF HDD
- 1.2 TB 10K rpm 6Gb SFF HDD
- 400 GB 12 Gb SSD

Specific ECKD-related Configuration Details – DS8870

Attachment

- The DS8870 will be directly attached to the zEnterprise server – no FICON directors
- For Enterprise Cloud System solutions delivered with DS8870, a minimum of six FICON cards (12 ports) will be required
 - FICON Express 8S SX cards – all configured for EMIF
 - Shortwave only (distance of up to 100 meters for copper cables and 300 meters for fiber optic cables)
 - One port from each of four cards will connect to DS8870 – 4 channels
 - 4 CTC connections will be established (taking up 8 ports)
 - no spare ports

Format

- The disks in the DS8870 storage device will be formatted as ECKD disks, even though SCSI is supported by these devices
- As part of Enterprise Cloud System, RAID 5 (distributed parity) + Spare technique will be employed
- All disks in the DS8870 will be configured as 3390 Model 9 – even those for customer use

Specific ECKD-related Configuration Details *(cont)*

Capacity

- For Enterprise Cloud System, only full drive sets (16 drives) will be orderable
- The software configuration in Enterprise Cloud System will utilize 1.5 TB (assuming 4 LPARs)

Other

- Predefined address range for 4096 ECKD devices in I/O configuration
- No additional features in DS8870 will be configured
 - No PAV (Parallel Access Volumes) or Hyper PAV, no thin provisioning, no mirroring, no compression, no Flashcopy – These features can be turned on afterwards by customer

Models Supported

- 2423-961 – 3 year Service & Support

Drives Supported

- 600 GB 10K rpm 6Gb SAS HDD
- 1.2 TB 10K rpm 6Gb SAS HDD
- 400 GB 12 Gb SAS Flash Drive

Other Implementation Details

Security Implementation for Enterprise Cloud System

- Enable RACF to:
 - Establish password rules
 - Control minidisks
 - Control virtual switches

- Protect the following resources with generic resource classes for:
 - DATASET
 - USER
 - GROUP
 - VMMDISK
 - VMCMD
 - VMLAN
 - FACILITY
 - SURROGAT
 - VXMBR
 - VMXEVENT
 - XFACILIT
 - GXFACILI

- Create generic profiles for:
 - VMMDISK - to control who can link to z/VM minidisks
 - VMLAN - to protect virtual switches
 - SURROGAT

- Activate list-of-groups checking

Security Implementation for Enterprise Cloud System (cont)

- Establish a single RACF group for Linux guests
- Establish a single RACF group for z/VM Administrators
- Other user ID related items:
 - Implement mixed case password
 - Establish the password change interval to be 90 days
 - Don't establish password syntax rules – Customer can establish these themselves later
 - Ensure RACF issues a warning message before a password expires – 30 days
 - Password change interval is 254 days
 - Do not automatically revoke inactive user IDs
 - Passwords for user IDs of service machine should not expire
 - Revoke user IDs after multiple (3) invalid password attempts
- Encryption related items
 - Don't establish SSL encryption for TCP/IP communications – because we don't know customer environment
 - DS8870 – Even though the encryption feature (no extra cost) will be part of a DS8870 order, we will not encrypt the drives during integration
 - To encrypt, we would require the customer to provide their key during integration. Customers are unlikely to provide their key, but customers may want to encrypt their drives afterwards.

Consoles Implementation for Enterprise Cloud System

- Configure one port in each of the two OSA Express 5S 1000BASE-T cards as OSA ICC for each LPAR
- This will then allow these connections to be used for the z/VM console and forgo the need to run the z/VM console on the HMC

Customer Tasks

Items that Customer Needs to Provide

- One to four x86-64 server with 620 GB of available Hard Disk space, 44 GB of available memory and 14 vCPU
 - Must be running KVM as the virtualization technology
 - 64-bit enabled
 - *14 vCPU* translates into a 4 processor server with 6-cores each (or equivalent)
 - Example: An x3650 server with two 12-cores processors and 96 GB of memory (room for future growth) and 3 TB of internal disk (to also account for KVM install)

- One Ethernet switch or Hub to connect HMC to zEnterprise with following characteristics:
 - 16 auto-negotiation ports
 - 10/100/1000 Mbps data rate
 - Full or half duplex operation
 - Auto-MDIX on all parts
 - Port Status LEDs

- Two Ethernet switches with a total of 8 available shortwave (short range) ports that connect to customers Intranet
 - Four RJ-45 modules
 - Four 1GbE optic modules
 - If configuration contains more OSA cards or 10GbE OSA cards, then appropriate number and type of modules are required

Items that Customer Needs to Provide *(cont)*

- Four copper cables to connect OSA 1000BASE-T ports to RJ-45 modules in Ethernet switches
 - EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cables with a maximum length of 100 meters

- Four fiber optic cables (shortwave) to connect OSA 1GbE ports to optic modules in Ethernet switches
 - Multimode fiber optic cables (50 or 62.5 micron) with LC Duplex connector type and maximum length of 300 meters

- Two copper cables to connect Bulk Power Hub to Ethernet switch (for HMC)
 - EIA/TIA Category 5 Unshielded Twisted Pair (UTP) cables with a maximum length of 100 meters

- License for RHEL for all IFLs in Enterprise Cloud System
 - License for RHEL for x86 server also

- License for SLES for all IFLs in Enterprise Cloud System (if applicable)

Items that Customer Needs to Provide *(cont)*

If Storwize V7000 is part of Enterprise Cloud System

- Two SAN switches with a total of 8 available shortwave (short range) ports
 - Four ports in each switch – two ports in, two ports out

- Space for 2U in a 19-inch industry-standard rack – for Storwize V7000 control enclosure

IBM Enterprise Cloud System – integrated solution for Cloud

A simple, highly efficient, server designed for a dedicated Linux environment

- Supports 1000s of Linux virtual servers
- Includes all aspects (HW, SW, and Services) necessary to establish a Cloud environment on System z
- A solution that offers for future growth
- Built on Open Source technology that is becoming defacto standards
- Solution is customized and personalized for each customers needs



Thank You!

