EDB Postgres on IBM Power Systems and IBM Spectrum Scale

Lenley Hensarling, EDB
Doug O'Flaherty, IBM

Digital Transformation

Organizations have a business imperative to transform digitally.

- Data must be harvested for greater value.
- Enterprises must use digital assets to enrich their own customers' experience.
- Operations and business processes must become more automated.
- Applications must become more intelligent.

Systems must be modernized to respond in real time to customer needs and increase profit margins.

EDB™
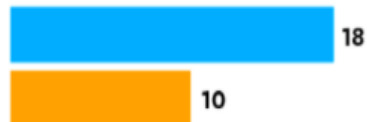POSTGRES

Success in Transformation Means Market Success

'Digital Leaders' outperform 'Digital Laggards'

Based on 2012–2014 data from 344 enterprises listed on U.S. exchanges.

3-YEAR AVERAGE GROSS MARGIN

Digital leaders (top 25% of companies) — 55%
Digital laggards (bottom 25% of companies) — 37%

3-YEAR AVERAGE OPERATING MARGIN

18
10

3-YEAR AVERAGE PROFIT MARGIN

11
7

Harvard Business Review

"Leading organizations are more likely to have a comprehensive data acquisition strategy and differentiate themselves from competitors based on their data platform."

Source: Harvard Business Review, January 2017

EDB POSTGRES

The Data Platform

The Data Platform is the fundamental element of digital transformation in the enterprise infrastructure.

Organizations need data platforms that can integrate data from across multiple sources to better support modernization and the creation of actionable insights.

EDB™
POSTGRES

Data Platform: Open Standards

Solutions based on open standards are better able to integrate disparate databases from across the data center.

Data platforms that utilize open standards are better positioned for digital transformation initiatives.
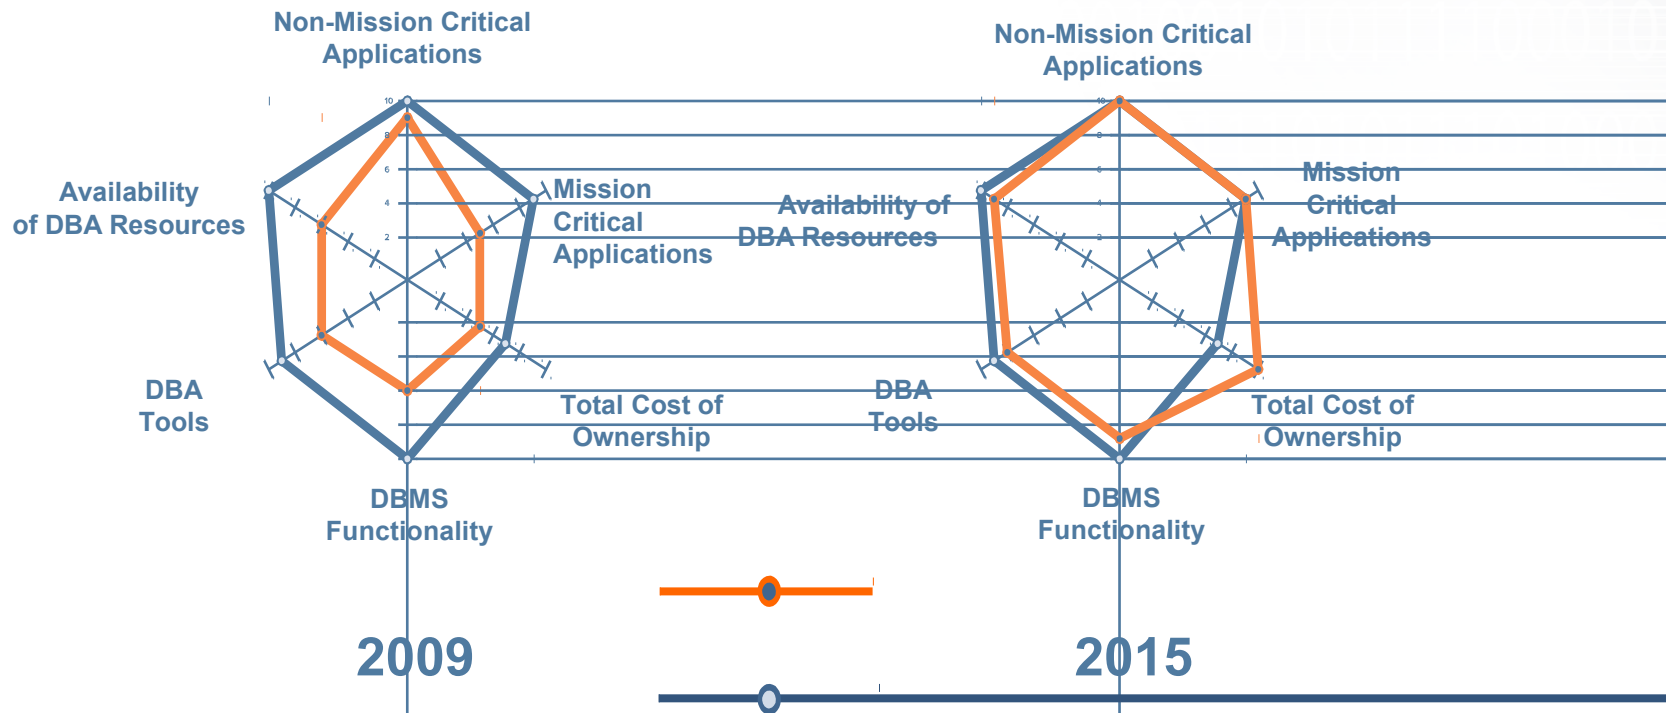
By 2018, more than 70% of new in-house applications will be developed on an OSDBMS, and 50% of existing commercial RDBMS instances will have been converted or will be in process.

Data Platform: Open Standards

Gartner: Relational OSDBMS has matured and today can be considered as a standard infrastructure choice for a large majority for new enterprise applications.

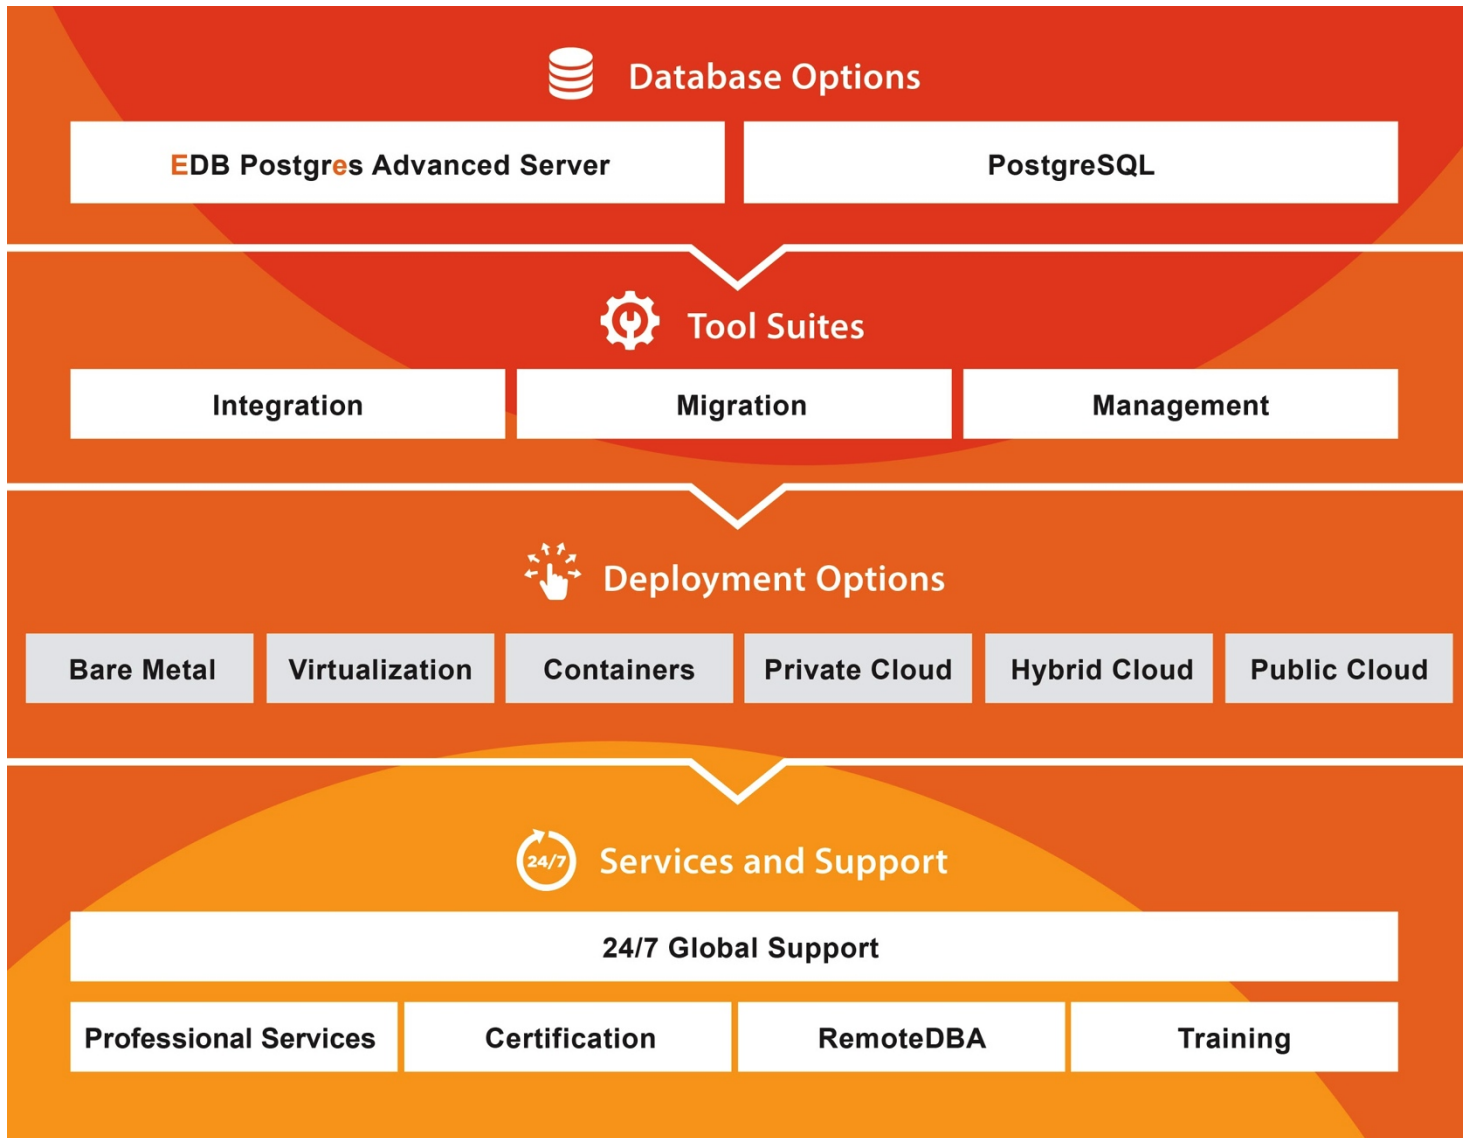*Source: Gartner, State of Open-source RDBMs, 2015, published April 2015*

Mission:

To deliver the premier open source multi-model data platform for new applications, cloud re-platforming, application modernization, and legacy migration. Our platform integrates the EDB Postgres platform with adjacent technologies for hybrid cloud management, data integration and data warehouse. Our customers benefit from the highest performing, most reliable, flexible, open, and cost-effective data management platform available. We envelop our customers with deep expertise and best practices supporting achievement of their goals.

Vision:

To be the strategic multi-model open source data platform for enterprises driving toward digital transformation

EDB
POSTGRES

What we sell: EDB Postgres Platform

Database Options
| EDB Postgres Advanced Server | PostgreSQL |

Tool Suites
| Integration | Migration | Management |

Deployment Options
| Bare Metal | Virtualization | Containers | Private Cloud | Hybrid Cloud | Public Cloud |

Services and Support
24/7 Global Support

| Professional Services | Certification | RemoteDBA | Training |

What we sell: EDB Ark
DBaaS for Hybrid Clouds

Product Value Map

Reliability
- EDB EFM
- EDB BART
- EDB Resource Manager
- EDB Replication Server
- EDB PEM
- Postgres

Inner ring (Reliability):
- High availability
- Disaster Recovery
- Predictable Execution
- Geographic Distribution
- Management & Monitoring
- ACID Transactional

Innovation
- JSONB, KVP, GIS
- Node.js, Python, Java ORB, …
- JSONB & PostGIS
- EDB Ark & EDB Containers
- BRIN, Cube, Set, Rollup, Parallel query
- EDB Data Adapters

Inner ring (Innovation):
- Multi model
- New development paradigms
- Mobile & geo aware
- DevOps, PaaS & Cloud
- Analytics in context of OLTP
- Open source/polyglot data integration

Leverage Existing
- EDB Replication Server
- EDB Postgres Advanced Server
- EDB Migration Tool Kit

Inner ring (Leverage Existing):
- Replication with existing DBMS
- Existing DBMS compatibility
- Automated migration from existing DBMS

Center: **DBMS Platform for Digital Business**

EDB POSTGRES
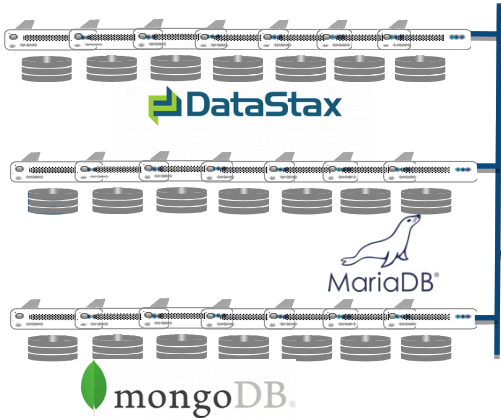
Data Management Landscape

Event Data Capture

Highly Available & Scalable
ACID Multi-model DB
- SQL
- Key/value
- JSON
- Geospatial

Integration
- Hadoop
- Mongo
- MySQL

Oracle compatibility
Replication to/from
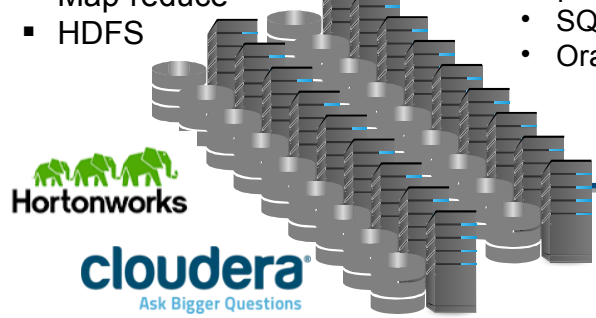- SQL – Server
- Oracle
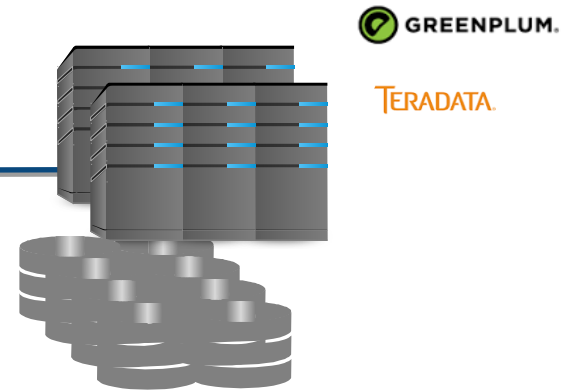
Legacy
Systems
- ERP
- CRM

Big Data
- Hadoop
- Map-reduce
- HDFS

Data Warehouse

Pricing Model

- Flexible, stratified, core-based (2015) pricing model
  - UniCore (universal core) pricing allows universal deployment: on-premises, Virtual/Container, Public, Private, Hybrid Clouds
  - EDB Postgres Enterprise ($1750/uc), Standard ($1225/uc)
  - EDB Ark – DBaaS ($0 EDB; $250/uc PGSQL)
  - Cloud pricing - starting at hourly
  - Custom OEM Embedded and SaaS deals
  - "Fair exchange of value" negotiating principle

- Category leading price-performance
  - Prices enable 80% of enterprise workloads on EDB Postgres for 10-20% of cost EDB Ark $0 with EDB Postgres (PGSQL=$250/uc)
  - Favorable subscription model - pay only for usage
  - Take your subscription anywhere

Recent EDB Postgres Platform Announcements

EDB Postgres Advanced Server 9.6

- Migrate more applications from Oracle and support ISVs

- Manage larger data sets

- Faster analytic queries

- Build more robust clustering, scale out & integration solutions

- Improved monitoring

pgAdmin 4

- New web based administration and development tool

FDW for Hadoop

- Certification with Spark

Backup and Recovery 2.0 Beta

- Faster backups with block level incremental change capture

Replication Server 6.1

- Updated support for Oracle 12c, SQL Server 2012

- Parallel replication for multi master improves performance with multiple active nodes

EDB Postgres Ark 2.1

- Available on over 10 AWS regions

- Support for OpenStack Mitaka Release (Newton / Ocata underway)

Upcoming EDB Postgres Platform Updates

EDB Postgres Platform Containers for Docker and OpenShift

- Easy to configure Postgres clusters running in Containers with read scale and HA

EDB Postgres for Pivotal Cloud Foundry

- Service Broker API for creating, binding/unbinding, archiving and terminating databases

- Option to use EDB Postgres Ark to manage Highly Available and custom defined database engines
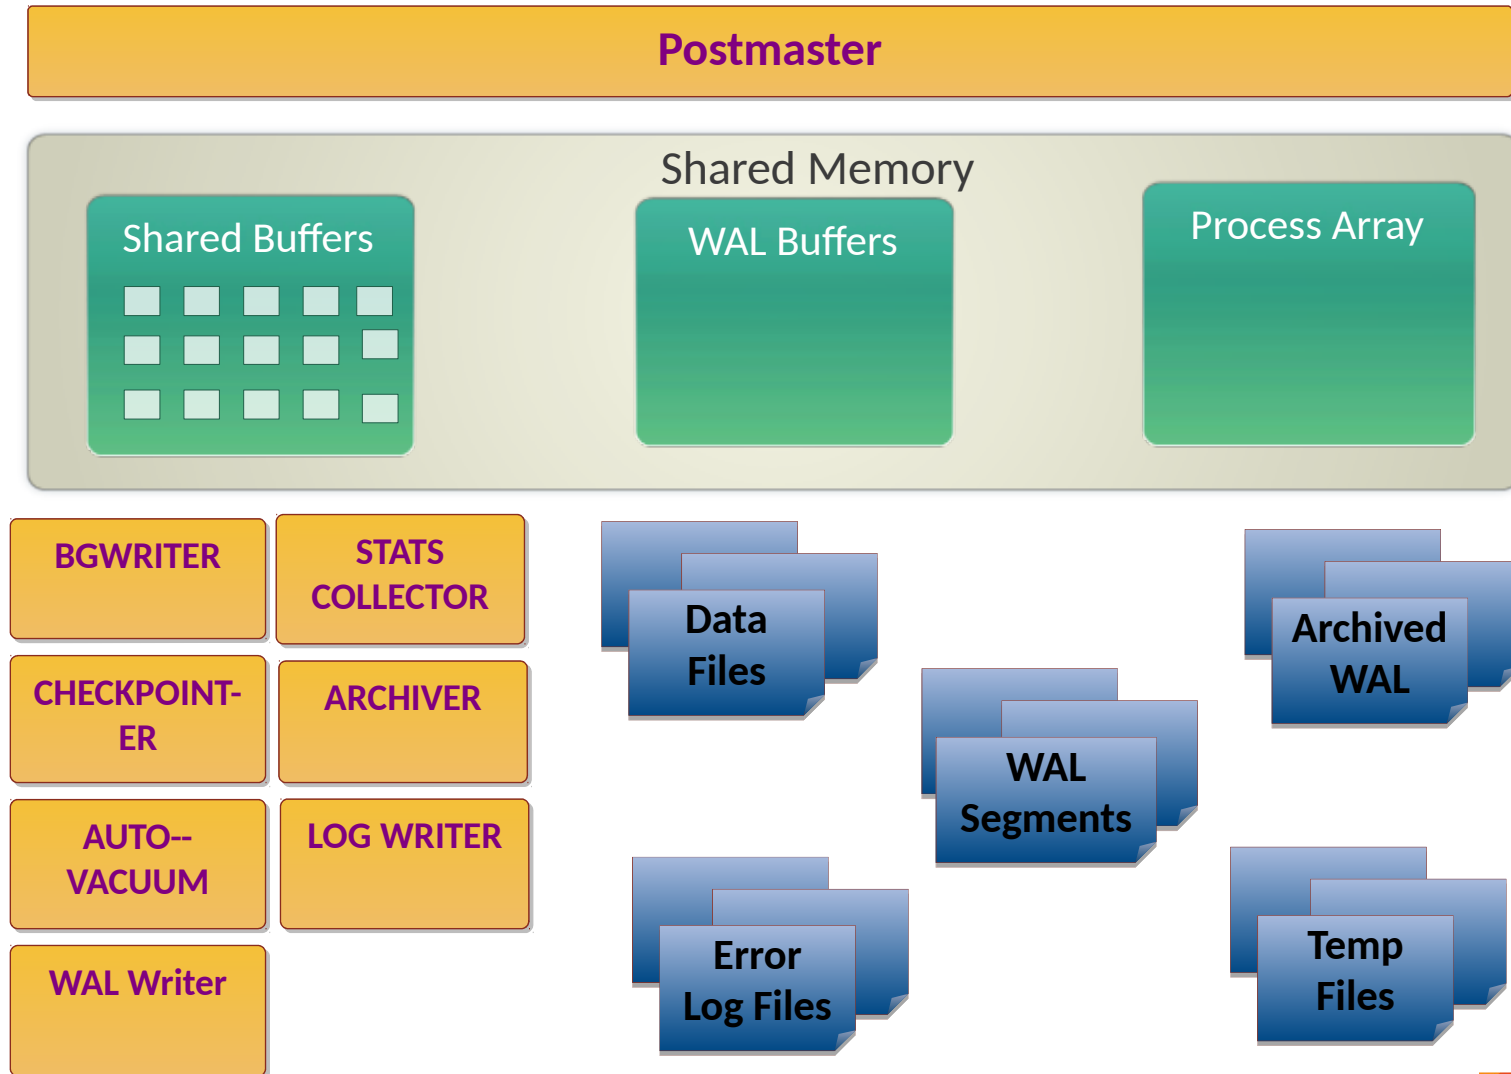
EDB Postgres Enterprise Manager 7

- Easier to use tools and wizards for Managed Servers

EDB Postgres Ark 2.2, 2.3

- Simplified Setup Wizard

- Support for RHEL Subscription Services

- Deploy on MS Azure

- Resource Utilization Reporting

- License Management

Process and Memory Architecture

Postmaster

Shared Memory

Shared Buffers	WAL Buffers	Process Array

BGWRITER

STATS COLLECTOR

CHECKPOINT-ER

ARCHIVER

AUTO--VACUUM

LOG WRITER

WAL Writer

Data Files

WAL Segments

Archived WAL

Error Log Files

Temp Files

Flexible Deployment onto Disk

- Split onto separate physical disks with a matching I/O pattern.

- The Write Ahead Log (WAL) has a sequential I/O pattern.

- Postgres Data Files have a random I/O pattern.

- Consider separating data files, WAL and the OS.

- First move WAL to fastest disk, then your active data.

- Use tablespaces for tiering of the data.

Disk #	Type	Purpose	File System	Profile
1	7.2k or 10k SAS	Operation System	ext4 or XFS	Sequential,Random
2	10k or 15k SAS with 1GB write-back battery backed cache or PCIe Flash SSS	Write Ahead Logs (WAL)	ext4 or XFS	Sequential
3	Tier-1	Postgres Data Files [0-6 months]	ext4 or XFS	Random
4	Tier-2	Postgres Data Files [6 months - 1 year]	ext4 or XFS	Random
5	Tier-3	Postgres Data Files [1 - 7 years]	ext4 or XFS	Random
6	PCIe Flash or Solid State Drive	Postgres Indexes	ext4 or XFS	Random

Query Planner Settings

- **random_page_cost** (default 4.0) - Estimated cost of a random page fetch, in abstract cost units.
    - Reduce this value to be closer to 1 (equal to seq_page_cost) to take let Planner know to prefer index scans. This will take advantage of SSD performance for random access.
    - Can be controlled at a tablespace level so each disk can be optimized for it's technology

- **seq_page_cost** (default 1.0) - Estimated cost of a sequential page fetch, in abstract cost units. May need to be reduced to account for caching effects. Must always set random_page_cost >= seq_page_cost.

- **effective_cache_size** (default 4GB) - Used to estimate the cost of an index scan. Rule of thumb is 75% of system memory.

Memory Settings

- **shared_buffers** (default: 128MB) - Size of shared buffer pool for a cluster.

- **temp_buffers** (default: 8MB) - Amount of memory used by each backend for caching temporary table data.

- **work_mem** (default: 4MB) - Amount of memory used for each sort or hash operation before switching to temporary disk files.

- **maintenance_work_mem** (default: 64MB) - Amount of memory used for each index build or VACUUM.

- **wal_buffers** (default: -1, autotune) - The amount of memory used in shared memory for WAL data. The default setting of -1 selects a size equal to 1/32nd (about 3%) of shared_buffers.

Database Compatibility for Oracle in Advanced Server

POSTGRES

SQL extension support

- Decode, NVL, Substr, NVL2
- Date/time functions: add_months, extract, next_day

DBMS_AQ advanced queuing support **NEW**

PL/SQL support

- REF Cursors, Implicit and explicit cursors
- Looping, variable declarations, conditional statements **NEW**
- Nested Sub Procedures & Functions
- Collections: Associative Arrays, Varrays, Nested tables
- Bulk binding
- Named parameters
- User Defined Exceptions
- Explicit Transaction Control
  - within a stored procedure

Tools

- EDB*Plus – SQL*Plus look-a-like
- EDB*Loader – SQL*Loader equivalent
- EDB*Wrap – similar to the PL/SQL wrapper
- EDB Resource Manager – similar to Resource Profiles

- Features
  - Packages
  - Stored procedures
  - Functions
  - Triggers
  - Hints & Parallel Hints
  - Database Links
  - Hierarchical Queries
  - Synonyms – Public and Private **NEW**
  - Sequences
  - OLD and NEW **NEW**
  - Rownum
  - Object types
    - Create type … as object
    - Create type … as table
    - Create type …as varray
    - Constructor and collection methods
  - Users/Roles
  - Dynamic SQL
  - Password Profiles

Database Compatibility for Oracle in Advanced Server

Data Types
- Integer, number, char, double precision, float, varchar2, blob, clob, xmltype, rowid

Oracle-like Data Dictionary
- ALL_, DBA_, USER_ views
- Most commonly accessed views

Diagnostics - DRITA
- System and session waits
  - Not exposed in PostgreSQL
  - Part of Advanced Server
- Statspack-like reporting

Support for Functions:
- REGEXP_INSTR
- REGEXP_COUNT
- REGEXP_SUBSTR

Oracle compatible partitioning syntax

Oracle compatible Materialized Views

- Package Support for:
  - DBMS_ALERT
  - DBMS_AQ   **NEW**
  - DBMS_CRYPTO
  - DBMS_JOB
  - DBMS_LOB
  - DBMS_LOCK.sleep
  - DBMS_MVIEW
  - DBMS_OUTPUT
  - DBMS_PIPE
  - DBMS_PROFILER
  - DBMS_RANDOM
  - DBMS_RLS
  - DBMS_SCHEDULER
  - DBMS_SQL
  - DBMS_UTILITY

- Package Support for:
  - UTL_HTTP for web server communications
  - UTL_URL
  - UTL_TCP
  - UTL_FILE
  - UTL_MAIL
  - UTL_SMTP
  - UTL_ENCODE

Our Development Focus

EDB continues to driver core enhancements to the internals of the PostgreSQL engine, e.g. parallelism, parallel I/O, and partitioning.

EDB Postgres Ark is delivering on the multi-cloud self-service provisioning vision.

While moving forward with Oracle compatibility, we are also experimenting with and commercializing new application models.

EnterpriseDB on IBM Power Systems Price-Performance Guarantee

IBM Power Systems guarantees the S822LC for Big Data system built with POWER8 delivers at least a 1.8X price-performance advantage versus x86 based servers when running a virtualized customer application/workload based on EnterpriseDB Postgres 9.5.

1.8X price-performance means that the customer's documented throughput performance on the S822LC POWER8 divided by the sum of the price of the system and associated EnterpriseDB licenses will be at least 1.8 times that of the customer's documented throughput performance on the x86 based system divided by the sum of the price of the comparable x86 system and associated EnterpriseDB licenses

EX: If transactions per second on the S822LC are 18,000 and 10,000 on the x86 based system, while the price of the S822LC and associated EnterpriseDB licenses is $10,000, and the price of the x86 based system and associated EnterpriseDB licenses is $10,000, then the Throughput Performance Per Price would be exactly 1.8 times advantaged and the guaranty would be met."

The IBM Power S822LC for Big Data server (20-core/2.92 GHz 256GB memory, 4 TB SATA Storage) must be purchased from IBM or an authorized IBM Business Partner prior to June 30, 2017. The guarantee period is valid for three (3) months from the date of purchase. The x86 based systems must be comparably configured branded servers from Cisco, Dell, or HP and the client is responsible for all EnterpriseDB licenses.

1.8 X price-performance means that the customer's documented throughput performance on the S822LC POWER8 divided by the sum of the price of the system and associated EnterpriseDB licenses will be at least 1.8 times that of the customer's documented throughput performance on the x86 based system divided by the sum of the price of the comparable x86 system and associated EnterpriseDB licenses

Remediation: IBM will provide additional performance optimization and tuning services consistent with IBM Best Practices, at no charge. If unable to reach guaranteed level of price-performance, IBM will provide additional equally configured systems to those already purchased to reach the guaranteed level of price-performance.

Notes:
1. Client's POWER8 Machine and the x86 Machine must be running at similar utilization rates. Eligible Machine and the Compared Machine must be partitioned with at least 4 equal sized partitions.
2. Client's POWER8 Machine's system performance cannot be constrained by I/O subsystem. Specifically, the I/O subsystem on the POWER8 Machines must achieve greater than or equal I/O bandwidth and operations per second than the x86 Machine.
3. Client's POWER8 Machine's physical memory must be the same or greater than the physical memory on the x86 Machine
4. Client is responsible for demonstrating comparable real-world representative workload between the POWER8 Machine and the x86 Machine through the use of the IBM provided tools and comparable tools on x86 systems.
5. 1.8x guarantee is based on list price for the x86 based server (Dell, Cisco,or HP) and list price for the IBM S822LC for Big Data.
6. EnterpriseDB Postgres Advanced Server 9.5 license are priced at $1750 per core - EDB 9.5 http://www.enterprisedb.com/products-services-training/subscriptions-power

Scalable Cloud: Reduce DBaaS operating costs with Virtualized IBM Power Systems

EnterpriseDB Postgres Advance Server 9.5 on IBM Power S822LC for Big Data delivers 2.12x *price-performance leadership* over Intel Xeon E5-2699 v4 Broadwell

EDB POSTGRES / Linux		IBM Power S822LC for BD (20-core, 256GB, 22 VMs) POWER8	HP DL380 (44-core, 256GB, 22 VMs) intel
Server price 3-year warranty		$13,341	$26,698
System Cost Server + RHEL OS + EDB Annual Subscription @ $1,750 per core (3yrs)		$119,640 ($13,341 + $1,299 + $105,000)	$258,997 ($26,698 + $1,299 + $231,000)
EDB pgbench Total Transactions per Second		577,671 tps	588,901 tps
TPS/$		4,828 tps/$	2,274 tps/$

2.12X better

2.1X per core performance

50% Lower HW & solution costs

2.12X Better Price-performance

- Results are based on IBM internal testing of single system in favor performance mode running multiple virtual machines with pgbench select only work load and are current as of August 25, 2016. Performance figures are based on running a 300 scale factor. Individual results will vary depending on individual workloads, configurations and conditions.
- IBM Power System S822LC for Big Data, 20 cores / 160 threads, POWER8; 2.9GHz, 256 GB memory, 2 x 1TB SATA 7.2K rpm LFF HDD, 10 Gb two-port, 1 x 16gbps FCA, EDB Postgres Advanced Server 9.5, RHEL 7.2 with KVM (22 VMs),
- Competitive stack: HP Proliant DL380 ; 44 cores / 88 threads; Intel E5-2699 v4; 2.2 GHz; 256 GB memory, 2 x 300GB SATA 7.2K rpm LFF HDD, 1 Gb two-port, 1 x 16gbps FCA , EDB Postgres Advanced Server 9.5, RHEL 7.2, KVM (22 VMs)
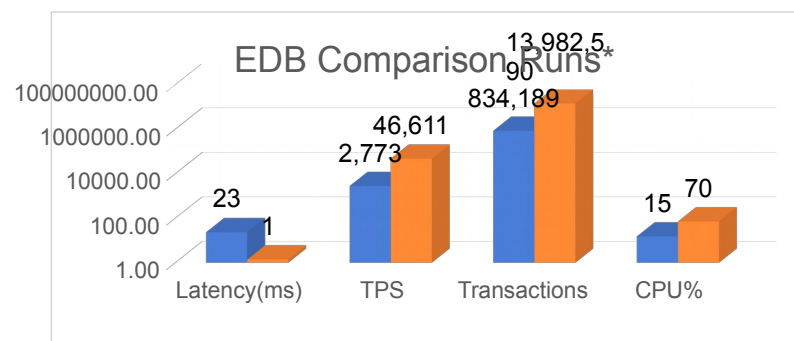- Pricing is based on: S822LC for Big Data http://www-03.ibm.com/systems/power/hardware/linux-lc.html, EDB 9.5 http://www.enterprisedb.com/products-services-training/subscriptions-power and HP DL380 https://h22174.www2.hp.com/SimplifiedConfig/Index

V7000 Flash – 16X more efficient than Traditional SAN

EnterpriseDB Postgres Advance Server 9.5 on IBM V7000 Flash delivers *16X more transactions for a given time period* and 5X *better price-performance* than traditional 15K RPM based SAN

	IBM V7000 Flash	Traditional 15K RPM 22 disk SAN
Storage price	$243,133	$70,711
License Cost 3 years support	$27,960 (3*9,320)	$11,304 (3*3,768.00)
EDB pgbench Total Transactions per Second	46,611 tps	2,773 tps
TPS/$	5.82 $/tps	5.08X better 29.57 $/tps

EDB Comparison Runs*

- 13,982.5*
- 90
- 834,189
- 46,611
- 2,773
- 23
- 1
- 15
- 70

Categories: Latency(ms), TPS, Transactions, CPU%

16x More Transactions and Lower Latency

55% Higher CPU efficiency

5x Better Price-performance

# Oracle vs. EDB TCO Comparison on IBM Power Systems

License Fee Per Core	Oracle Enterprise Edition (32 cores) (Power processor)	EDB Postgres Advanced Server (32 cores) (Power processor)
Database	$47,500	included in subscription
Partitioning	$11,500	Included
Data Guard	$11,500	Included
Diagnostics	$5,000	Included
Total License Fee per Core	$75,500	included in subscription
Total License Fee per Server (CapEx)	$2,416,000	$0
Annual support/subscription cost per core	22% of License Fee	$1,750 per core
Annual Support/Maintenance per Server (OpEx)	$531,520	$56,000
Total 3 Year License and Support Cost	$4,010,560	$168,000

No CAPEX • Annual OPEX reduction **90%** • 3 YR TCO cost savings **96%**

Modernize your Database
with POWER8/KVM and EnterpriseDB vs x86/KVM and Oracle EE

77%
reduction in SW licensing cost with fewer cores and running EDB
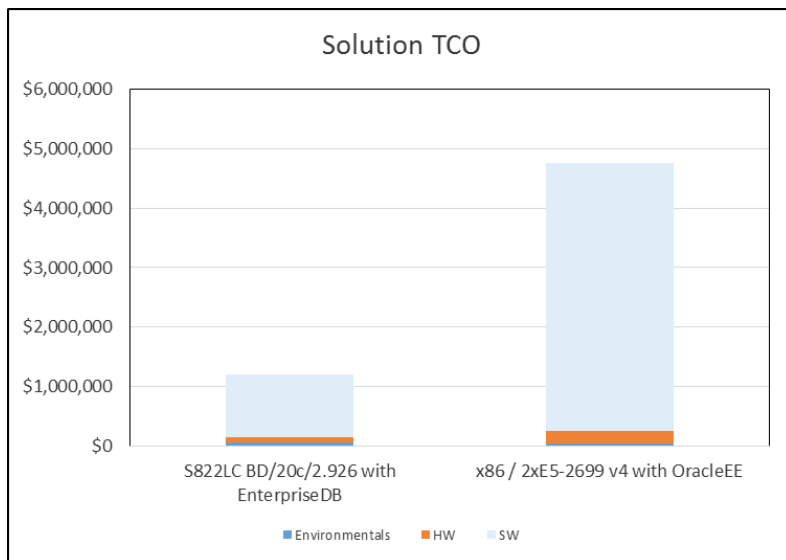
57%
reduction in HW costs and maintenance

EDB POSTGRES

75%
3-year Cost Reduction

$3.5M
3-year Savings
per 10 servers

Solution TCO

- Assumptions:
  - 10xPower S822LC for Big Data/20c servers with KVM (40% utilization) have equivalent performance as 10xHPDL380/E5-2699 v4/44c servers with KVM (40% utilization)
  - Discounts: PowerS822LC – 35%, HP DL380 – 35%, EnterpriseDB – 0%, OracleEE – 70%
  - Perfromance are based on IBM internal testing of single system in favor performance mode running multiple virtual machines with pgbench select only work load and are current as of August 25, 2016. Performance figures are based on running a 300 scale factor. Individual results will vary depending on individual workloads, configurations and conditions.

EDB POSTGRES

EnterpriseDB Postgres Advanced Server

DB Engines Ranking – April 2017

	Rank		DBMS
Apr 2017	**Mar 2017**	**Apr 2016**	
1.	1.	1.	Oracle ➕
2.	2.	2.	MySQL ➕
3.	3.	3.	Microsoft SQL Server ➕
4.	4.	⬆ 5.	PostgreSQL ➕
5.	5.	⬇ 4.	MongoDB ➕
6.	6.	6.	DB2 ➕
7.	7.	7.	Microsoft Access
8.	8.	8.	Cassandra ➕
9.	⬆ 10.	9.	Redis ➕
10.	⬇ 9.	10.	SQLite

- Most mature open source DBMS technology
- Enterprise-Class Features (built like Oracle, DB2, SQL Server)
- Enterprise-Class Support
- Strong, independent community driving rapid innovation

- **Fully ACID Compliant**
- **MVCC**
- **Point in Time Recovery**
- **Data and Index Partitioning**
- **Bitmap Indexes**
- **ANSI Constraints**
- **Triggers & Stored Functions**
- **Views & Data Types**
- **Nested Transactions**
- **Online Backup**
- **Online Reorganization**
- **Foreign Keys**
- **Streaming Replication**
- **Multi-Core Support**
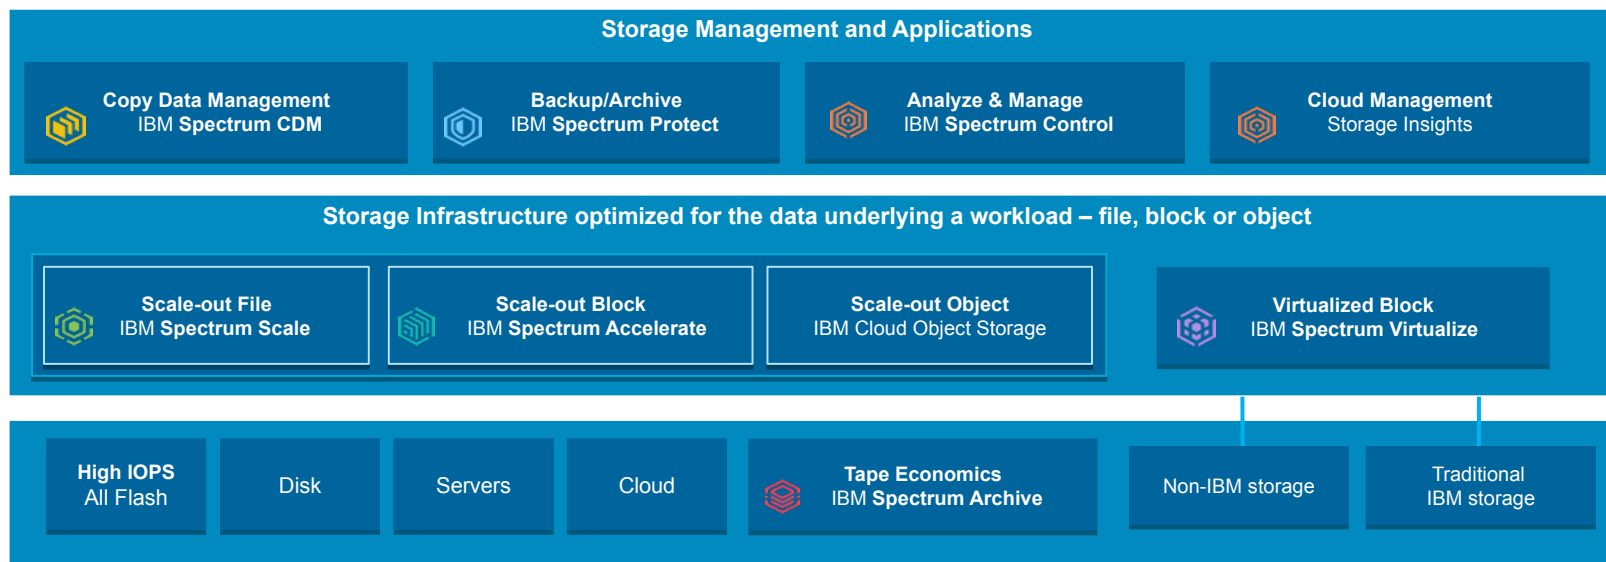- **JSON Support**
- **HStore**

Consider Software-Defined Storage
IBM Spectrum Storage Family

2014, 2015, 2016:
Ranked # 1 in Worldwide Software-Defined Storage Software Market

International Data Corporation (IDC) Worldwide Quarterly Storage Software Qview for Q4 2016 (March, 2017).
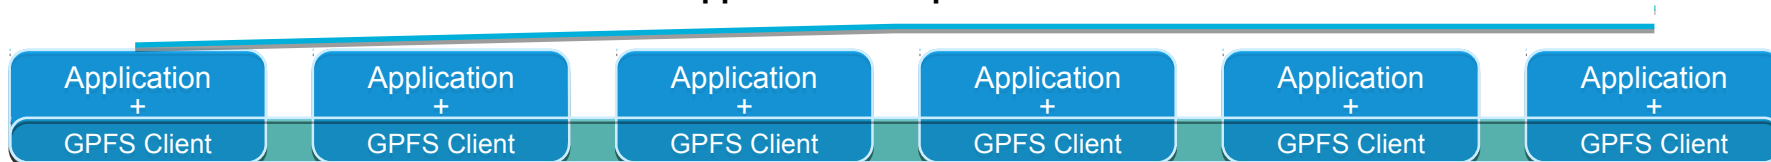
Storage Management and Applications

Copy Data Management IBM **Spectrum CDM**	**Backup/Archive** IBM **Spectrum Protect**
Analyze & Manage IBM **Spectrum Control**	**Cloud Management** Storage Insights

Storage Infrastructure optimized for the data underlying a workload – file, block or object

| Scale-out File IBM **Spectrum Scale** | Scale-out Block IBM **Spectrum Accelerate** | Scale-out Object IBM Cloud Object Storage | Virtualized Block IBM **Spectrum Virtualize** |

| **High IOPS** All Flash | Disk | Servers | Cloud | **Tape Economics** IBM **Spectrum Archive** | Non-IBM storage | Traditional IBM storage |

All of IBM SDS can adapt to high performance or high capacity needs by leveraging appropriate underlying storage media – Flash, NL-SAS or Tape

Available as software, appliance, or as a service in the cloud

IBM Spectrum Scale: The power of a parallel file system

Demonstrated 400 GB/s throughput, building to 2.5TB/s
Local caching for Read and Write

Application/Compute cluster

Application +	Application +	Application +	Application +	Application +	Application +
GPFS Client	GPFS Client	GPFS Client	GPFS Client	GPFS Client	GPFS Client

Direct connection between GPFS client and Spectrum Scale node for data transfer (no centralized metadata node)

Single client can see the full/aggregate performance of the Spectrum Scale cluster, unlike NFS/SMB

Spectrum Scale cluster

Spectrum Scale Server	Spectrum Scale Server	Spectrum Scale Server	Spectrum Scale Server	Spectrum Scale Server

All File System operations are fully distributed on both the
client and the server side for extreme of performance and capacity

IBM Spectrum Scale Performance

Data on Flash: Data is more important than Metadata
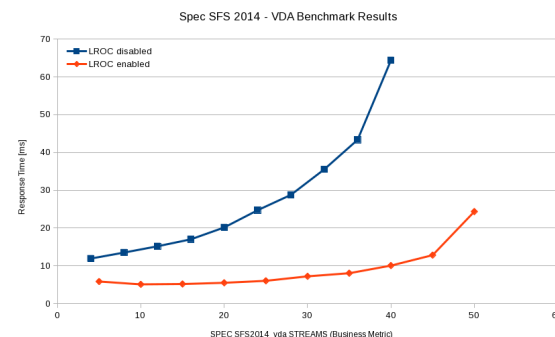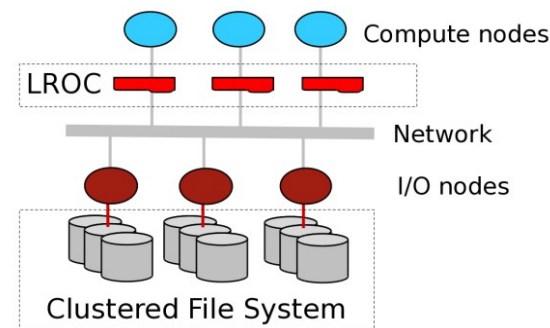
Take advantage of GPFS protocol
- Parallel File Access & Distributed data read/write
- Requires fast networks

Turn on local caching
- LROC: Local Read Only Cache : Highly recommended
- HAWC: High Availability Write Cache: optional.
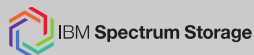 - Many storage systems, such as IBM ESS, have local caching

Consider other SDS options and uses
- IBM Spectrum Scale: High-performance backup target w/ unified File, Object & HDFS
- IBM Cloud Object: On-premises *and* as-a-service with IBM BlueMix
- IBM Spectrum Virtualize: SAN virtualization w/ tiering
- IBM Spectrum Accelerate: Scale-out block storage

https://www.intel.com/content/dam/www/public/us/en/documents/white-papers/performance-gains-ibm-spectrum-scale.pdf

IBM Spectrum Scale & Elastic Storage Server (ESS)

IBM Spectrum Storage

Spectrum Scale

IBM Power8 Linux Server

"Twin Tailed" JBOD Disk Enclosures X TB Drives

← Capacity →

Model GL2:
2 Enclosures, 12U
116 NL-SAS, 2 SSD

Model GL4:
4 Enclosures, 20U
232 NL-SAS, 2 SSD

Model GL6:
6 Enclosures, 28U
348 NL-SAS, 2 SSD

← Speed →

Model GS1
24 SSD

Model GS2
46 SAS + 2 SSD
or 48 SSD Drives

Model GS4
94 SAS + 2 SSD
or 96 SSD Drives

Model GS6
142 SAS
+ 2 SSD Drives

Running Red Hat Enterprise Linux

IBM and EDB

Bringing Cost Optimized Performance to the Enterprise

- Working together to optimize Postgres for IBM Power Systems

- Collaborating on deployment architectures

- Driving the value of open source on POWER

EDB™
POSTGRES