



IDUG

2022 EMEA Db2 Tech Conference

Db2 Warehouse Reference Architecture

John Bell, Les King

jwbell@us.ibm.com, lking@ca.ibm.com

IBM

Agenda

- Introduction
- PDOA, IIAS – Going Forward – Les King
- Cloud Rack for Db2U – John Bell
- Questions

IBM Db2

Portfolio of
database
solutions

***Built to run
the world's
mission
critical
workloads***

Cloud/SaaS

Db2

SaaS

—
Relational database
delivered as a
service



Cloud/SaaS

Db2 Warehouse

SaaS

—
Cloud data
warehouse delivered
as a service



Software

Db2

—
Relational database
built to run the
world's mission
critical workloads



Software

Db2 Warehouse

—
High-performance
data warehouse for
deep analytics and
machine learning



Software

Db2 BigSQL

—
SQL-on-Hadoop
engine, delivering
MPP and advanced
data query



Offering - Definitions

PureData System for Operational Analytics (PDOA) – this was a Power-based, AIX-based, Db2-based analytic appliance targeting ODS and Operational Analytic workloads.

IBM Integrated Analytics System (IIAS) – this was a Power 8-based, Power Linux-based, Db2-based analytic appliance which was positioned as: a follow-on to PDOA; a consolidated platform for both PDA (Netezza) and PDOA; used as a base for IDAA V7

NOTE: When NPS became available we focused on PDOA follow-on only

IBM Db2 Analytics Accelerator (IDAA) – this is the Db2 for z/OS variant of IIAS used for off-loading Db2 for z/OS analytic workloads and accelerating those workloads using the Db2 engine. The actual appliance is identical to IIAS.

End of Support

PDOA V1.0 – Already EOS with a few clients on service extensions.

PDOA V1.1 – Announced EOS in September 2023

IIAS M4001 – Announced EOS in September 2023

IIAS M4002 – Planned EOS in September 2026++

Next Generation

1. On-Premises

1. **X86 Reference Architecture** - customers can deploy a Db2 on CP4D reference architecture.
2. **Power Reference Architecture** – customers can deploy a Power-based reference architecture.

2. Cloud

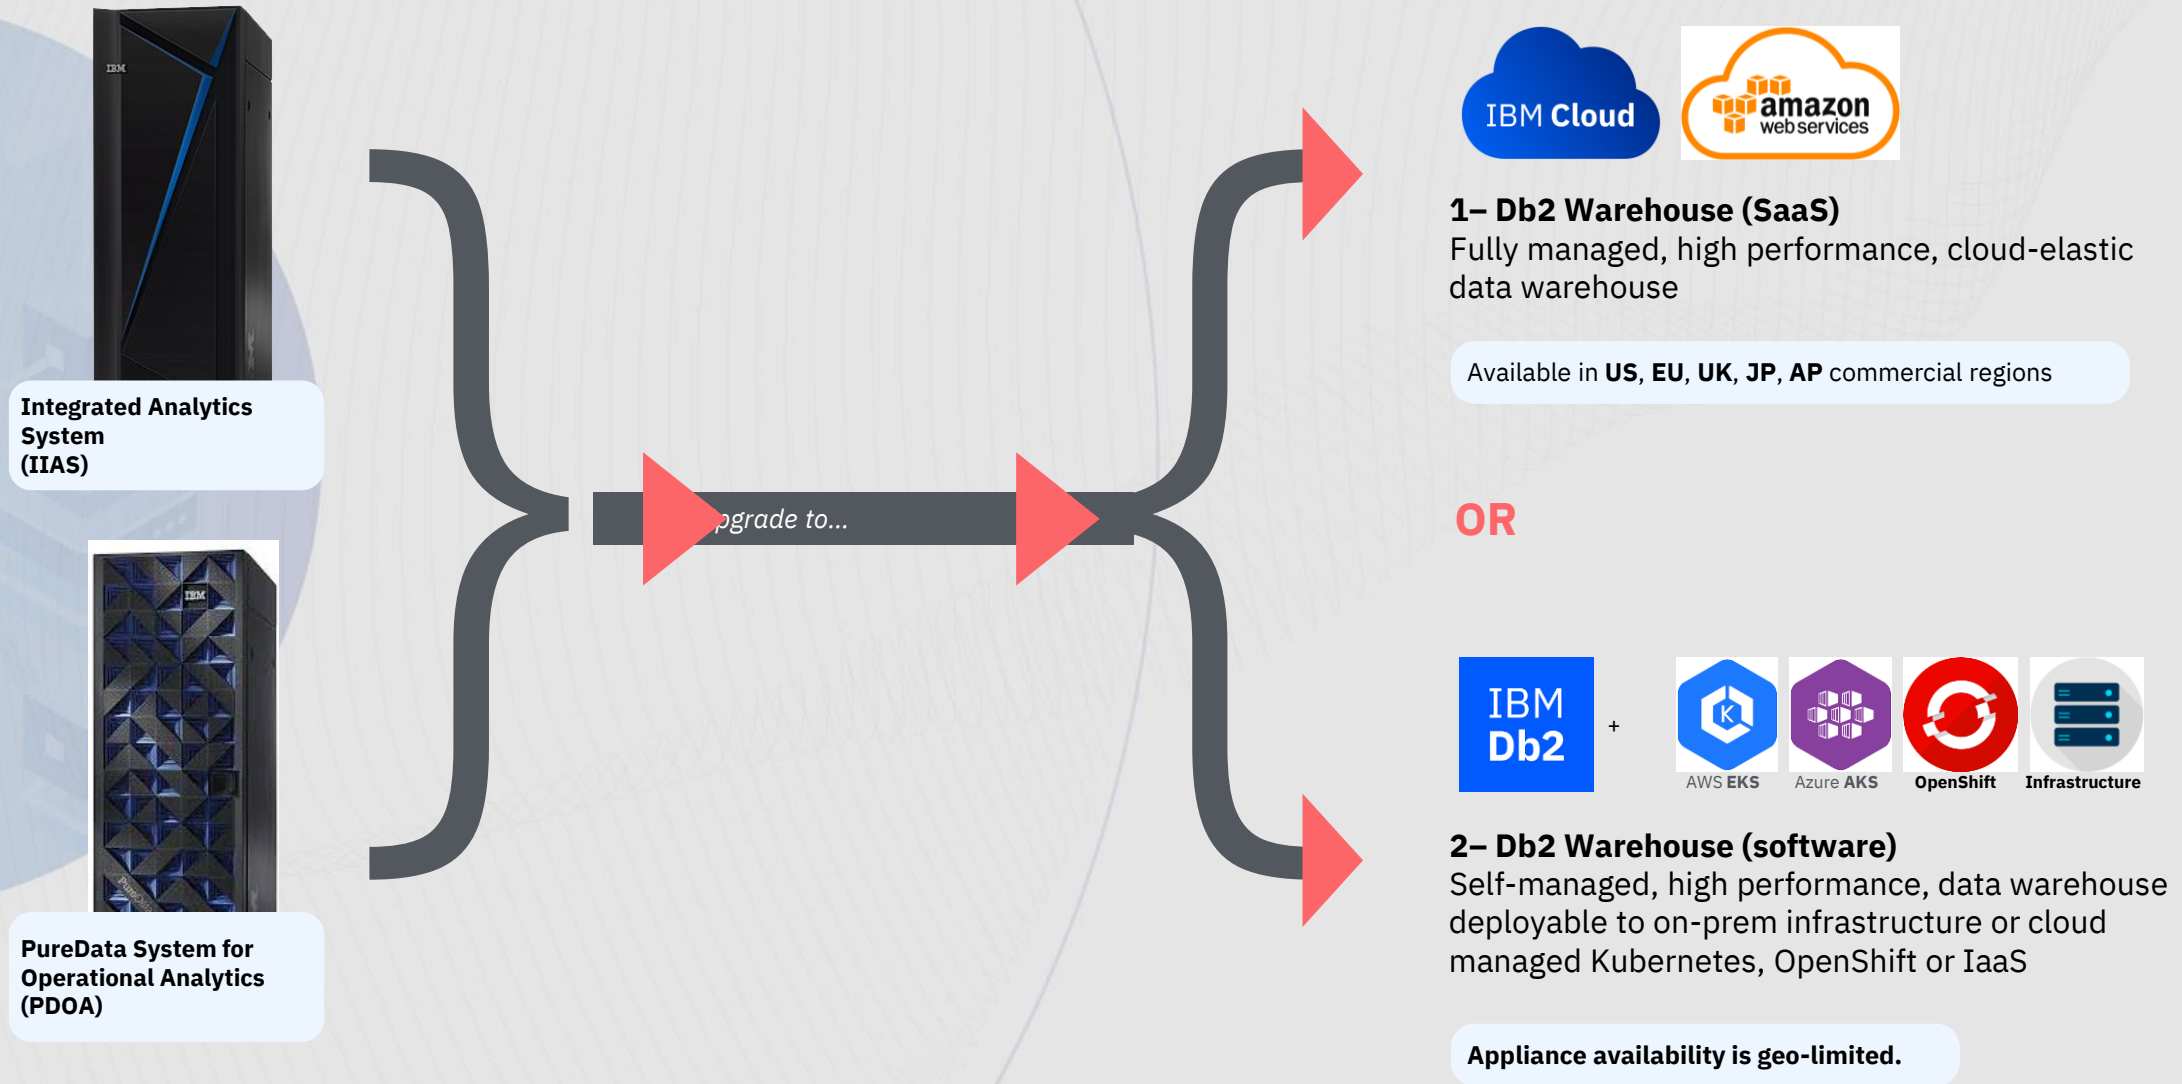
1. **Self-Managed Db2 Reference Architecture** – customers can deploy the x86 reference architecture on any of the key cloud providers (AWS, Azure, GCP)
2. **Fully-managed Db2 Warehouse on Cloud** – customers can deploy a SaaS solution on IBM or AWS

All fully compatible with PDOA and IIAS

Summary of Solution Options

Comparison Point	Db2 Warehouse Reference Architecture on X86	Db2 Warehouse Reference Architecture on Power	Cloud Self-Managed Best Practices	Cloud Fully-Managed SaaS
Hardware Platform	X86 – Lenovo	Power 10	X86 – Follow RA Best Practices	Db2 Warehouse on Cloud
Operating Platform	X86 Linux	Power Linux LE	X86 Linux	X86 Linux
Db2 Deployment Platform	Db2U	Db2U	Db2 AE or Db2U	Db2 CC / Db2 U (Gen 3)
Containerization	RHOS	RHOS	AKS, EKS, ARO, Rancher, ...	CC / RHOS (Gen 3)
Cloud Pak for Data	YES	NO – Can be a data source	OPTIONAL	NO – Can be a data source
Functional Compatibility	YES	YES	YES	NO / YES (Gen 3)
Cloud Provider Options	On-Premises	On-Premises	AWS, Azure, GCP, IBM	AWS, IBM
Support for PB Warehouse	YES	YES	YES	NO / YES (Gen 3)
Hardware Included	NO	NO	NO	YES
Support for Db2 Vnext	YES	YES	YES	YES
Performance Compatibility	YES	YES	YES	YES
Deploy on Cloud Pak System	NO	NO	NO	NO
Availability Target	4Q 2022	4Q 2022	1Q 2023+	Today / 1Q 2023 (Gen 3)

PDOA / IIAS Upgrade Options



1 - Db2 Warehouse SaaS



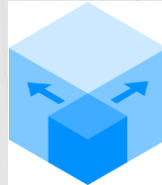
Fully managed / SaaS

Focus on the analytics, we'll take care of the rest



Blazing-fast Db2

Columnar-organized, memory-optimized data warehouse



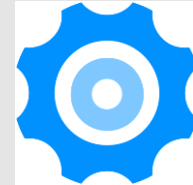
Scalable & elastic

Independently scale and manage compute & storage



Continuously available

Managed compute, highly available storage, cross-cloud replication



Reliable

Double protection with disaster recovery & self-service backup/restore

1 - Db2 Warehouse SaaS

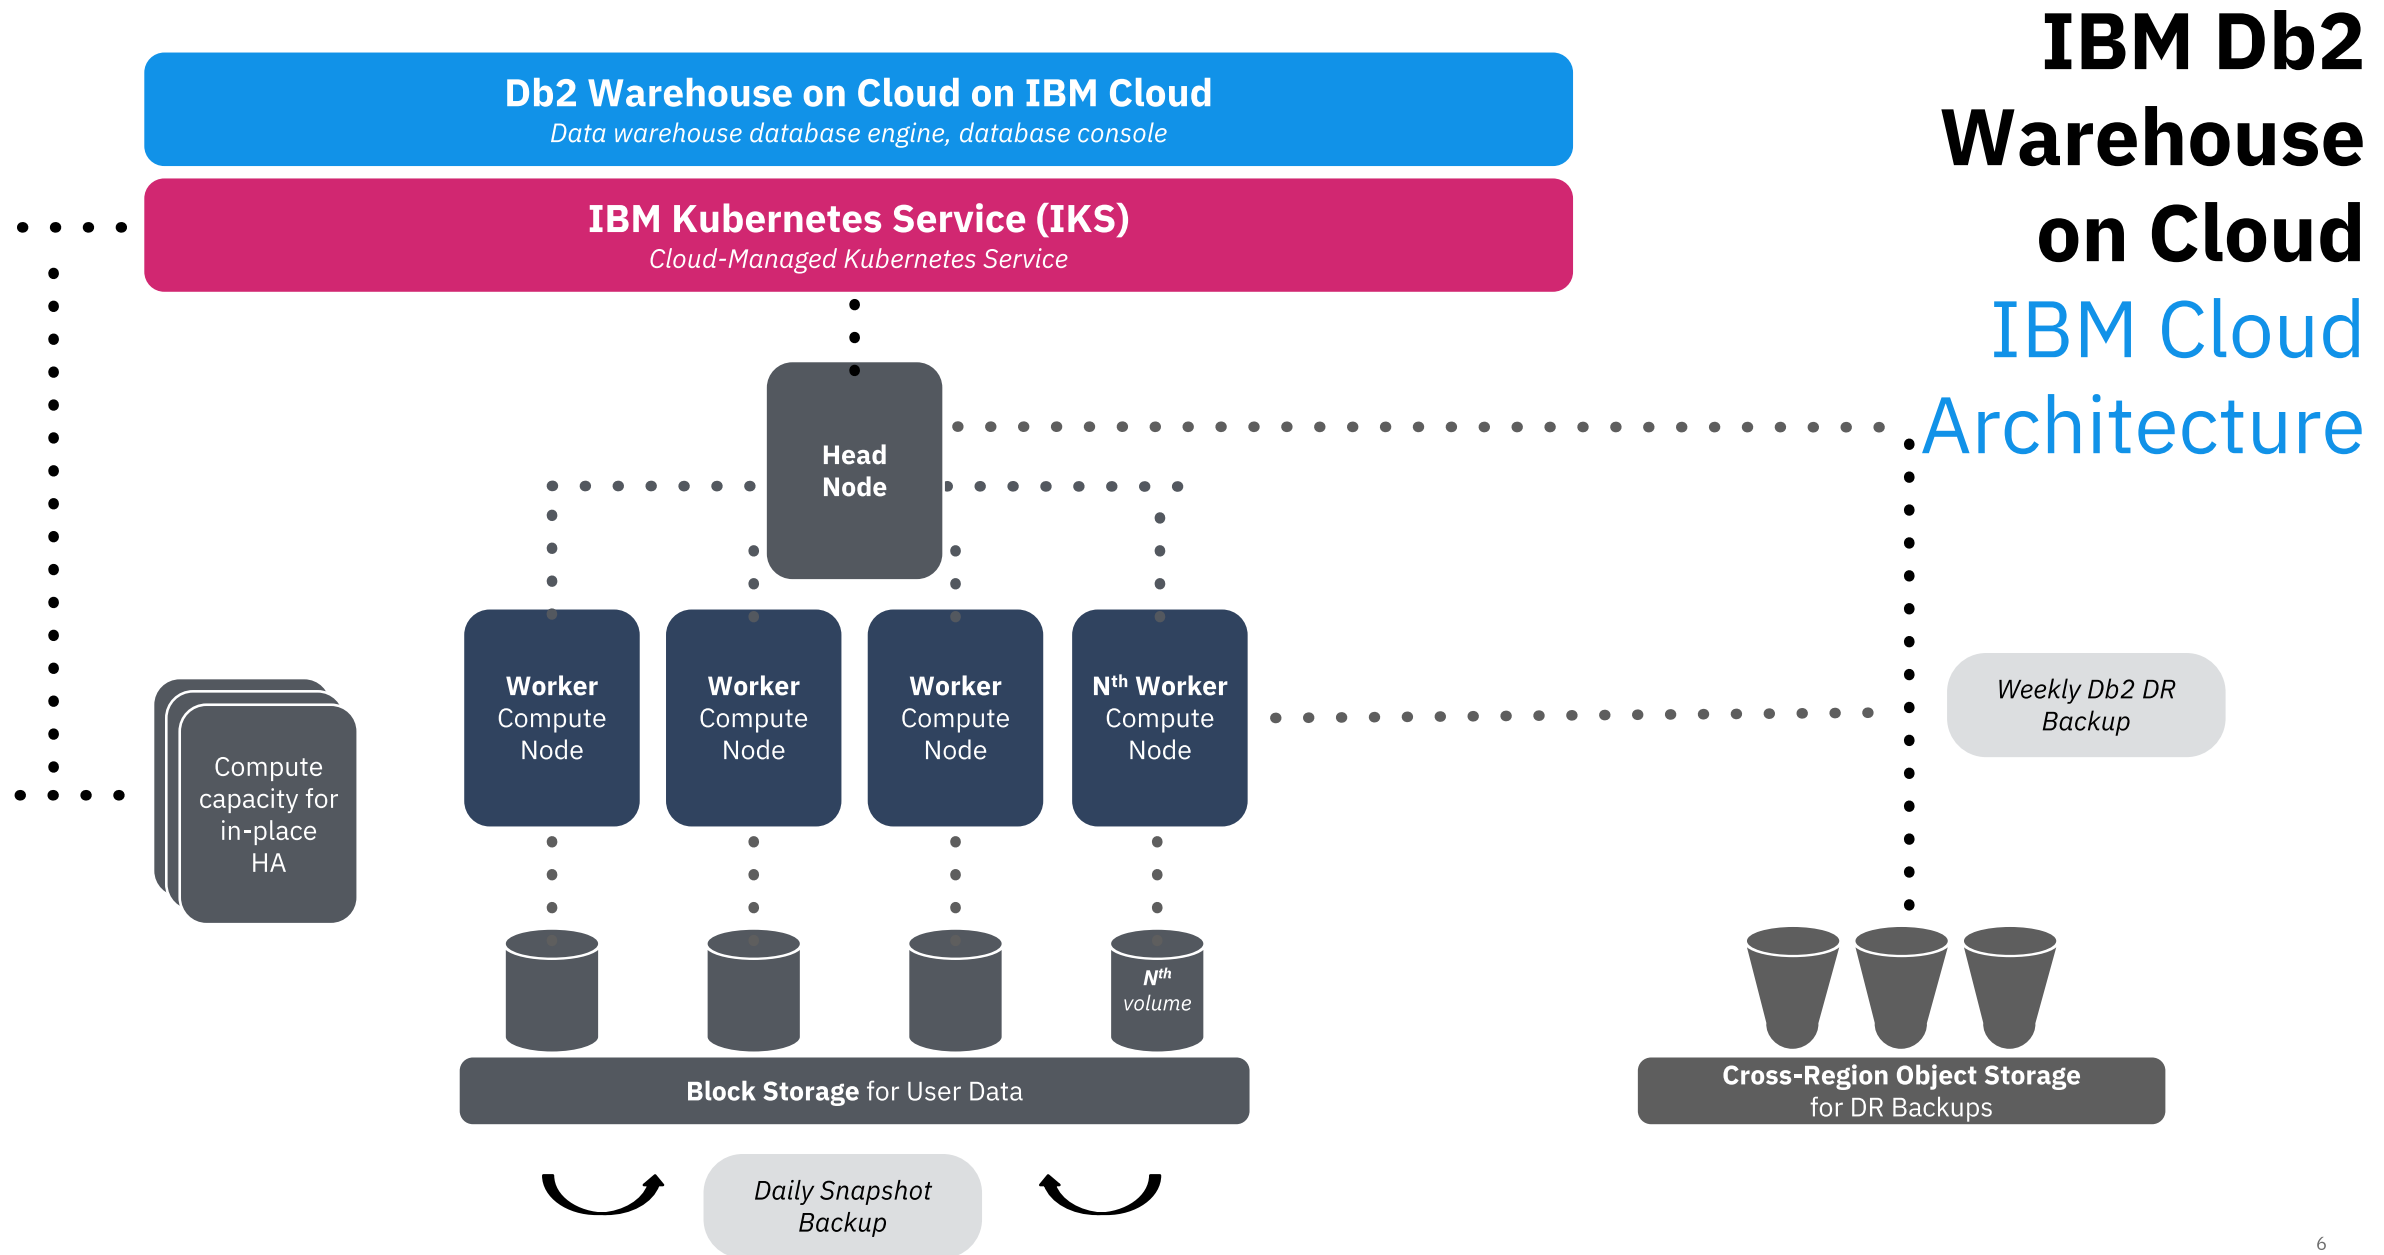
***Choose from two
cloud deployment options***



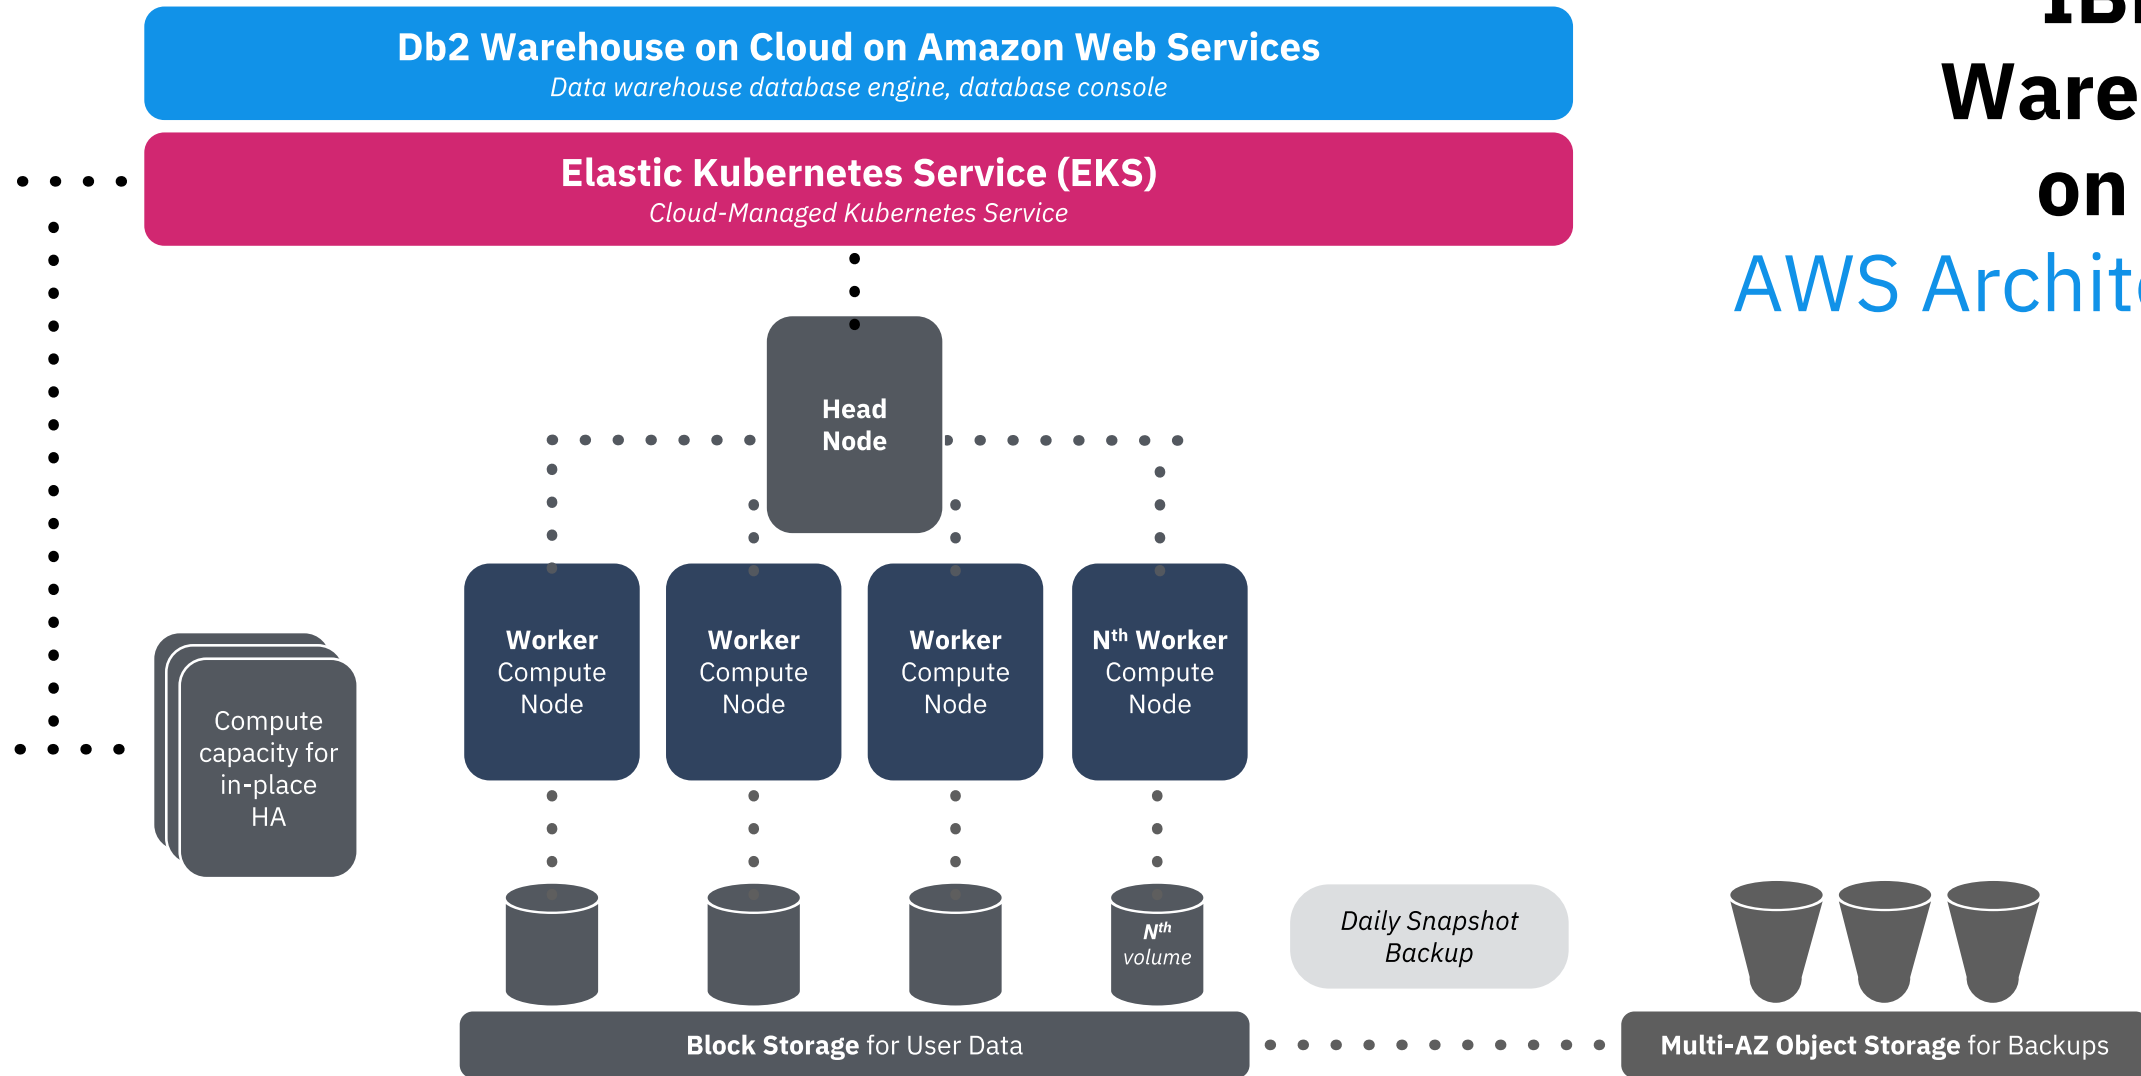
**IBM
Cloud**



1 - Db2 Warehouse SaaS



1 - Db2 Warehouse SaaS



IBM Db2 Warehouse on Cloud AWS Architecture

1 - Db2 Warehouse on Cloud (SaaS) Gen 3

- Historical monitoring support in Console
- AWS cross regional and zonal restore from snapshots
 - Recovery across zones today is done by Ops
 - Moving forward, customers can restore snapshots from a source system to a target system – different zone or region
- Granular/schema level backup and restore
- New configurations for IIAS customers on AWS
 - Alignment with Db2U containers and RHOS platform
 - Support for granular backup options
 - Support for multi-tiered storage
 - Native Cloud Object Store support
 - >1PB support

2 - Db2 Warehouse Software

2

X86 Reference Architecture for on-premises deployments. Equal or better performance than IIAS. Prescriptive HW and Db2 Configuration. Customer purchases SW and HW separately. Leverages **Cloud Pak for Data** and **Red Hat OpenShift**.

Check out
IBM Cloud
Pak for Data



1

Self-Managed Best Practices for cloud deployments. Equal or better performance than IIAS. Prescriptive HW and Db2 Configuration. HW provisioning through cloud provider. **Best cloud alternative** where we don't have SaaS offering.

3

Power Reference Architecture for on-premises deployments. Equal or better performance than IIAS. Prescriptive HW and Db2 Configuration. Customer purchases SW and HW separately. **Appliance-like**. Most similar to existing PDOA and IIAS appliances.



PDOA
IIAS



All Options are

- Fully Functional
- Leverage Db2 Common Engine
- Improved Performance
- Prescriptive (appliance-like)
- Include Services (appliance-like)

2 - Db2 Warehouse Software

2

X86 Reference Architecture for on-premises deployments. Equal or better performance than IIAS. Prescriptive HW and Db2 Configuration. Data & AI sellers sell software. Customer purchases HW separately. Leverages **Cloud Pak for Data** and **Red Hat OpenShift**.

Check out
IBM Cloud
Pak for Data



1

Self-Managed Reference Architecture for cloud deployments. Equal or better performance than IIAS. Prescriptive HW and Db2 Configuration. Data & AI sellers sell software. HW provisioning through cloud provider. **Best cloud alternative** where we don't have SaaS offering.



3

Power Reference Architecture for on-premises deployments. Equal or better performance than IIAS. Prescriptive HW and Db2 Configuration. Data & AI sellers sell software. Hardware sellers sell HW. **Appliance-like**. Most similar to existing PDOA and IIAS appliances.



PDOA
IIAS



All Options are

- Fully Functional
- Leverage Db2 Common Engine
- Improved Performance
- Prescriptive (appliance-like)
- Include Services (appliance-like)

2 - Db2 Warehouse Software

Comparison Point	Cloud Self-Managed RA	On-Premises X86 Ref Arch	On-Premises Power Ref Arch
Hardware Platform	X86	X86	Power 10
Operating Platform	X86 Linux	X86 Linux	Power Linux LE
Db2 Deployment Platform	Db2 AE (BYOL) or Db2U	Db2U	Db2U
Containerization	AKS, EKS, GKE, ARO, ROSA, Rancher	RHOS	RHOS
Cloud Pak for Data	OPTIONAL	OPTIONAL	OPTIONAL
Functional Compatibility	YES	YES	YES
Cloud Provider Options	AWS, Azure, GCP, IBM	N/A	N/A
Support for PB-Size Warehouse	YES	YES	YES
Hardware Included	SOLD SEPARATELY	SOLD SEPARATELY	SOLD SEPARATELY
Support for Db2 Vnext	YES	YES	YES
Performance Improvements	YES	YES	YES
Deploy on Cloud Pak System	NO	NO	NO
Availability Target	4Q 2022+	4Q 2022	4Q 2022
Appliance-Like	NO	PARTIALLY	YES

2 - Db2 Warehouse Software

Comparison Point	Cloud Self-Managed RA	On-Premises X86 Ref Arch	On-Premises Power Ref Arch
Hardware Platform	X86	X86	Power 10
Operating Platform	X86 Linux	X86 Linux	Power Linux LE
Db2 Deployment Platform	Db2 AE (BYOL) or Db2U	Db2U	Db2U
Containerization	AKS, EKS, GKE, ARO, ROSA, Rancher	RHOS	RHOS
Cloud Pak for Data	OPTIONAL	OPTIONAL	OPTIONAL
Functional Compatibility	YES	YES	YES
Cloud Provider Options	AWS, Azure, GCP, IBM	N/A	N/A
Support for PB-Size Warehouse	YES	YES	YES
Hardware Included	SOLD SEPARATELY	SOLD SEPARATELY - LENOVO	SOLD SEPARATELY – INTEGRATED SYSTEM
Support for Db2 Vnext	YES	YES	YES
Performance Improvements	YES	YES	YES
Deploy on Cloud Pak System	NO	NO	NO
Availability Target	4Q 2022+	4Q 2022	4Q 2022
Appliance-Like	NO	PARTIALLY	YES

2 - Db2 Warehouse Reference Architecture on X86

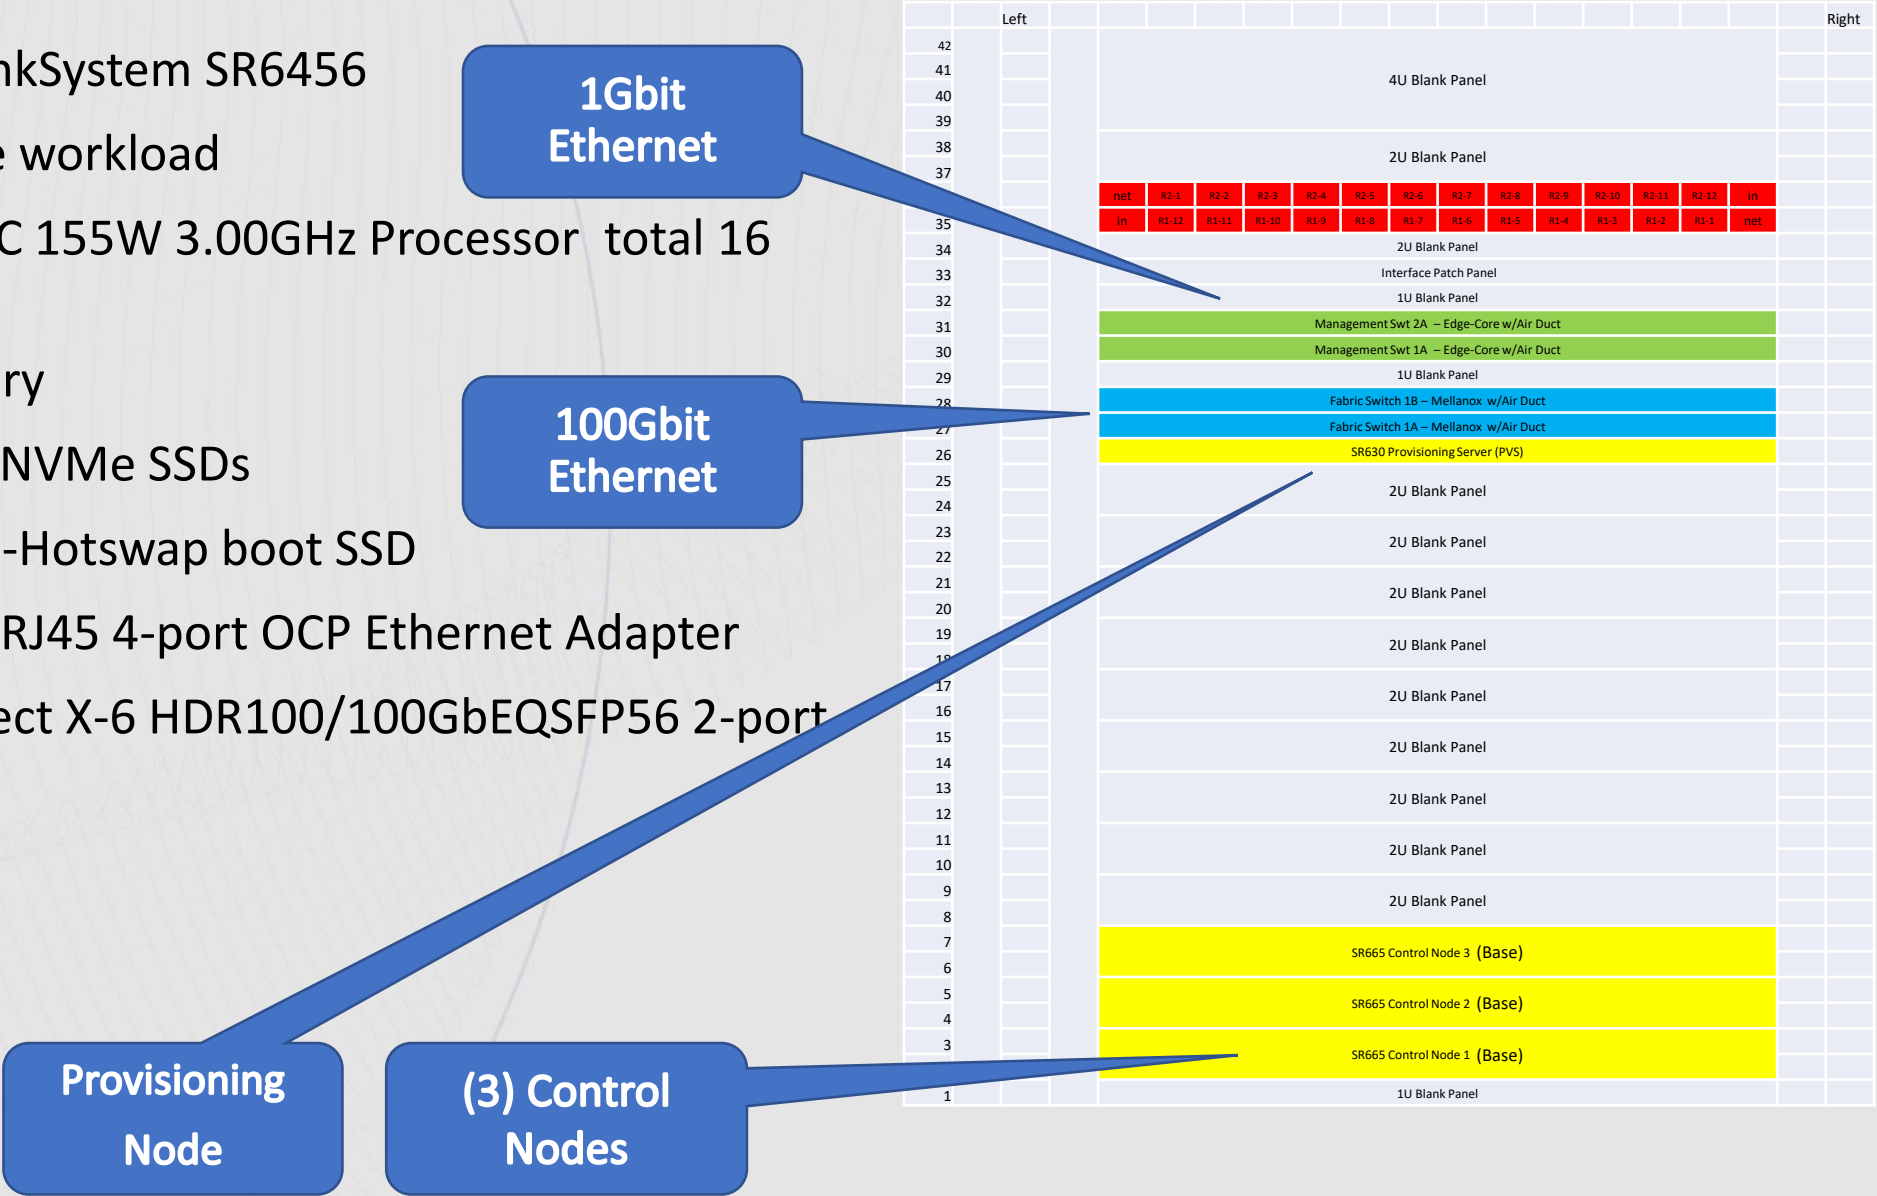
Comparison Point	Cloud Self-Managed RA	On-Premises X86 Ref Arch	On-Premises Power Ref Arch
Hardware Platform	X86	X86	Power 10
Operating Platform	X86 Linux	X86 Linux	Power Linux LE
Db2 Deployment Platform	Db2 AE (BYOL) or Db2U	Db2U	Db2U
Containerization	AKS, EKS, GKE, ARO, ROSA, Rancher	RHOS	RHOS
Cloud Pak for Data	OPTIONAL	OPTIONAL	OPTIONAL
Functional Compatibility	YES	YES	YES
Cloud Provider Options	AWS, Azure, GCP, IBM	N/A	N/A
Support for PB-Size Warehouse	YES	YES	YES
Hardware Included	SOLD SEPARATELY	SOLD SEPARATELY - LENOVO	SOLD SEPARATELY – INTEGRATED SYSTEM
Support for Db2 Vnext	YES	YES	YES
Performance Improvements	YES	YES	YES
Deploy on Cloud Pak System	NO	NO	NO
Availability Target	4Q 2022+	4Q 2022	4Q 2022
Appliance-Like	NO	PARTIALLY	YES

Db2 Warehouse Reference Architecture on X86

- Targeted for 4Q 2022
- Take advantage of pre-defined pre-integrated Db2U on OpenShift Private Cloud optimized solution
- Lenovo HW, RedHat and Lenovo Storage
- Db2 & OpenShift (OCP) Modern & Cloud-native Strategy
- Easy integration to CP4D or Optionally as Services within CP4D
- Ready For 10TB to PB+ Warehouse / Analytical Applications

Control Node Base Unit

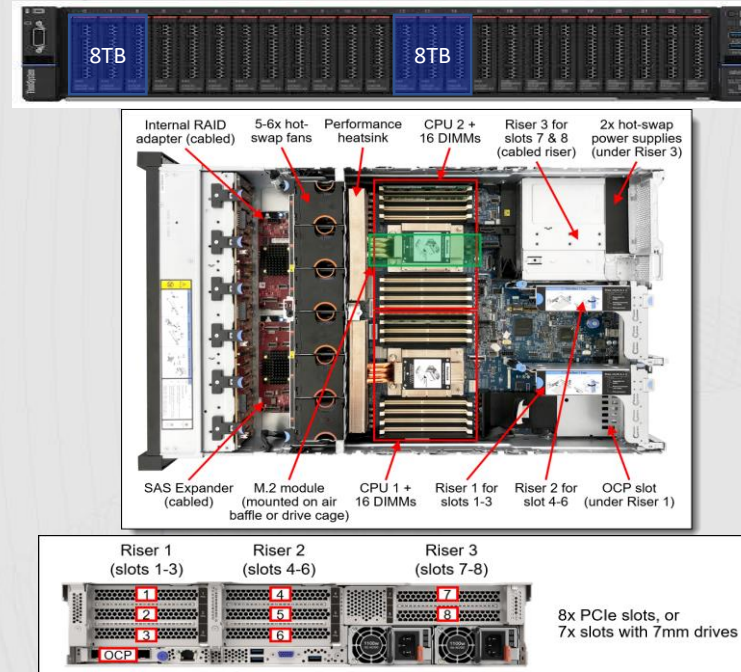
- (3) Control Node : ThinkSystem SR6456
- Sized for Control Node workload
- 1 AMD Epyc 7313P 16C 155W 3.00GHz Processor total 16 Cores
- 128GB 8X16GB Memory
- Four – 2TB PCIe Gen4 NVMe SSDs
- Two – 480GB m.2 Non-Hotswap boot SSD
- Broadcom 5719 1GbE RJ45 4-port OCP Ethernet Adapter
- One – Mellanox Connect X-6 HDR100/100GbEQSFP56 2-port PCIe VPI Adapter



Db2 Worker Node

Server View

- SR665
- Sized for OCP Workers and Database workloads
 - BOM ~\$33k (Assumes WC drive price)
- 2x AMD Epyc 7413 24C 180W 2.65GHz Processor total 48 Cores
 - 2 internal drives Non-hot swap 960GB
- 1TB DDR4 3200MHz 32X64GB Memory
- Six – 7.68TB PCIe Gen4 NVMe SSDs
- Broadcom 5719 1GbE RJ45 4-port OCP
- Ethernet Adapter
- Two – Emulex 32Gb FC Dual-port HBA
- Two – Mellanox Connect X-6
- 4X25Gb Bond Network



(12) Worker Nodes full Rack

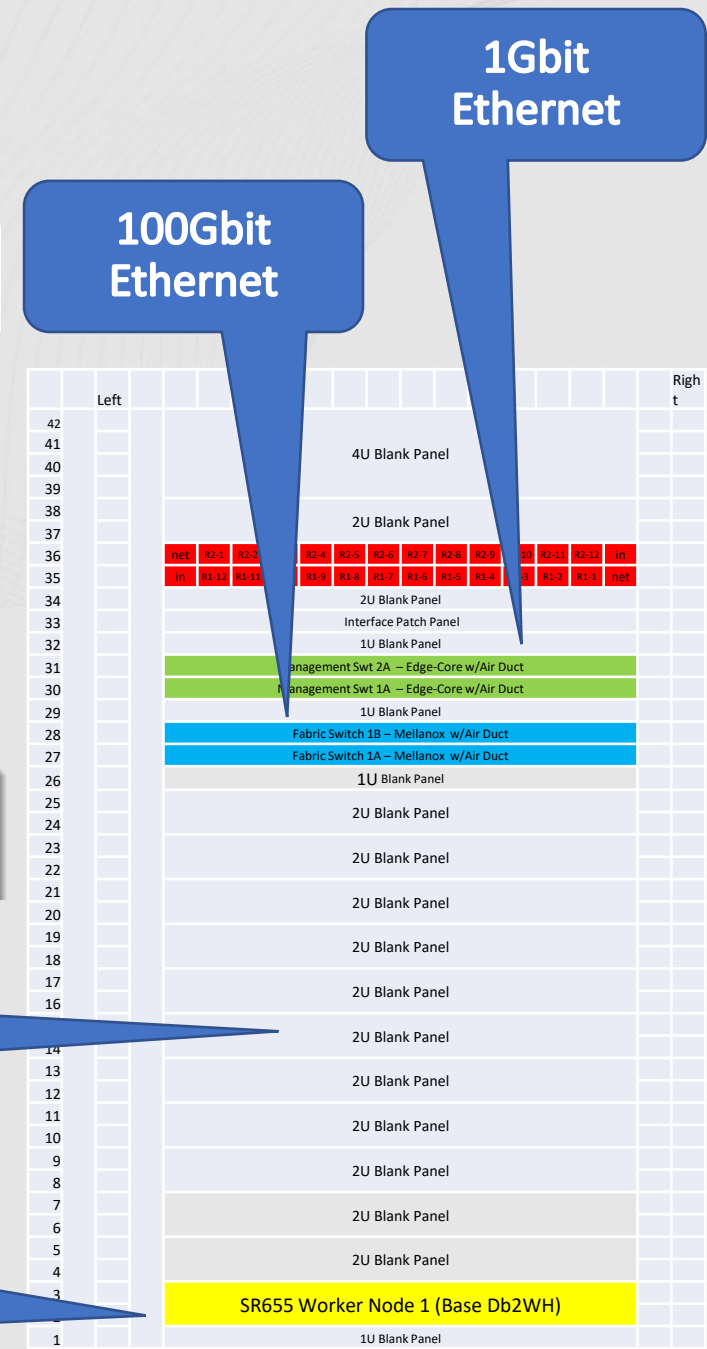
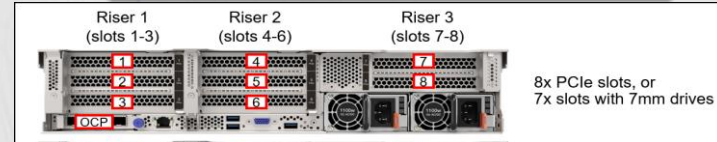
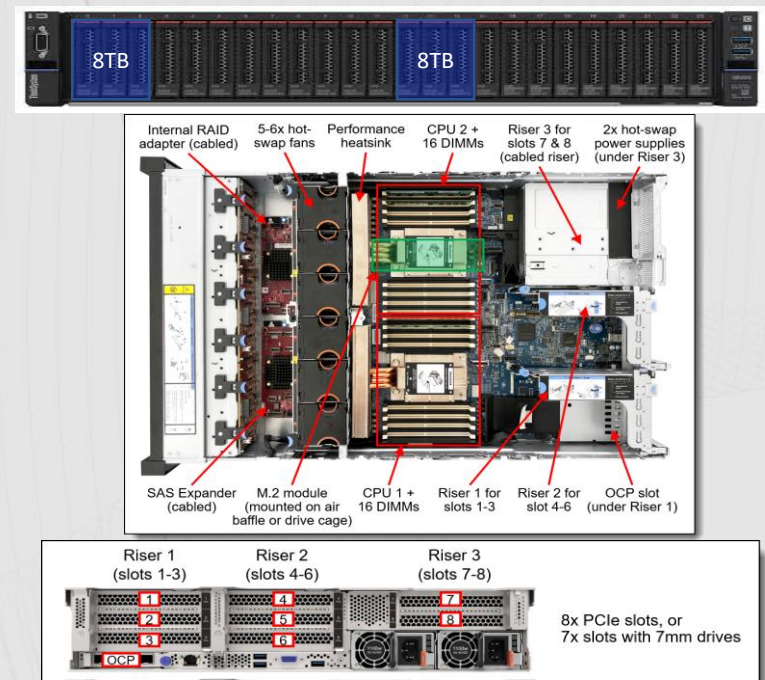
(1) Worker Nodes



Db2 Worker Node

MLN View

- 8 MLNs
 - 6 Cores per MLN
 - 120GB of Memory per MLN
 - ~2TB of Formatted MLN Space (NVMe with 2-way Mirroring)
 - Optional ~4TB 2nd Tiered Storage Formatted per MLN (SSD RAID5)
- Optional Customer Supplied SAN Storage

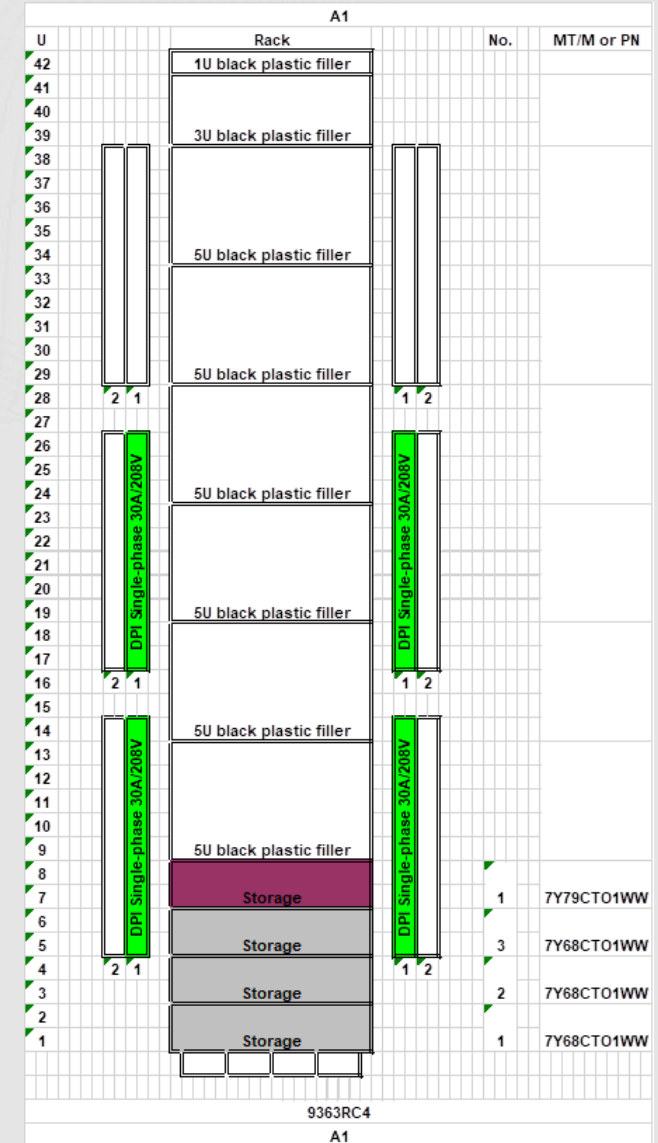


(12) Worker Nodes full Rack

(1) Worker Nodes

Tier 2 Storage Rack

- One IO Controller and up to Three Expansion Draws to support Tier 2 Storage for Four Db2Wh Worker Nodes
 - Controller 7Y79CTO1WW
 - (3) Draws 7Y68CTO1WW
- The Controller and Each Expansion Draw
 - (16) SSD's 3.8TB – Formatted ~49.4TB
 - RAID6



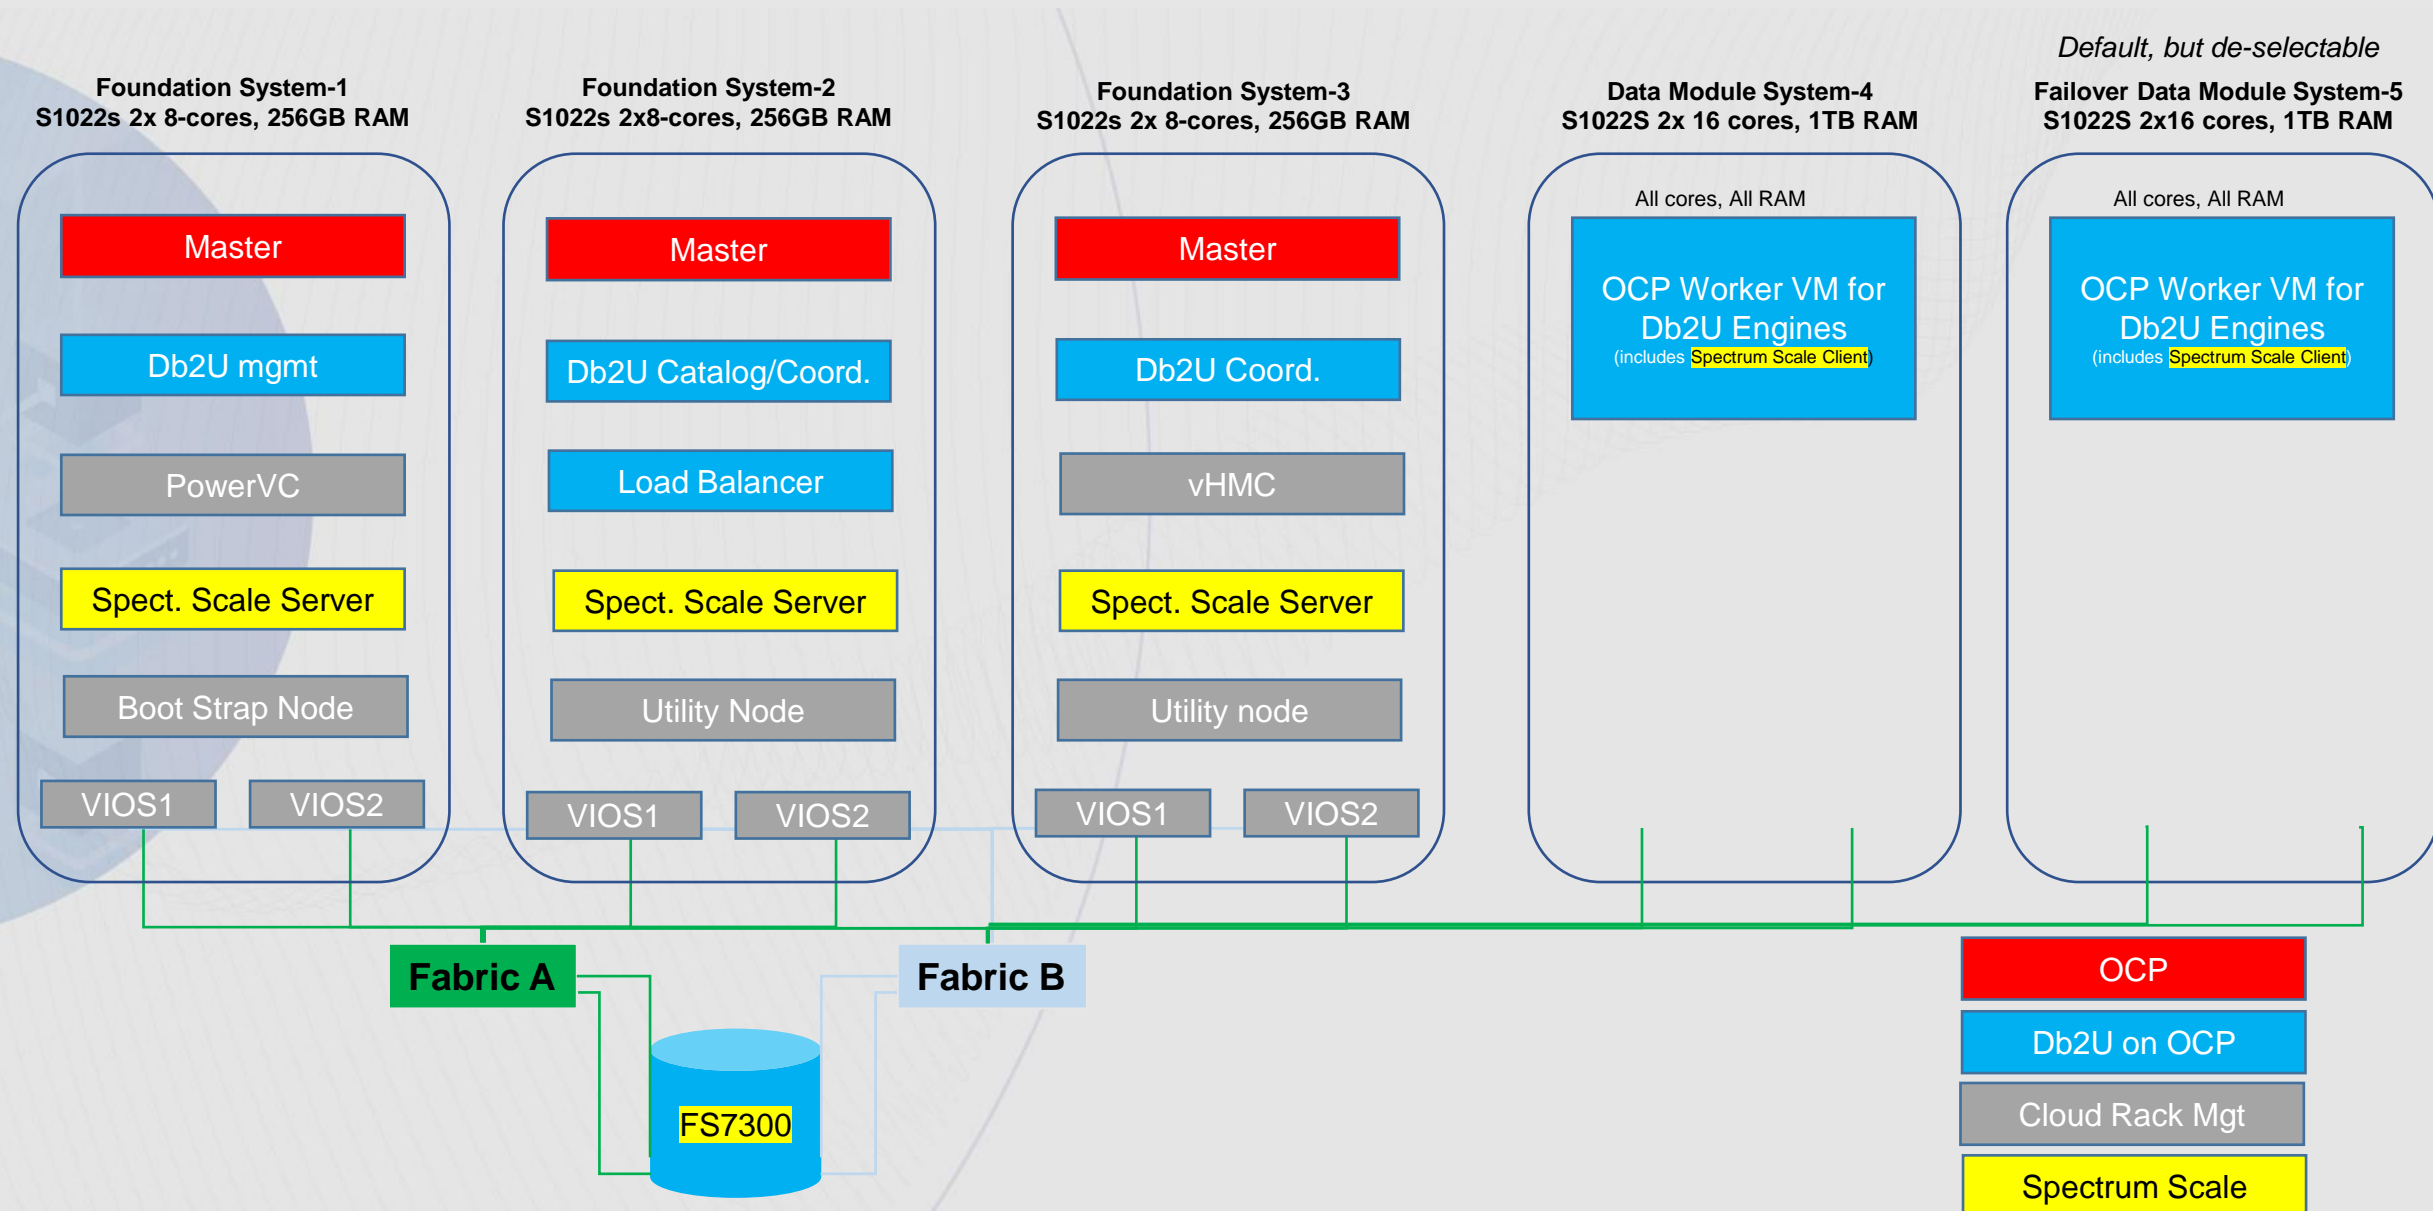
2 - Db2 Warehouse Reference Architecture on Power

Comparison Point	Cloud Self-Managed RA	On-Premises X86 Ref Arch	On-Premises Power Ref Arch
Hardware Platform	X86	X86	Power 10
Operating Platform	X86 Linux	X86 Linux	Power Linux LE
Db2 Deployment Platform	Db2 AE (BYOL) or Db2U	Db2U	Db2U
Containerization	AKS, EKS, GKE, ARO, ROSA, Rancher	RHOS	RHOS
Cloud Pak for Data	OPTIONAL	OPTIONAL	OPTIONAL
Functional Compatibility	YES	YES	YES
Cloud Provider Options	AWS, Azure, GCP, IBM	N/A	N/A
Support for PB-Size Warehouse	YES	YES	YES
Hardware Included	SOLD SEPARATELY	SOLD SEPARATELY - LENOVO	SOLD SEPARATELY – INTEGRATED SYSTEM
Support for Db2 Vnext	YES	YES	YES
Performance Improvements	YES	YES	YES
Deploy on Cloud Pak System	NO	NO	NO
Availability Target	4Q 2022+	4Q 2022	4Q 2022
Appliance-Like	NO	PARTIALLY	YES

Db2 Warehouse Reference Architecture on Power

- Targeted for 4Q 2022
- Leverage Cloud Rack base – Cloud Rack for Db2U
- Take advantage of pre-defined pre-integrated Db2U on OpenShift Private Cloud optimized solution
- IBM Power10, RedHat and IBM Storage FlashSystem 7300
- Db2 & OpenShift (OCP) Modern & Cloud-native Strategy
- Easy integration to CP4D or Optionally as Services within CP4D
- Ready For 50TB to PB+ Warehouse / Analytical Applications

Cloud Rack for DB2 Minimum Configuration



Foundation Compute Module

S1022s

9105-22B

- 2X8 cores per system
- 8 SMT8 cores/socket
- 16 DDIMM slots
- 256GB Memory
- 10 PCIe HHHL slots (8 Gen5 capable)
- 8 NVMe U.2 bays
- Max of 51.2 TB of internal storage

X 3

Control Nodes

- OCP 4.8 Masters
- Db2U Catalog and Coordination
- Cloud Rack Management
- Spectrum Protect Server



Data / Worker Compute Module

S1022s

9105-22B

- 2X16 cores per system 2U
- 8 SMT8 cores/socket
- 16 DDIMM slots
- 1000GB Memory
- 10 PCIe HHHL slots (8 Gen5 capable)
- 8 NVMe U.2 bays
- Max of 51.2 TB of internal storage

Db2U Worker Nodes

- OCP 4.8 Work
- Db2U Data Modules
- Spectrum Protect Client

Db2U Pods

- 8 MLNs
- 33TB Formatted Tier 1 Storage
- Recommended up to 166TB Tier 2 & 3 Storage
- Per MLN
 - 3.+ TB formatted space
 - 100GB Memory
 - 4 Power10 Cores



Storage IBM FlashSystem 7300

FS7300

- 40 cores per System
- 1.5TB of Cache
- Targeting 120,000 Sustained IOPS, 45GB/s
- 1 to 3 Storage Pools
- 12 Storage Class Memory(SCM)
- 24 NVMe SSDs RAID6

SAN Switches

SAN Switches

1 to 3 Worker Nodes

OCP Worker VM for
Db2U Engines
(includes [Spectrum Scale Client](#))

OCP Worker VM for
Db2U Engines
(includes [Spectrum Scale Client](#))

OCP Worker VM for
Db2U Engines
(includes [Spectrum Scale Client](#))



Appliance-Like Experience

A complete set of services-based offerings will be available

Some may be built into special parts for our strategic offerings such as [migration services](#) and [initial deployment services](#) and [deploying DR environments](#) and [maintenance updates](#) and [expansions](#).

Some will be available on demand such as [health checks](#), [performance tuning](#), and [capacity planning](#).

NOTE: Some services offerings above are only applicable to on-premises deployments. Fully managed public cloud deployments take care of some of these activities already.

IDUG

2022 EMEA Db2 Tech Conference



John Bell, Les King
jwbell@us.ibm.com, lking@ca.ibm.com



Please fill out your session evaluation!