Linux on IBM Z / LinuxONE Open Source Ecosystem Status and Strategy

for NY/NJ Linux Council Meeting on March 1, 2019

Enyu Wang

Program Director, Ecosystem Strategy and Business Development

enyuw@ca.ibm.com

As an enterprise platform

WHY ARE WE INVESTING IN OPEN SOURCE ECOSYSTEM?

TREND: Enterprise Going Open Source

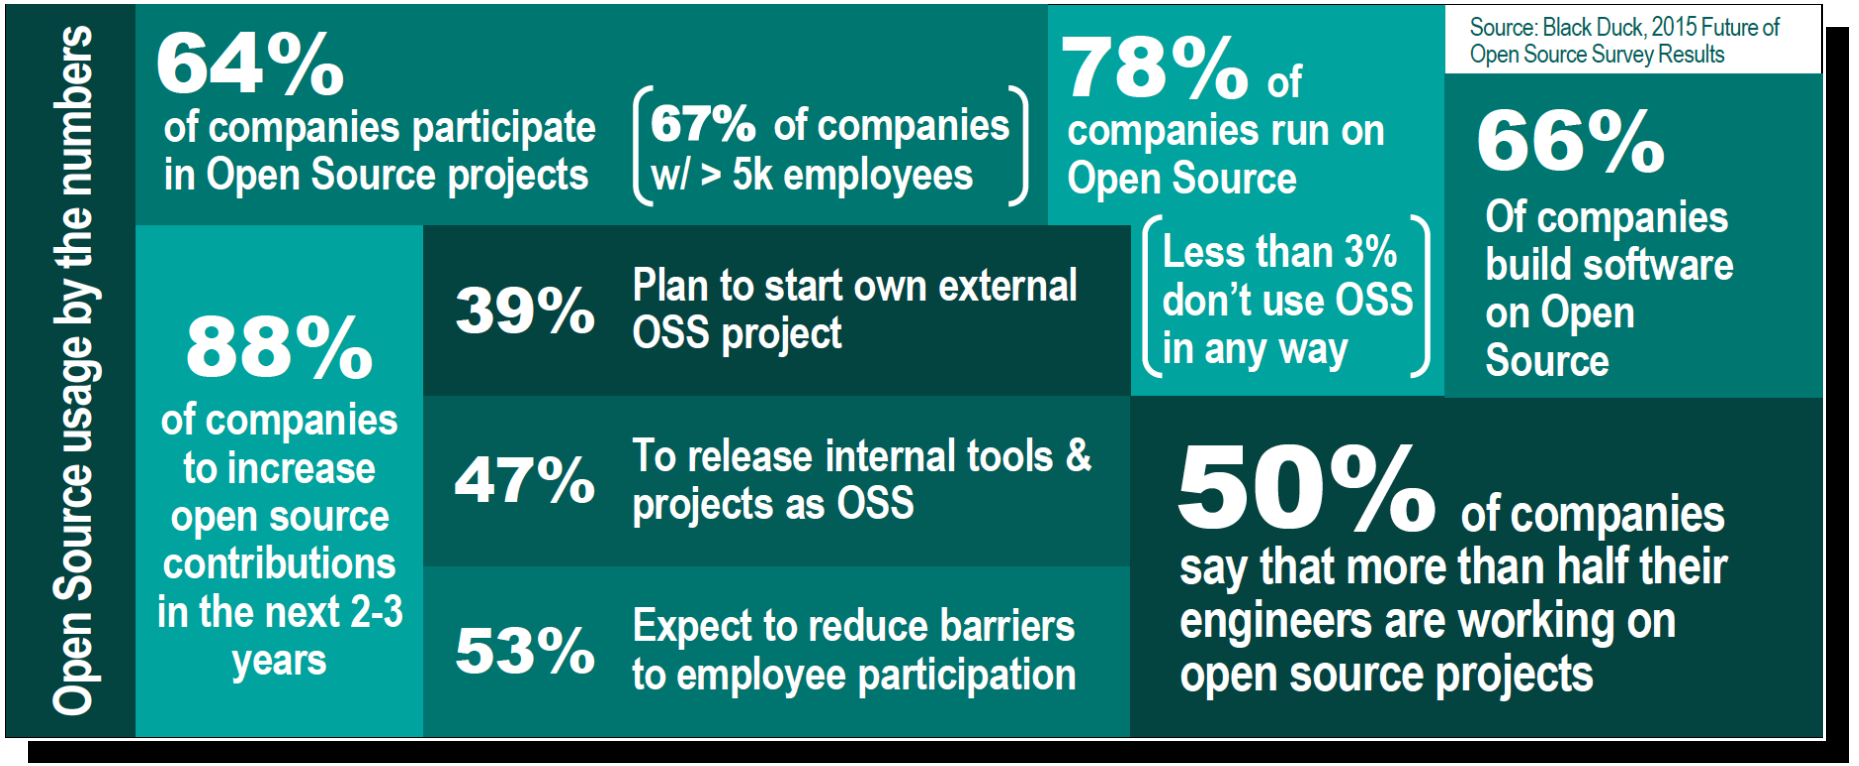

- **83%** hiring managers surveyed for the 2018 Open Source Jobs report said hiring open source talent was a priority this year
- Some of the biggest trends in enterprise IT, such as **containers** and **hybrid cloud**, rely on open source technologies including **Linux** and **Kubernetes**

OPEN SOURCE
Building Blocks for Enterprise Digital Transformation

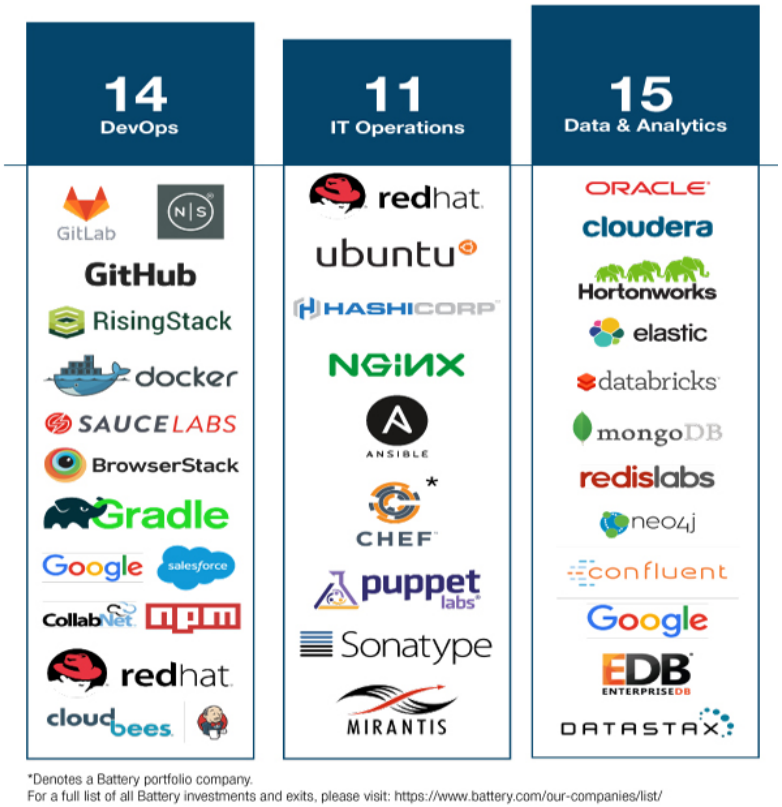

OUR MISSION

Provide a Rich and Robust Ecosystem to Clients. Help Accelerate their Digital Transformation

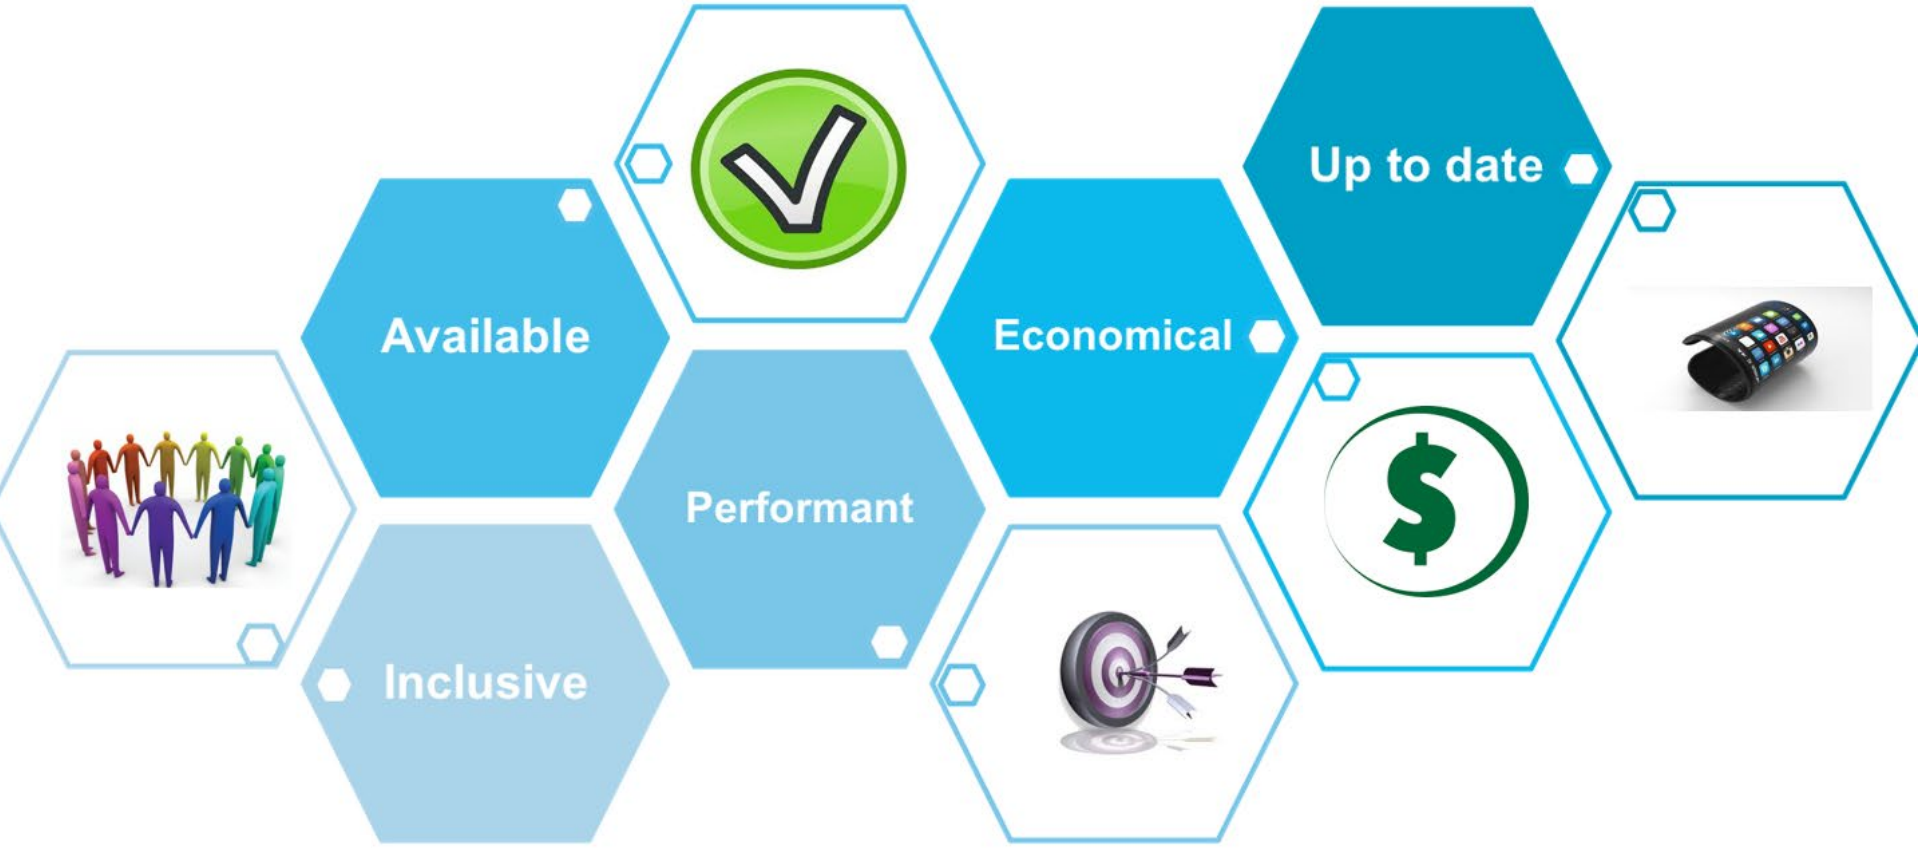

Rich Open Source Ecosystem on Linux on Z/LinuxONE

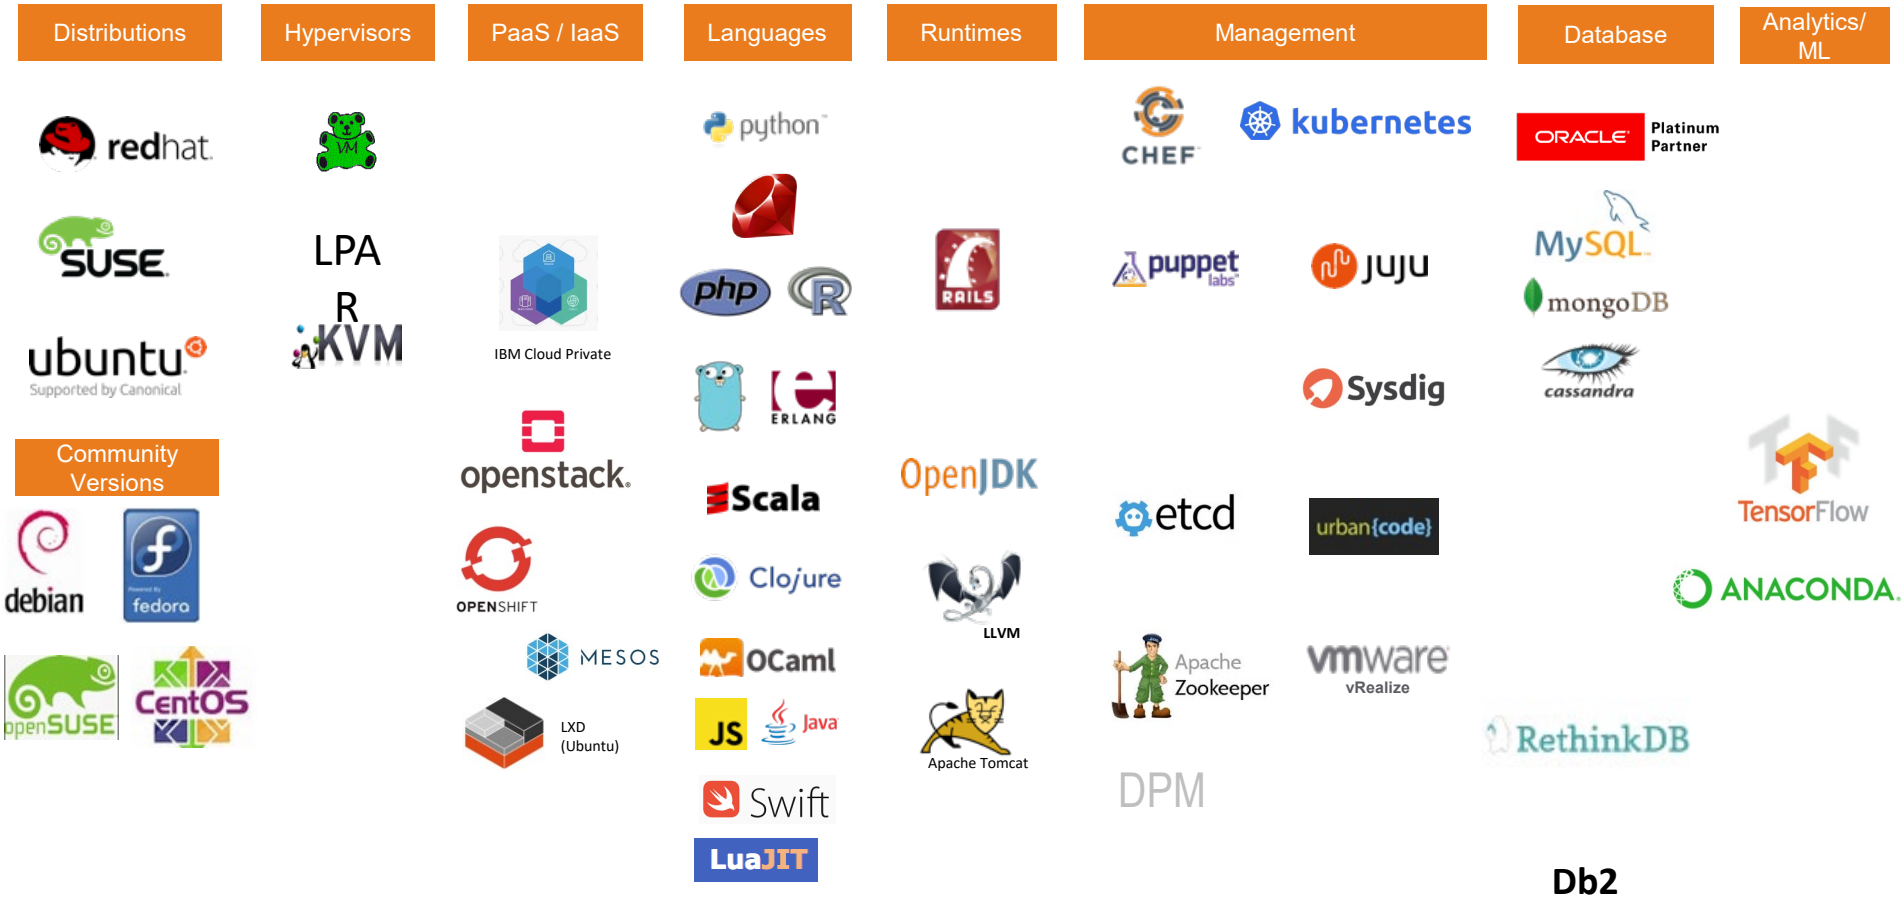

Building an Open Ecosystem Isn’t Just Porting…

Composition of Open Source Ecosystem on Z

– a combination of community based projects and vendor governed projects

https://www.ibm.com/developerworks/community/groups/community/lozopensource

Objective for technical enablement and support on Z is making Z part of supported architecture for the project

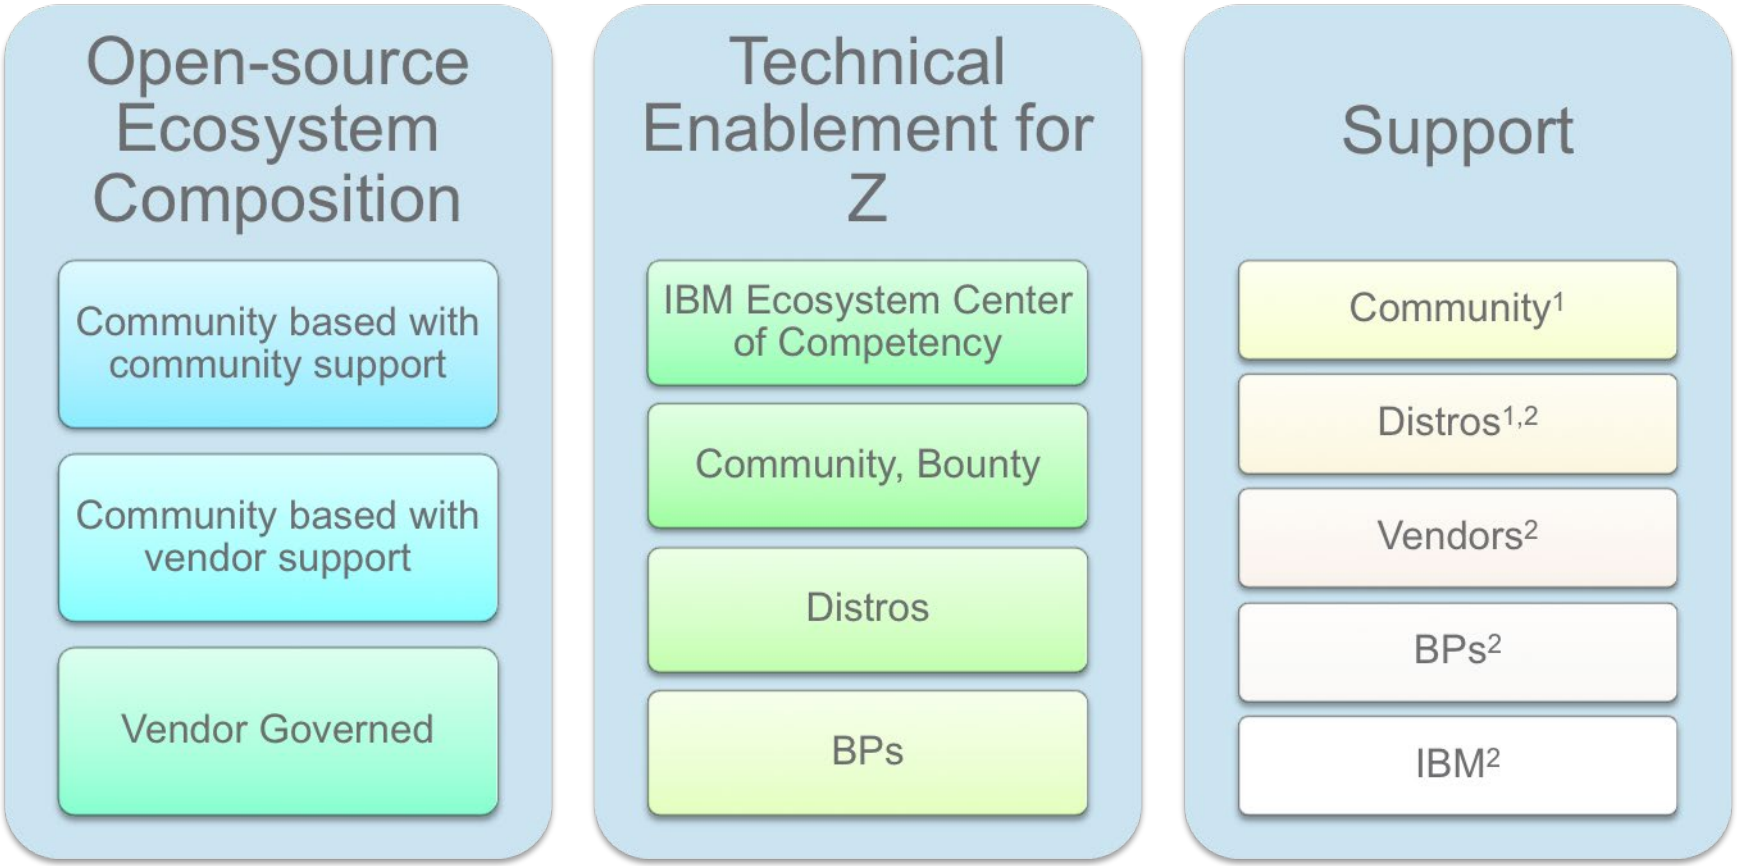

[1]Community based support with no SLA

[2] Enterprise support (9x5, 24x7, etc)

Open Source Characteristics and What It Means to Z

Wild Wild West	• Tons of open-source projects out there - need to identify technology areas and top projects that are good match for z client segment. • Rapid innovation and disruption – requires vigilant monitoring of trend • Deep market research is required
Community driven	• Developers not aware of Z architecture (e.g. big endian, scale up) • We need to follow open-source development model and contribute to communities • We need to keep z active in the mainline code. Stay engaged with the community. Do not fork.
Fast paced	• Making keeping currency a challenge • Enterprise clients tend to stick to stable version or want the latest for a specific feature
Business model	• Freely available. Subscription and support are cheap relying on volume. No huge software licensing saving like Oracle • Who to drive and how to track client adoption

Z Open Ecosystem Work Flow

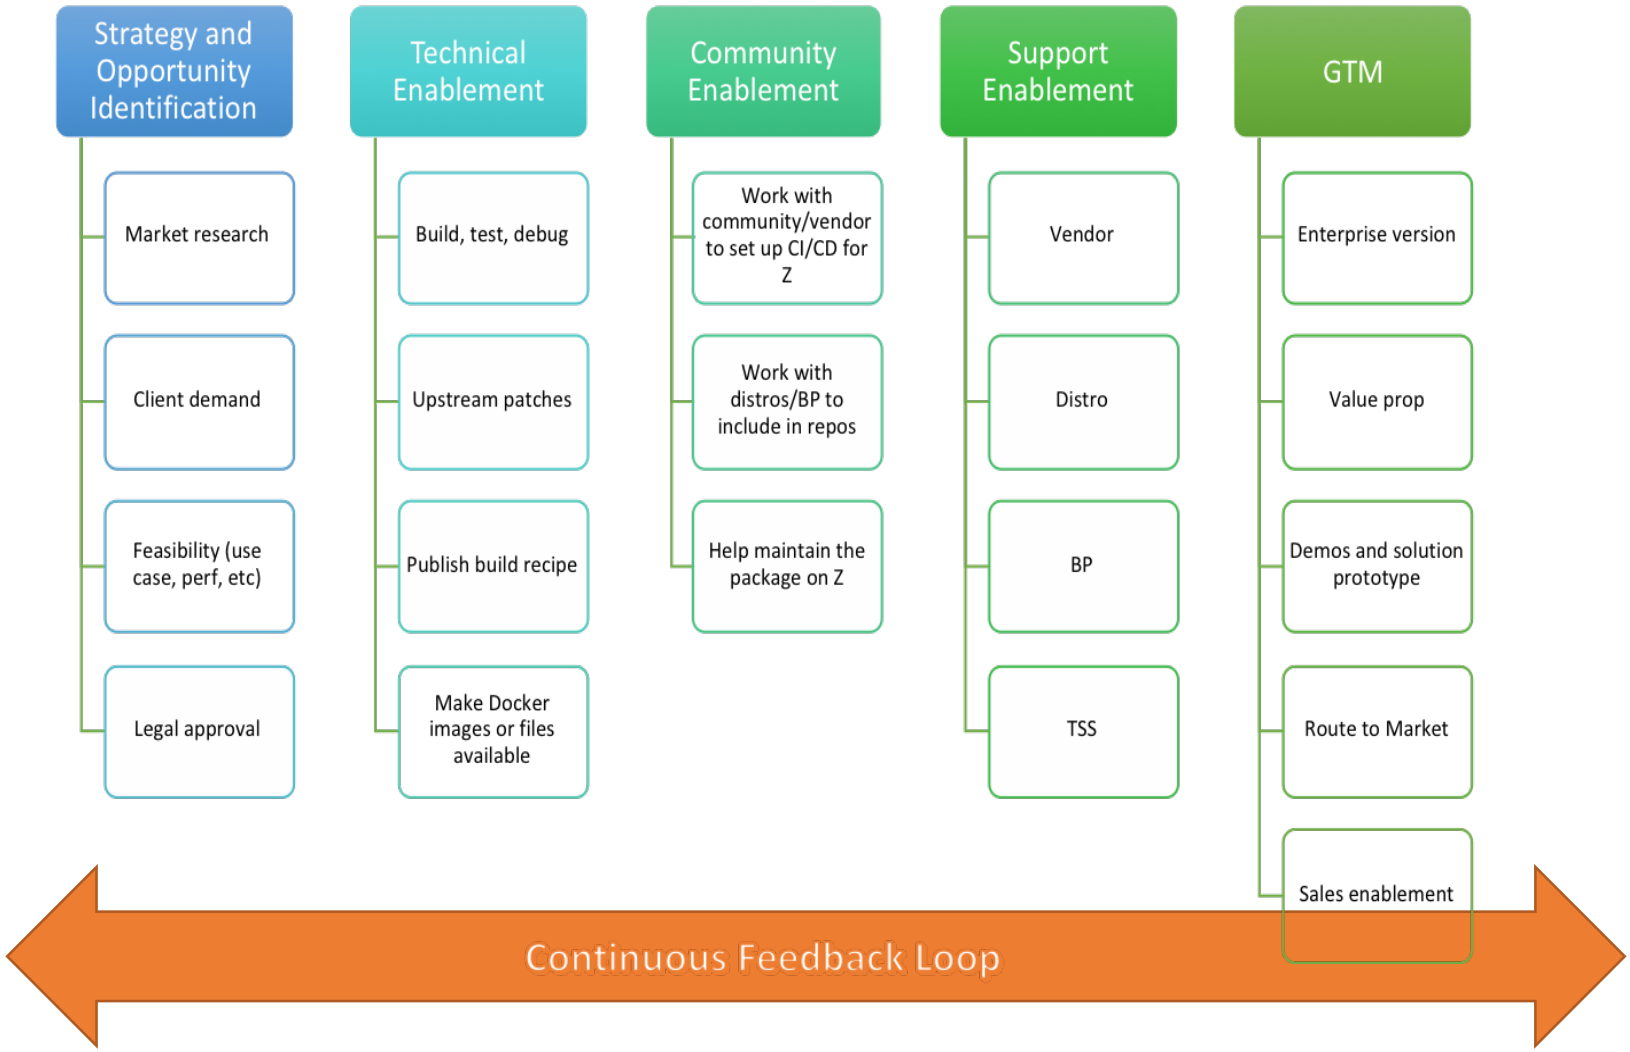

Make More Binaries Available on Z

Request for Binaries

- Customers don't always want to build open source binaries from source, especially for complex Often times binaries can be downloaded from a website or install from repos
- Challenges for Z: Z hardware isn't pervasively used as x86 for building packages

Multi-proned Approach on Z

- Leverage distro repos
- Invite communities to build on Z on LinuxONE Community Cloud and Dallas ISV Center
- Make z QEMU available for packages that are built using emulation
- Enable build services (Jenkins, Travis-CI) on Z
- Publish Docker images on Docker Hub through a Docker build tool hosted at UIUC

Support for Open Source

Support options on x86

- Community based support
- Distro support
- Vendor support

Support options on Z

- Community based support
- Distro support
- Vendor support for a selection of software
- Business partners provide enterprise support for popular packages
  - RogueWave, Rocket Software, SNA
- IBM (TSS) support

Technology Focus Areas

What's enabled, what is being worked on, what is in the pipeline

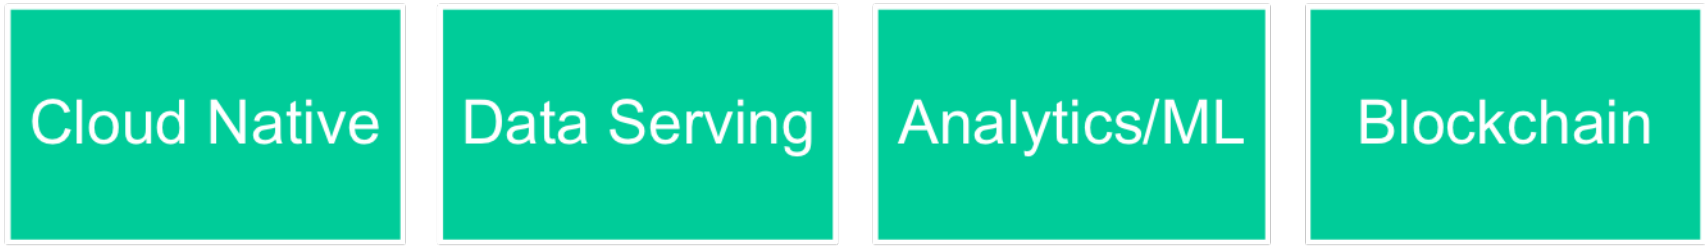

Market Challenge with growth opportunities

Today's multicloud reality presents new opportunities for clients, as well as new challenges

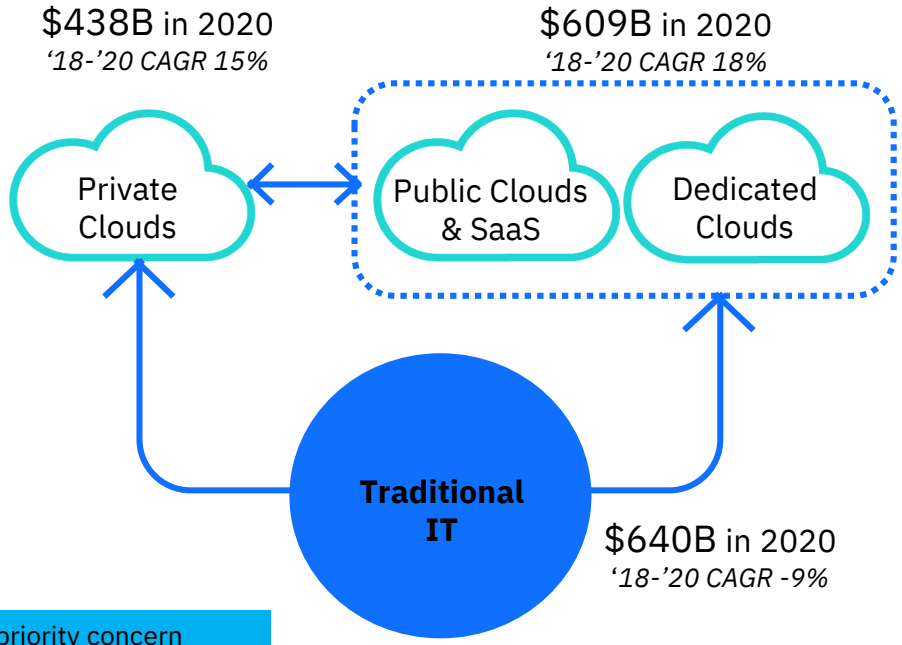

A real world look at multicloud

94% **of enterprise customers are using multiple cloud environments**
(public and/or private)

67% **of enterprise customers are using more than one public cloud provider**
(expected to remain constant or increase by 2022)

priority concern

Movement between clouds	73%
Connectivity between clouds	82%
Consistency of management	67%

Source: IBM MD&I; BCG and McKinsey research

Containers Are the Foundation for Cloud Native Apps

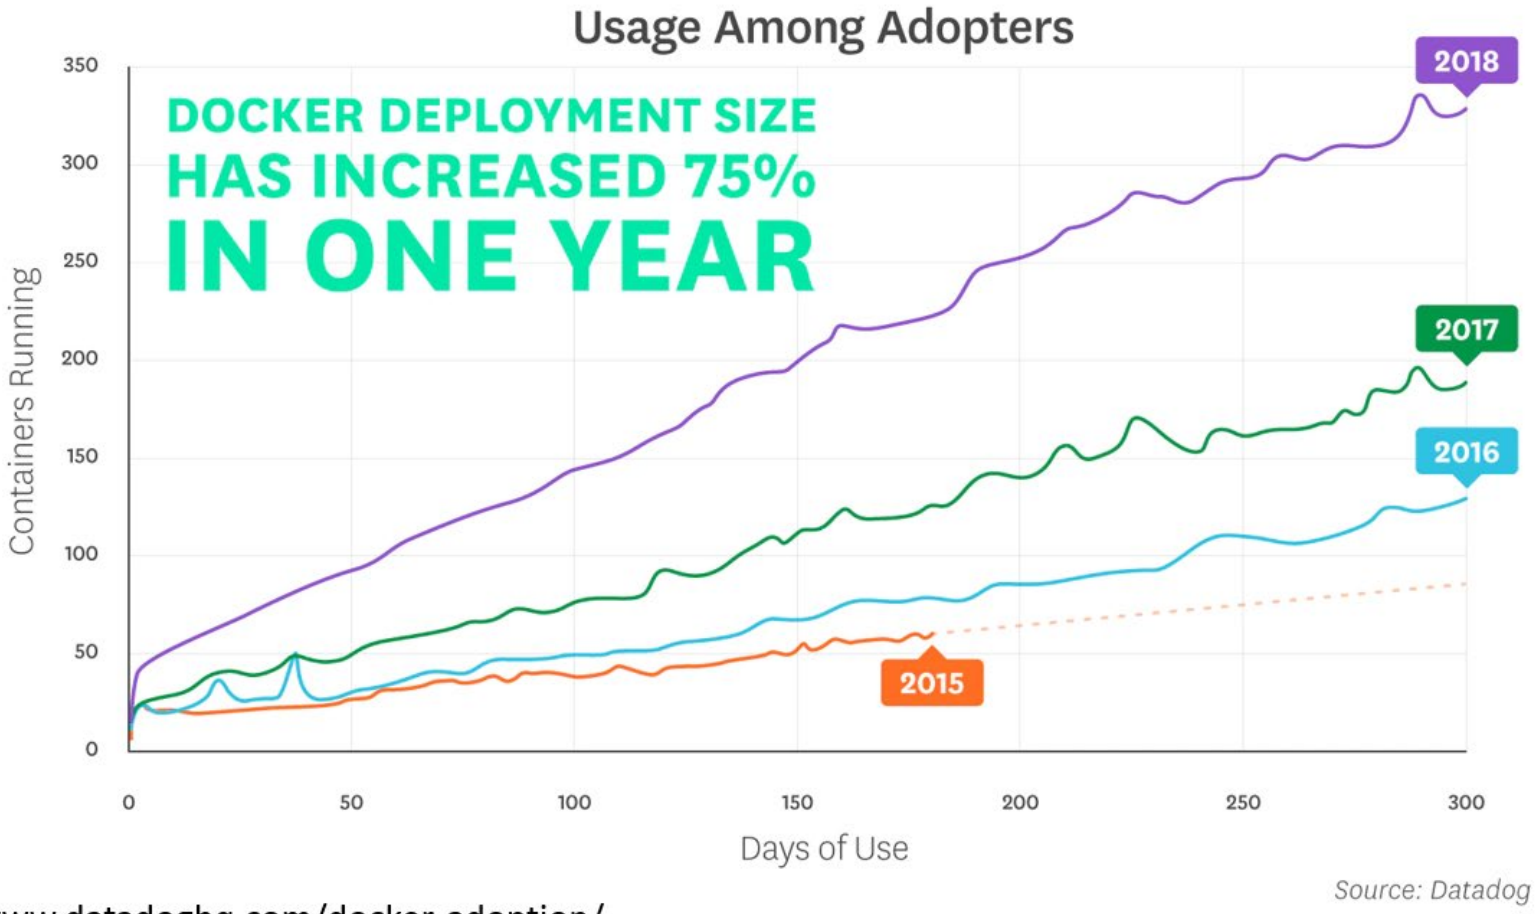

https://www.datadoghq.com/docker-adoption/

Where Is Z in Cloud Native Technology Adoption

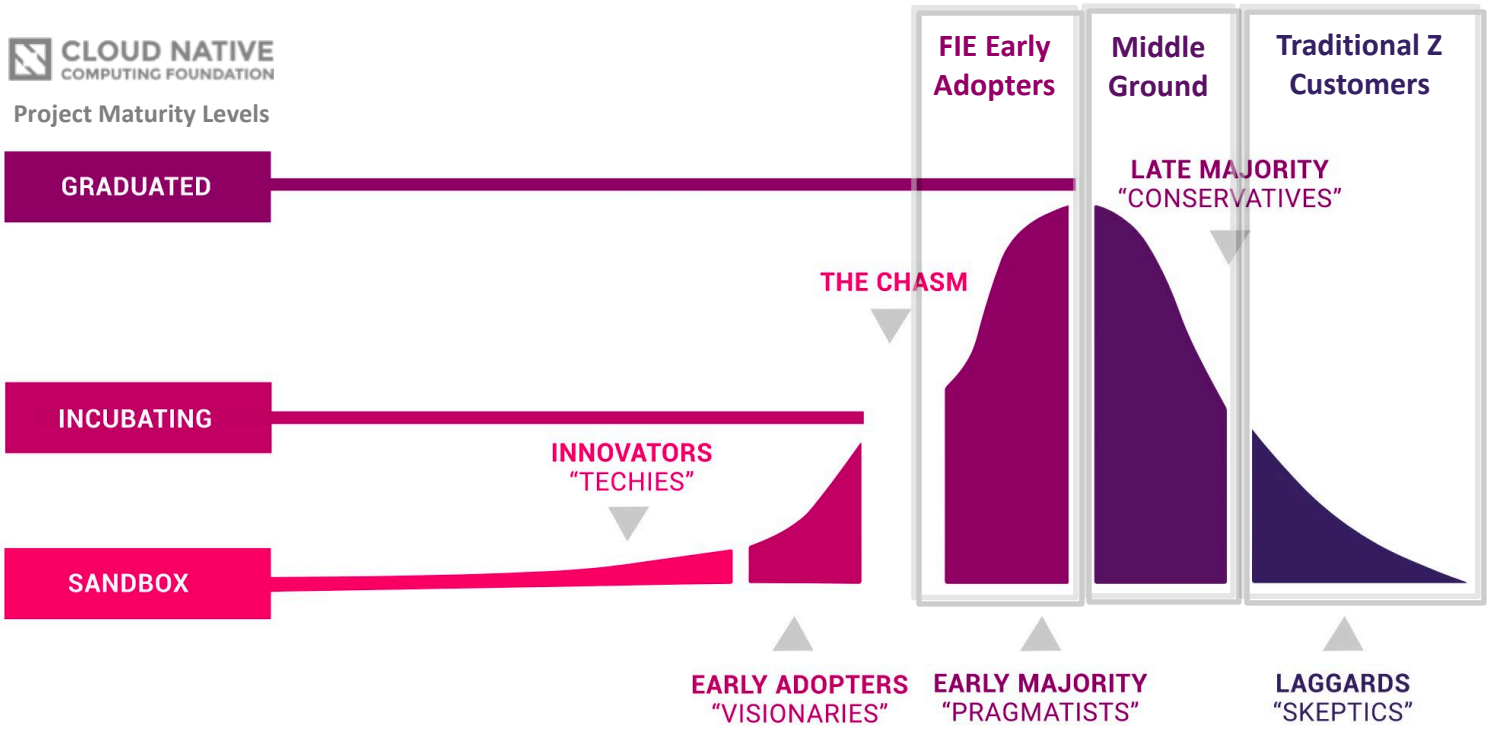

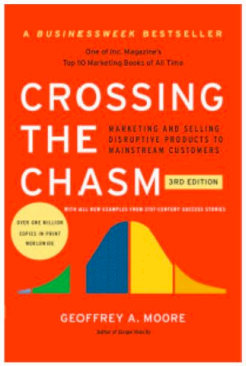

Cloud Native Ecosystem Strategy

Align with Tech Trend

- Shift from VM to Containers
- Kubernetes are the new Linux
- as a Service is the trend
- Hybrid Cloud goes mainstream
- Customers requires cloud like experience for on prem, ie. life and shift on converged infrastructure
- Open-ness: Rometty said the battle for cloud computing will be 'open versus proprietary', and she has placed all her backing behind open source.

Make Z the most secure server for cloud

Key Differentiator: Security

No other Linux server can deliver more protection against both internal and external cyber threats – faster and without changes to applications.

Other Differentiator

Performance

scalability

HA/DR

Consolidation

zOS Cloud Broker

What a Cloud Platform Needs to Deliver

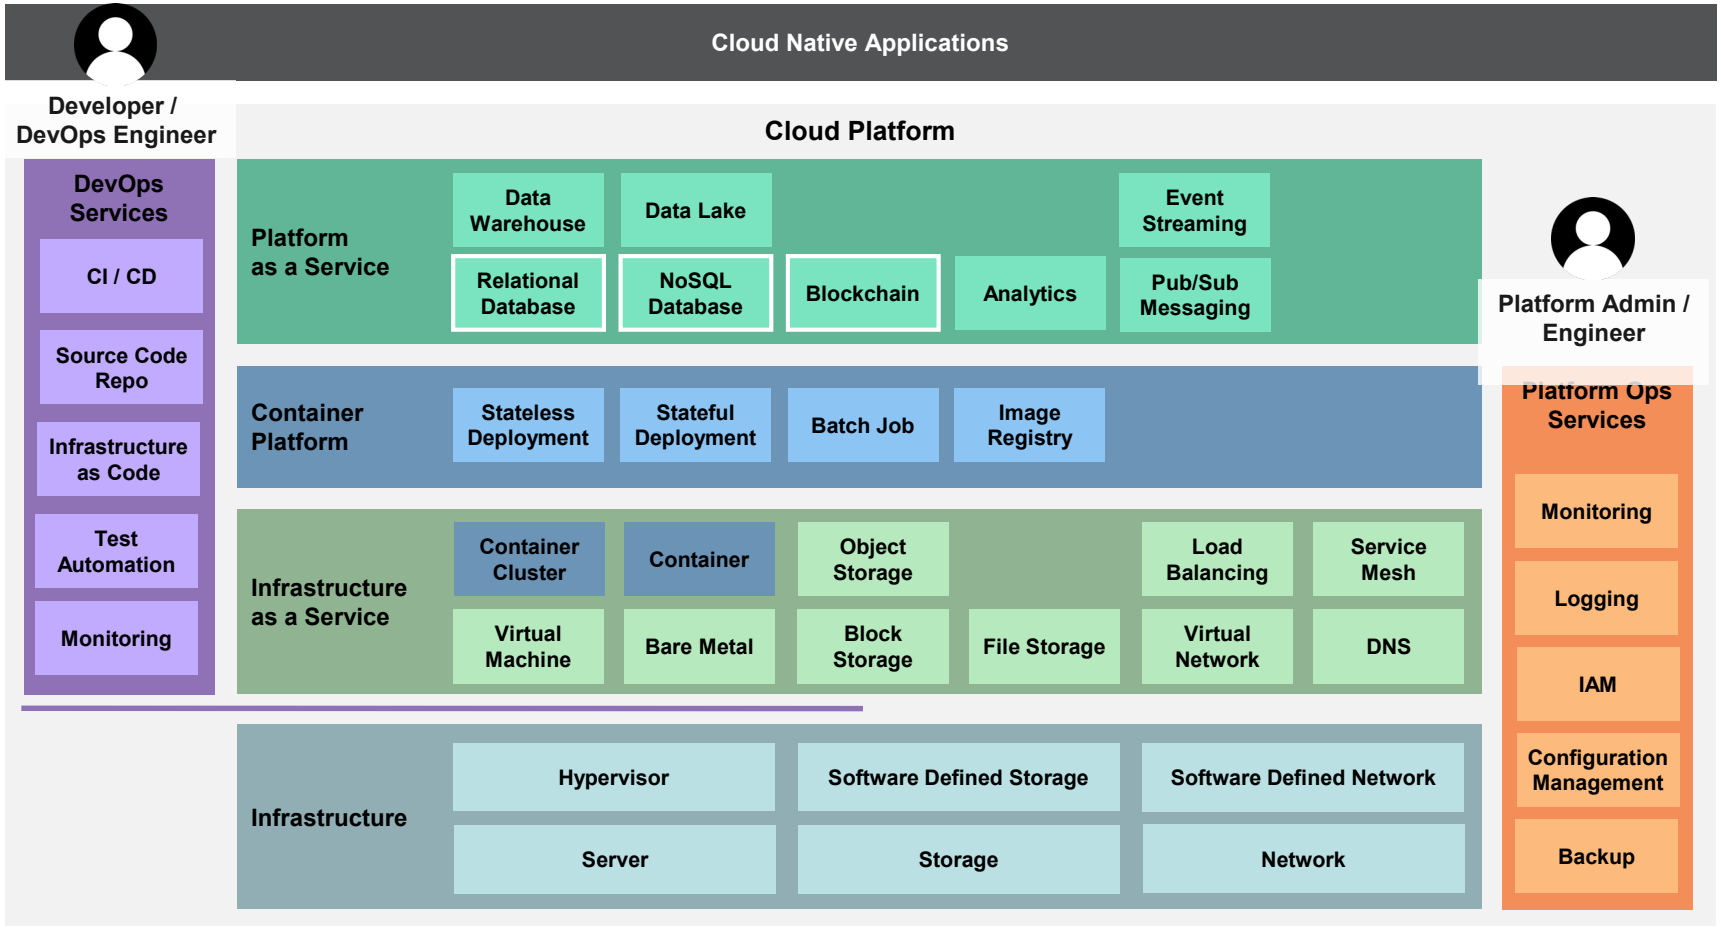

IDC Market Glance: Cloud Platform Landscape

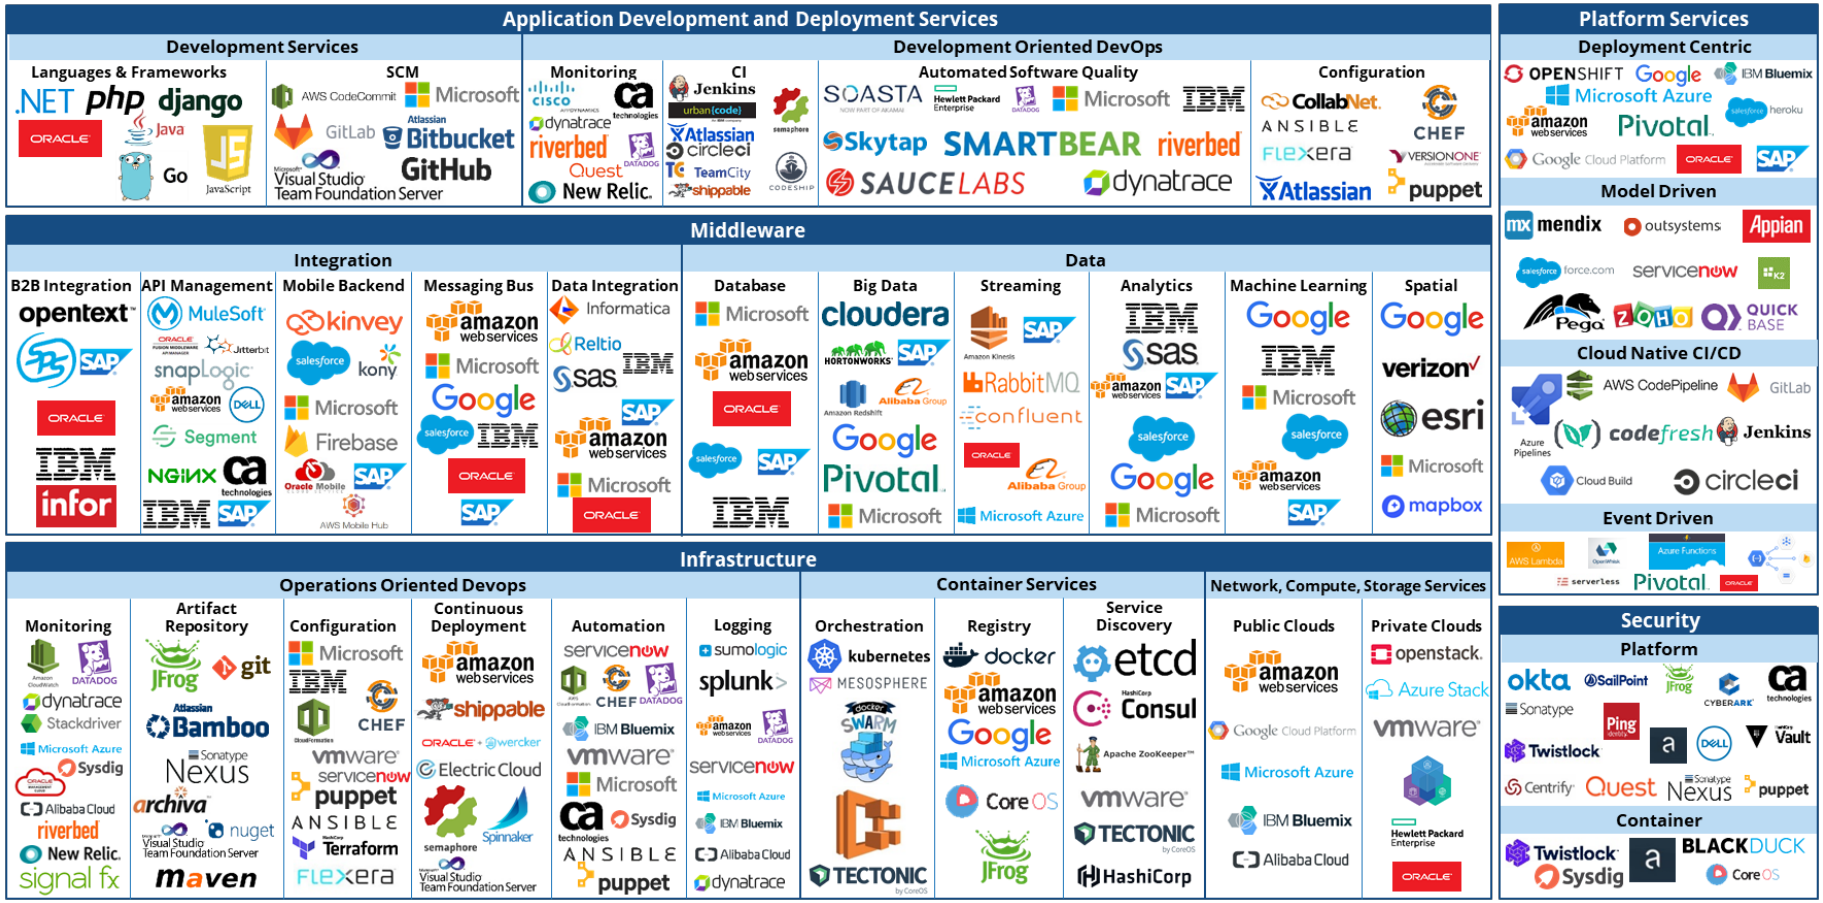

Source: IDC, 1Q19

For areas on which IDC publishes market share data, the top 3–5 market share leaders are represented. For areas on which IDC does not publish market share data, vendor selection is up to analyst discretion.

Cloud Native Ecosystem Partners

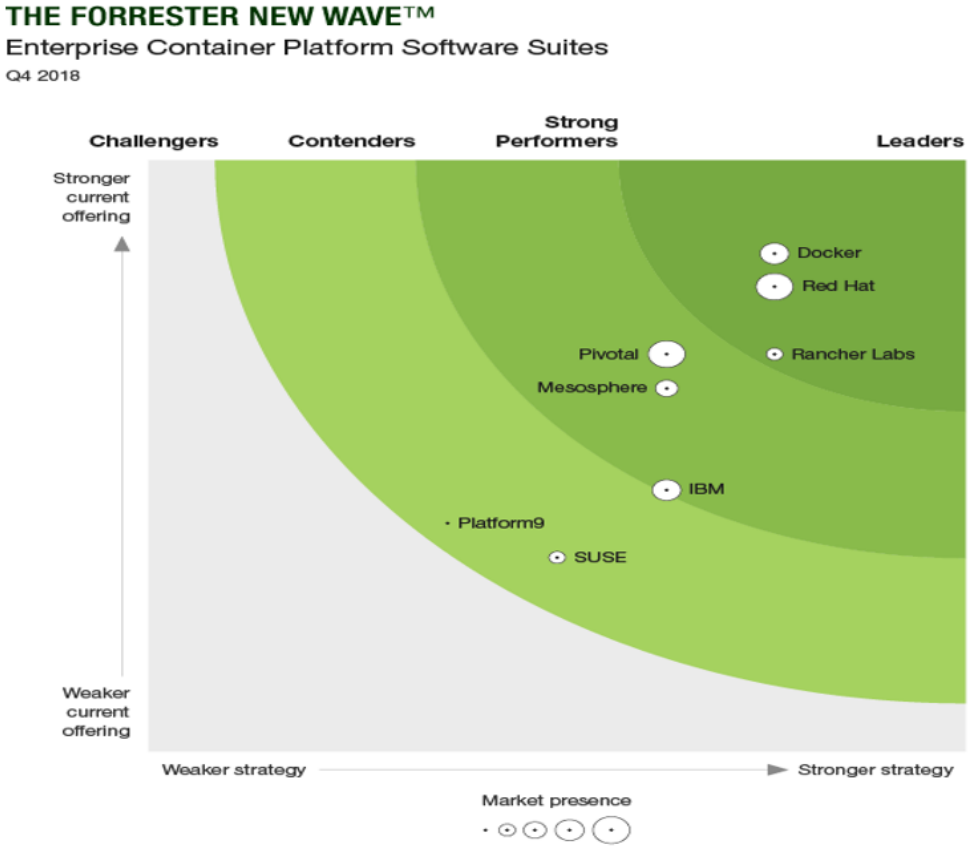

•https://reprints.forrester.com/#/assets/2/431/RES141562/reports

Company	Runtime and orchestration	Image management	Operations management	Security features	User experience	App life-cycle management	Integrations and APIs	Vision	Road map	Market approach
Docker	▲	▲	▲	▲	▲	=	▲	▲	=	▲
Red Hat	▲	▲	=	=	▲	▲	▲	▲	▲	=
Rancher Labs	▲	▲	=	▲	=	=	=	▲	▲	=
Pivotal	=	▲	▲	▲	=	▲	▼	▲	=	=
Mesosphere	▲	=	▲	=	=	▼	▲	▲	=	=
IBM	▼	=	=	▲	=	=	▼	=	▲	=
SUSE	=	=	▼	=	▼	▼	=	=	=	=
Platform9	=	▼	▲	▼	=	▼	=	=	=	▼

▲ Differentiated = On par ▼ Needs improvement

141562 Source: Forrester Research, Inc. Unauthorized reproduction, citation, or distribution prohibited.

https://www.redhat.com/en/blog/red-hat-among-vendors-lead-pack-forrester-new-wave-enterprise-container-platform-software

Linux on Z/LinuxONE Cloud Platform Enablement Status

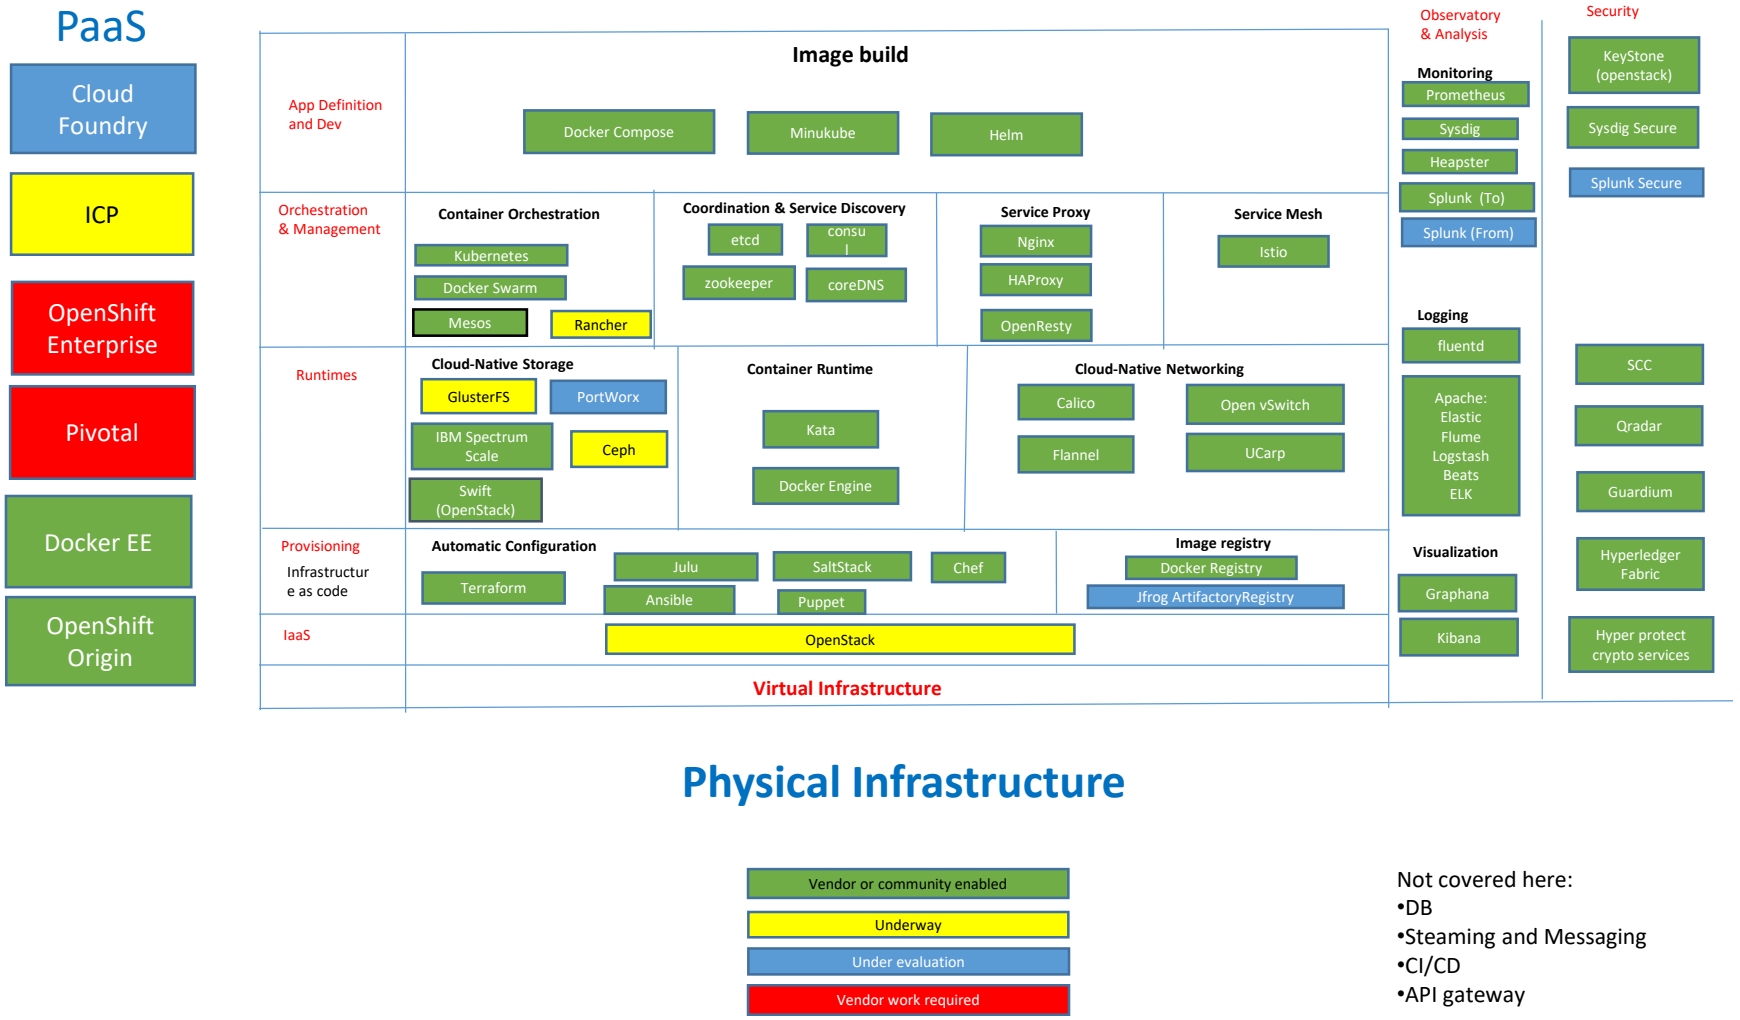

Physical Infrastructure

Vendor or community enabled

Underway

Under evaluation

Vendor work required

Not covered here:
- DB
- Steaming and Messaging
- CI/CD
- API gateway

Cloud Native Ecosystem Focus Areas

Align the content with “converged infrastructure”

- CI/CD
- Monitoring tools
- Security tools
- IaaS
- Software defined storage, network, etc

Continue to enlist ecosystem partners

Red Hat Cloud Stack

Data Serving Ecosystem

Linux on IBM Z/LinuxONE for Open-Source Databases

The most scalable, reliable and securable hardware and the most popular open-source database management system

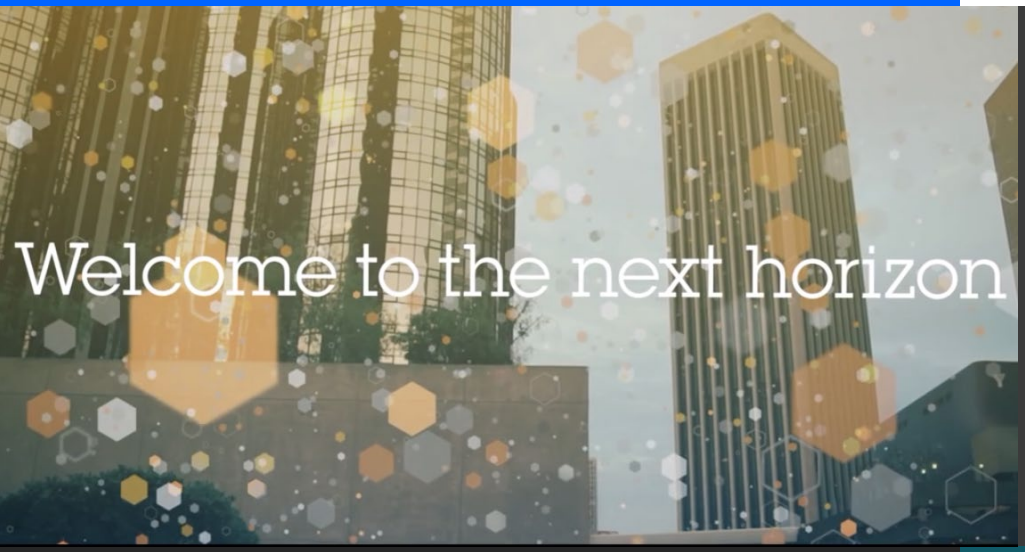

OSDBs are preferred
for more than half of the database types:

- Graph DBMS
- Search engines
- Key-value stores
- Document stores
- Time series DBMS
- Wide column stores

You now have as many choices for OSDBs as for commercial DBs
(170 vs. 173)

Over the past six years, OSDBs have taken 13% market share away from commercial DBs

DB-Engines.com top 5 open source systems
January 2019
All of them run on Linux on Z/LinuxONE

Rank	System	Score
1.	MySQL	1154
2.	PostgreSQL	466
3.	MongoDB	387
4.	Redis	149
5.	Elasticsearch	143

Data-serving challenges: disruption caused by Big Data

IDC and EMC projected
By 2020 data will grow to 40 zettabytes;
50-fold growth from 2010

Scalability
Server Sprawl
Availability
Ensuring Privacy

In Jan 2017 ... **27,000 MongoDB databases were stolen! 25% of the 100,000 deployed MongoDB databases in the world have been compromised.**

Why Linux on Z/LinuxONE for Open Source Database?

Why Linux on Z/LinuxONE for Open Source Database

Security

- ✓Pervasive encryption across the entire system
- ✓Protected keys
  - Safer than open-key
  - **Up to 28x faster** than secure-key
- ✓Isolation on IBM Z/LinuxONE starts in the hardware ... *"What happens in a VM stays in a VM"*

Why Linux on Z/LinuxONE for Open Source Database

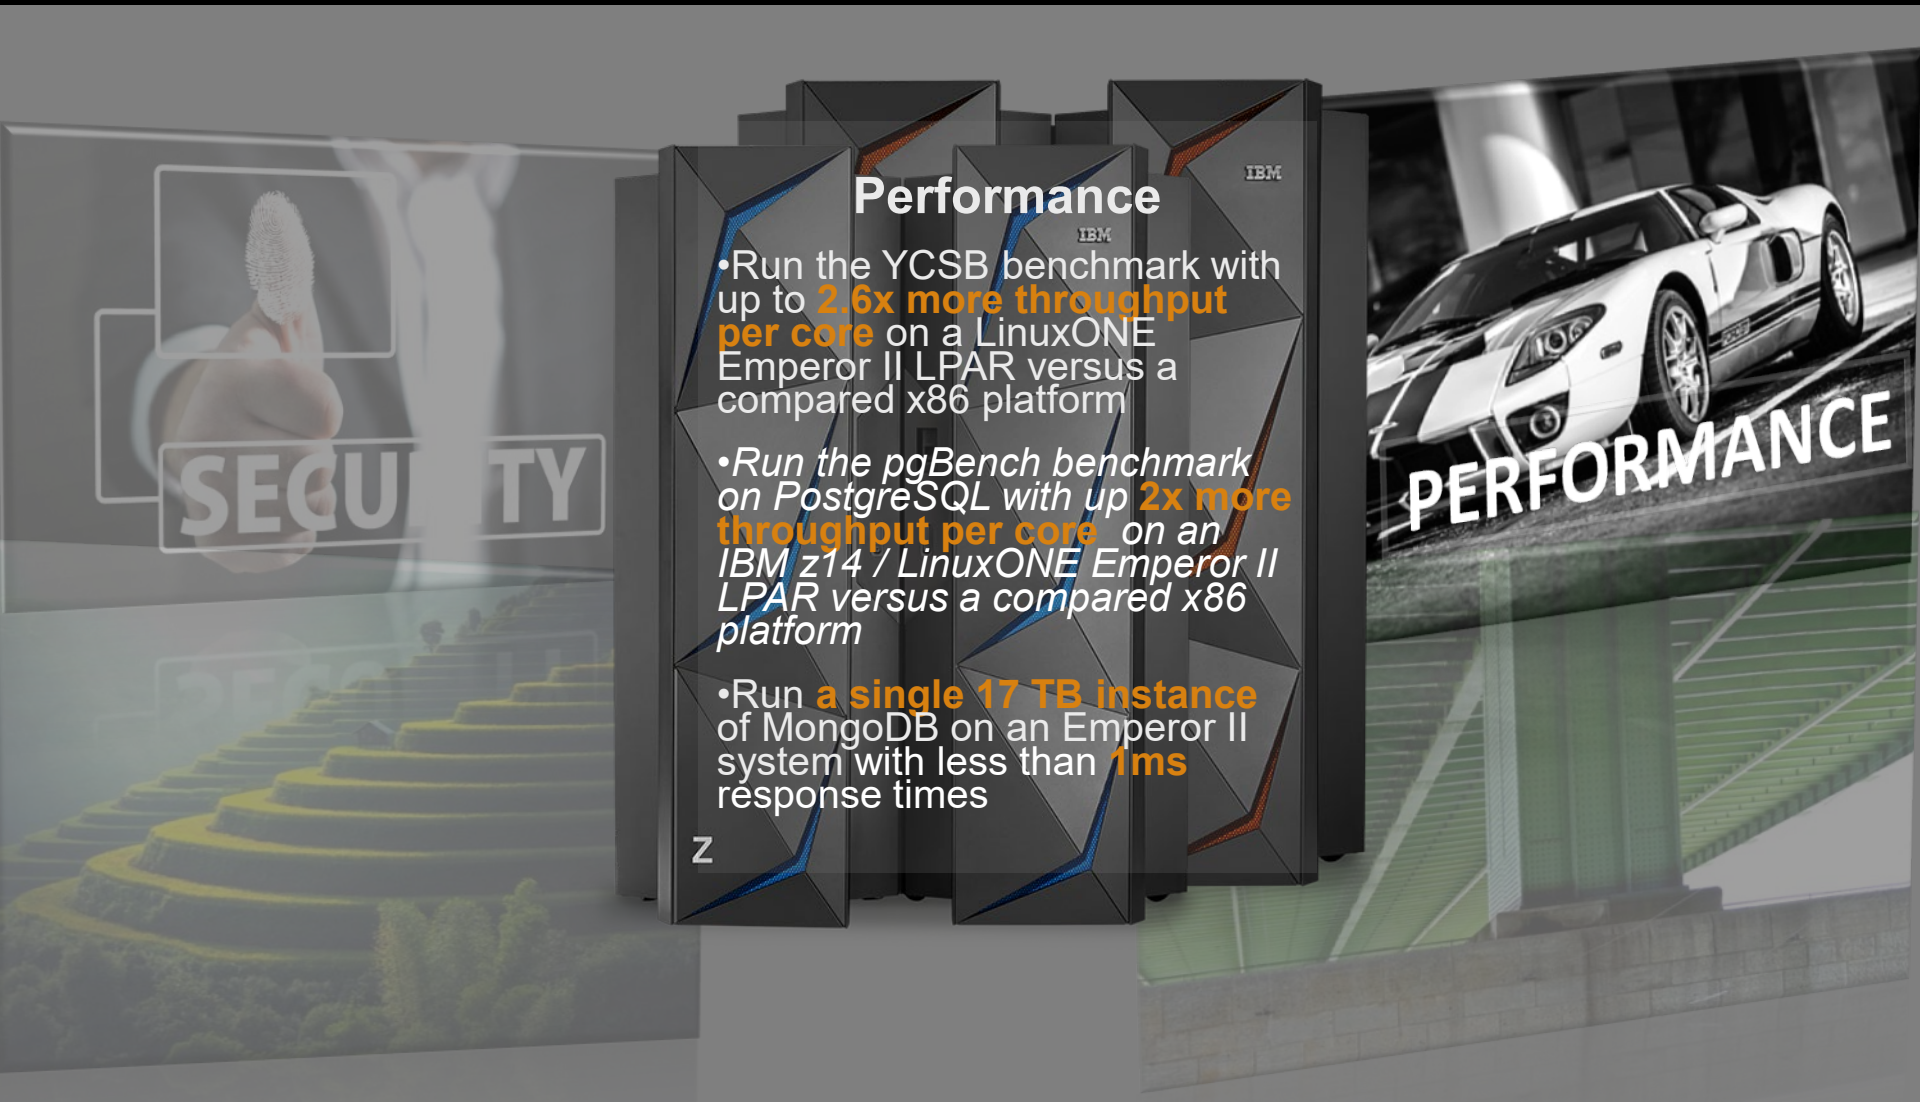

Performance

- Run the YCSB benchmark with up to **2.6x more throughput per core** on a LinuxONE Emperor II LPAR versus a compared x86 platform
- *Run the pgBench benchmark on PostgreSQL with up* ***2x more throughput per core*** *on an IBM z14 / LinuxONE Emperor II LPAR versus a compared x86 platform*
- Run **a single 17 TB instance** of MongoDB on an Emperor II system with less than **1ms** response times

Why Linux on Z/LinuxONE for Open Source Database

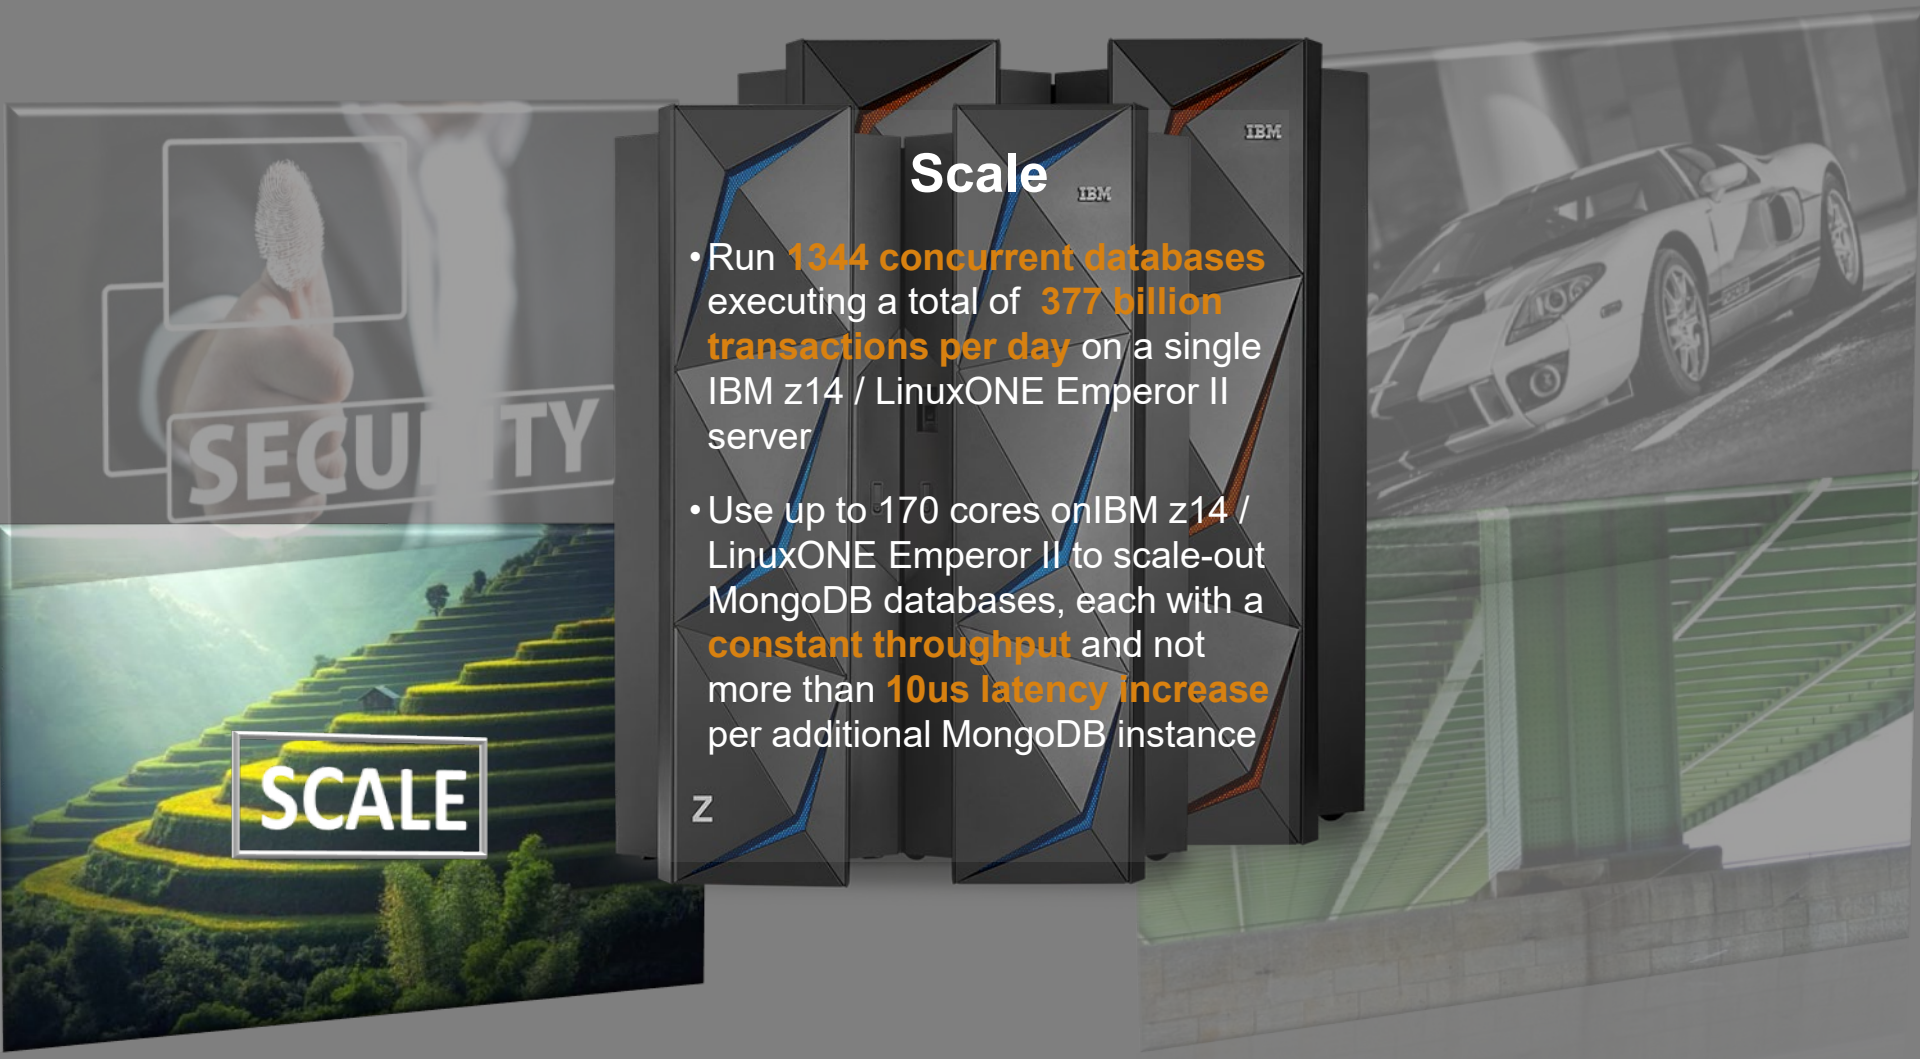

Scale

- Run **1344 concurrent databases** executing a total of **377 billion transactions per day** on a single IBM z14 / LinuxONE Emperor II server
- Use up to 170 cores onIBM z14 / LinuxONE Emperor II to scale-out MongoDB databases, each with a **constant throughput** and not more than **10us latency increase** per additional MongoDB instance

Why Linux on Z/LinuxONE for OSDBs?

*2017 ITIC survey shows IBM Z / LinuxONE had **no unplanned system downtime** due to inherent flaws in hardware*

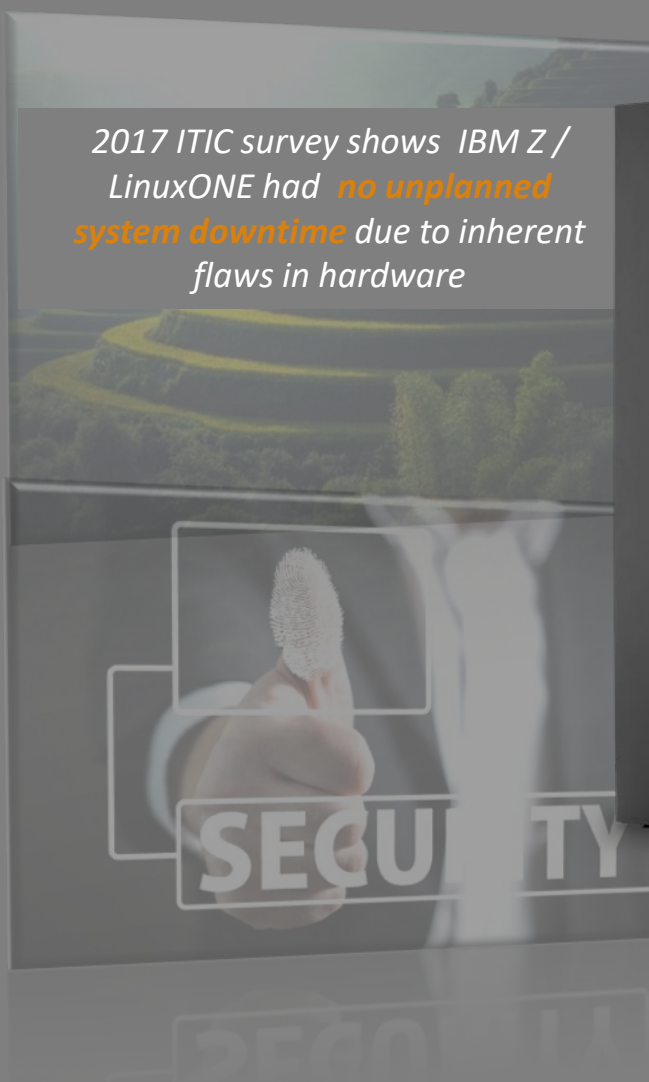

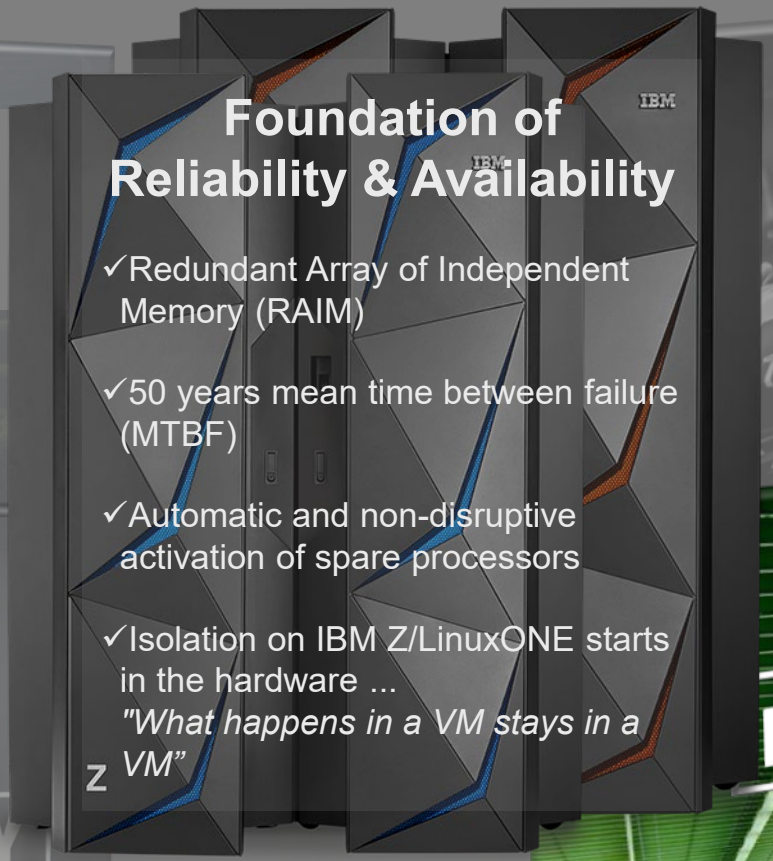

Foundation of Reliability & Availability

- ✓Redundant Array of Independent Memory (RAIM)
- ✓50 years mean time between failure (MTBF)
- ✓Automatic and non-disruptive activation of spare processors
- ✓Isolation on IBM Z/LinuxONE starts in the hardware ... *"What happens in a VM stays in a VM"*

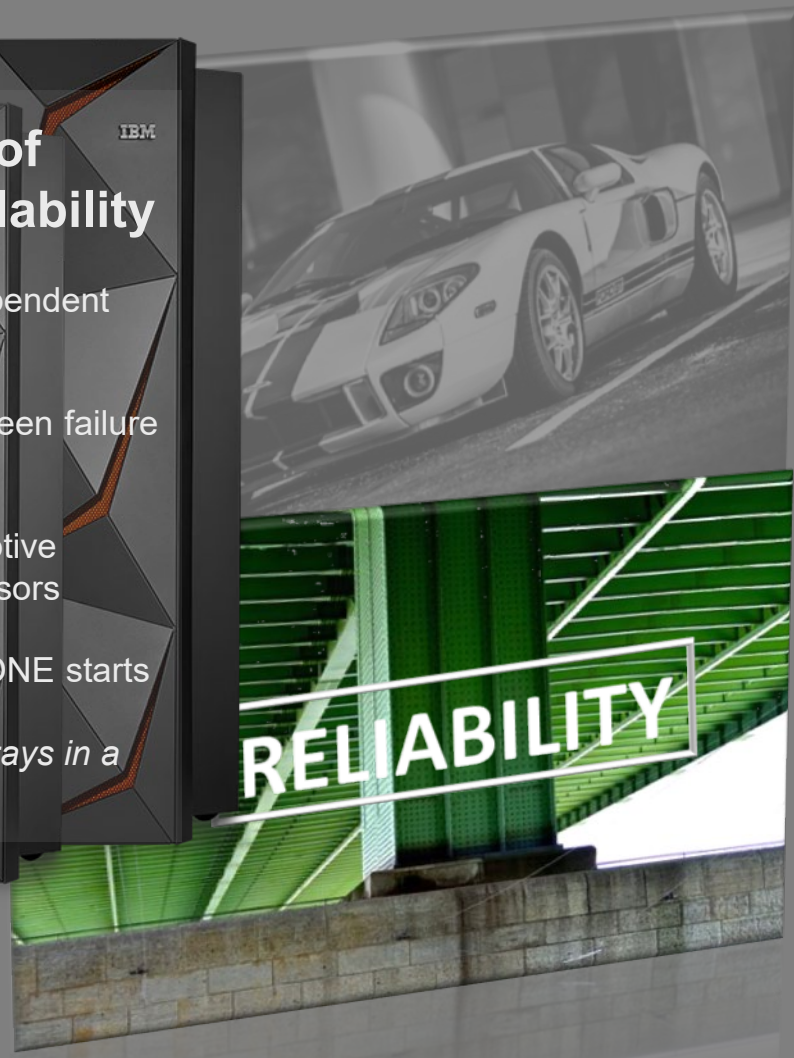

Data Serving Ecosystem Focus Areas

Assist clients with open source database adoption

- Growing interest in MongoDB, PostGreSQL, and MySQL (MariaDB)
- Provide expertise in configuration and tuning
- Liaison with partners

Next frontiers

- Identify other popular database trend and use cases
- Explore partnership with ScyllaDB, ElasticSearch, Redis

Analytics and Machine Learning

Analytics Ecosystem Strategy

- IBM analytics products: DSX, ICP4D
- Open source: Spark, Solr, ELK, etc
- Partnership with commercial solutions (Splunk, AppDynamic, Circinus, etc)

ML/AI Ecosystem Strategy

Foundation

Priority enable ML libraries and languages on Z

Content
TensorFlow, Anaconda, Keras, python, R, PyPy, Eigen, OpenBlas, etc

Use case and why Z

Use Case
Inference

Value Prop
- colocation with System of Record
- ML in the Z security Z, computation on encrypted data

ISVs and PoC

ISVs enaged
- SparkBeyond
- Fatbrain

PoC
- fraud detection with a credit card company

Blockchain Ecosystem

2019 Enterprise Projections

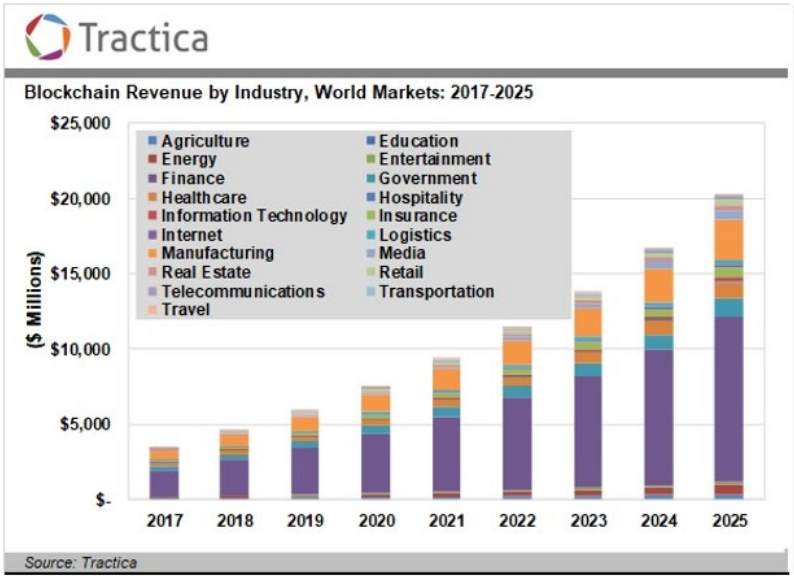

Blockchain Revenue by Industry

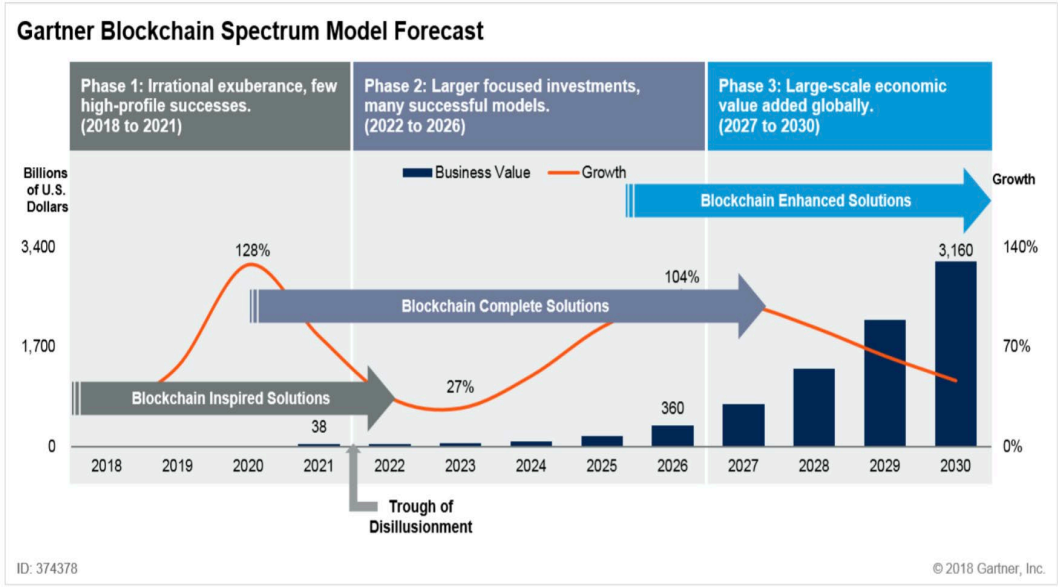

Blockchain Spectrum Model Forecast

Key Takeaway: Finance is projected to be the biggest sector for Blockchain adoption

Market POV detail

Blockchain provider ecosystem

ILLUSTRATIVE LISTS, NOT COMPREHENSIVE

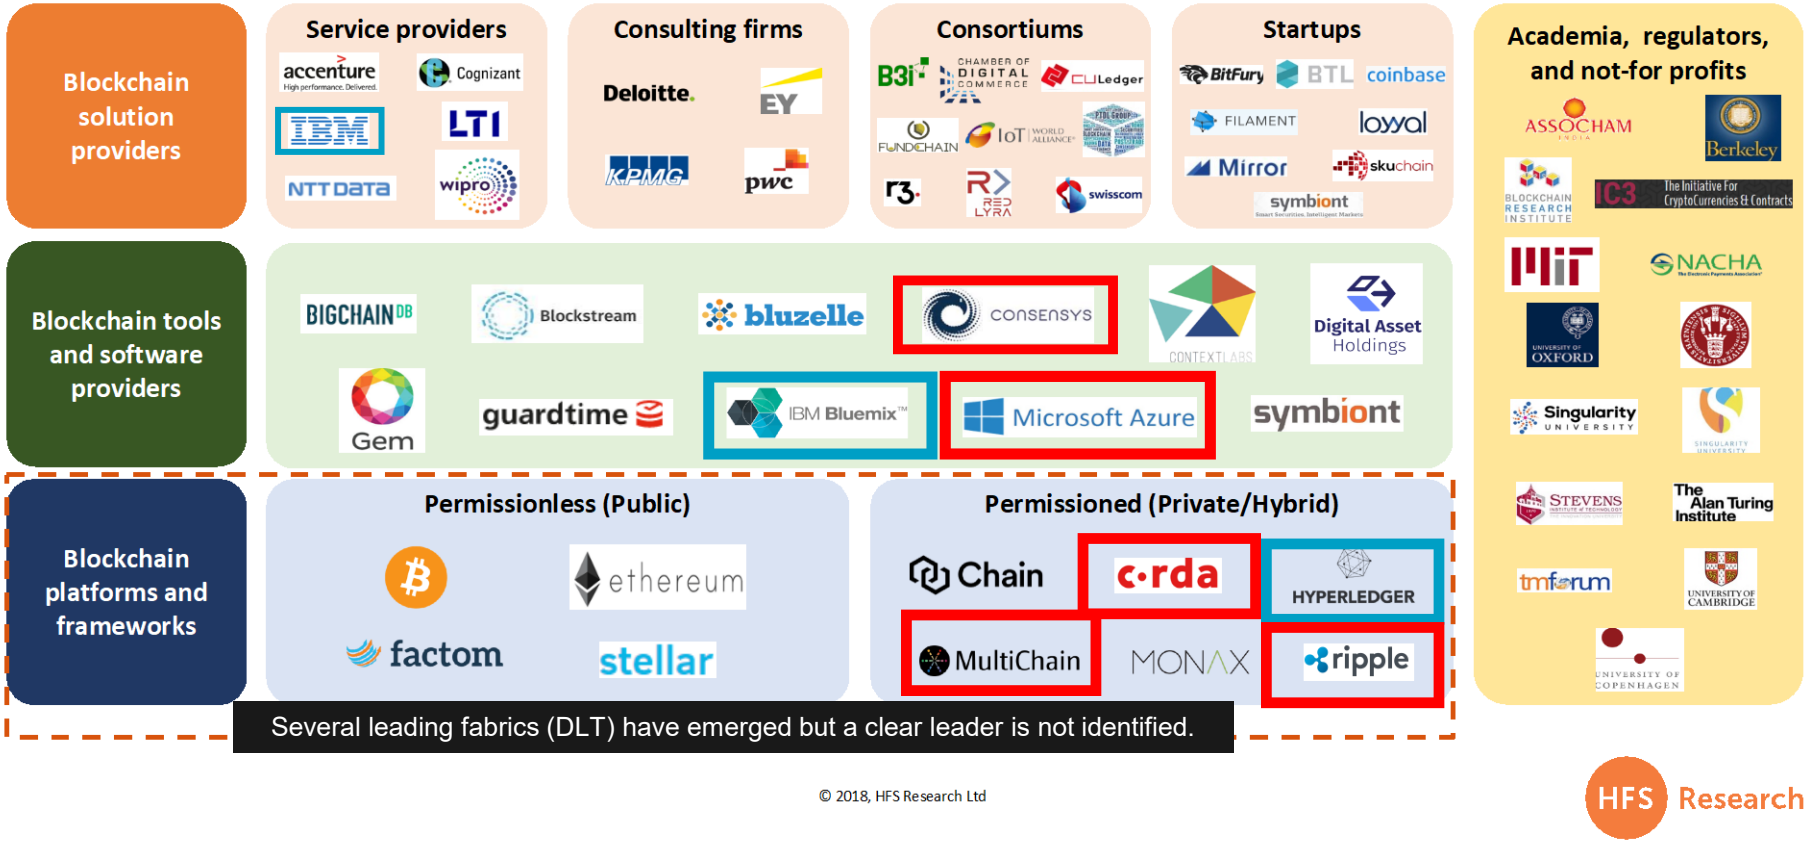

HFS Research

IBM Blockchain Offerings

IBM Blockchain Platform (Hyperledger)

- provides an easy way to deploy blockchain solutions
- deploy option: IBM Cloud (including Hyper Protect), AWS, and on-prem with ICP
- target client who wish to development their own blockchain network
- many different use cases
- business network can be built on Z

IBM Blockchain Solution (HyperLedger) running on IBM Cloud, optionally Hyper Protect Services (Z)

Food Trust

- food supply management network

World Wire

- interbank cross-border payments (With annual growth expected to average 7%, the payments industry is expected to be a $2-trillion-dollar industry by 2020.[1])

TradeLens

- Supply Chain use case

IBM BlockchainTrusted Identity

- Identity network (1,339 breaches impacting over 170 million records in the U.S. in 2017)
  - secure, digital wallet that holds cryptographically secured credentials. You control what information is shared and who can access it.
  - Blockchain facilitate a secure, trusted exchange of information between parties who do not typically trust each other.

Blockchain Ecosystem Strategy on Linux on Z/LinuxONE

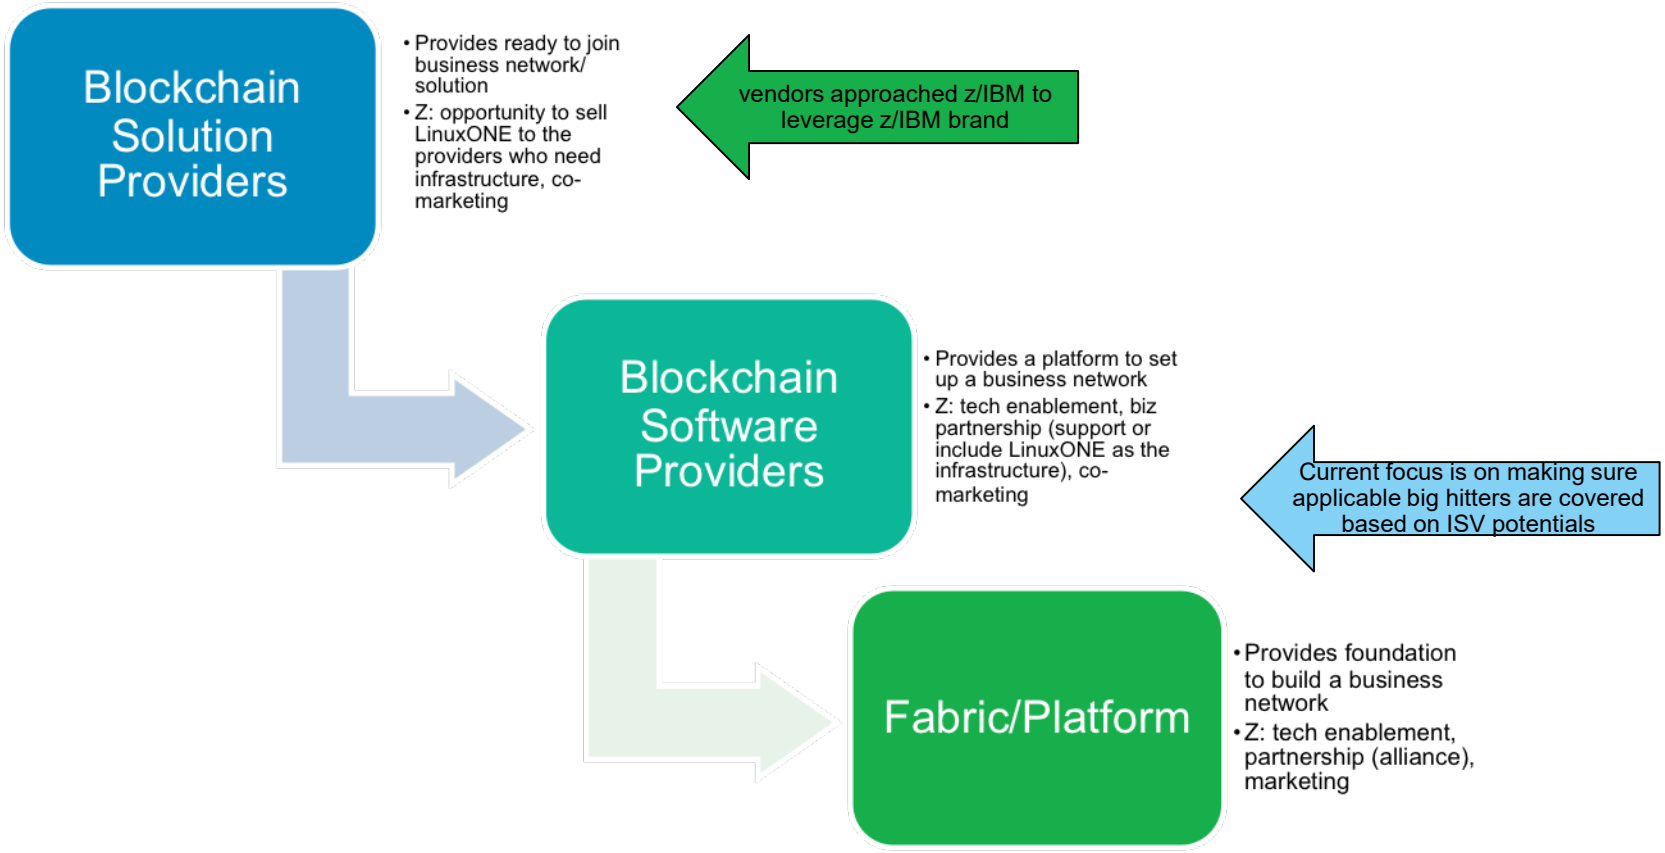

Blockchain Ecosystem Candidates

Evaluate/Enable Blockchain Fabric and Platform

- Ethereum
- Corda
- Multichain
- Quorum

Explore partnership with Vendors

- Primechain/Infibeam
- EdgeVerve/Infosys
- Geoblue Lab
- Billon/FIS

Ported and Validated on Linux on Z/LinuxONE as of Dec 2018

Up-to-date list: http://ibm.biz/LoZOpenSourceVerifiedList

Backup

Open Source & ISV Linux SW Capability

Ported - verified
Work in progress (target quarter/half)
Work under Evaluation/not started

Languages and Dev Environment	Databases	Messaging & Streaming
Node.js	MySQL	RabbitMQ
Ruby	PostgreSQL	Apache Kafka
Rails	MariaDB (w MaxScale)	Logstash (ELK)
Python, PyPy Python JIT	MongoDB	Fluentd
LLVM	Redis	Apache Flume
OpenJDK, OpenJDK 9 (w JIT)	Apache CouchDB	Apache ActiveMQ
GCCGO, Golang compiler	Cloudant^ (CouchDB based)	Apache Camel
oCaml, oCaml native compiler	Apache Geode	Graylog2*
Erlang	CouchBase	Apache Apex(Data Torrent)*
PHP/Zend	Apache HBase	Apache NiFi*
R	ScyllaDB	IronMQ*
Clojure	RethinkDB	
Scala	XtraDB*	**Graph DBs**
Swift (Apple)	OrientDB*	Spark GraphX
BIRT (Eclipse plug-in)	Hazelcast*	Neo4j
Mono(C#...) [open source for .NET]	MemSQL*	Pegasus*
GOLD(alternate Linux linker)	Aerospike*	Titan*
ATLAS BLAS(z13 SIMD Vector Extension)	Druid*	
LuaJIT	Apache Ignite*	
JIT* for PHP, Dart*		

* Pending input on priority
^ IBM offering

** Content and priority are subject to change*

Various sources of input: e.g. BlueMix, Github stats, feedback from: direct client input, IBM client reps, on going research

Open Source & ISV Linux SW Capability

Ported - verified
Work in progress (target quarter/half)
Work under Evaluation/not started

Orchestration/ Deployment	Deployment Management & Config - Monitoring	Machine & Deep Learning	Operating Systems
Docker	Chef	Spark MLLib	SLES <enterprise>
Docker Swarm	Puppet	Word2vec	RHEL <enterprise>
Docker Compose	Ansible	TensorFlow	Ubuntu <enterprise>
Kubernetes	virt-install	EigenBLAS(z13 SIMD Vector Extension)	Debian
LXD	Consul	SystemML*	Fedora
Apache Mesos	cAdvisor	H2O*	ClefOS (CentOS)
Marathon	Etcd	Neon*	OpenSUSE*
Chronos	Apache ZooKeeper	Apache SINGA*	CoreOS*
Mesosphere DC/OS	Sysdig	**FaaS, IaaS & PaaS**	RancherOS*
SaltStack	Heapster	OpenStack	**Networking(SDN, DNS)**
Exechealthz	Prometheus	Cloud Foundry local	CoreDNS
Flannel	Service-catalog	OpenShift	Onos
kube-dnsmasq	Terraform	Juju	StatsD
Calico	Zenoss*	Deis	UCARP
Rex-Ray	ElasticBox*	Apache OpenWhisk*	OpenDaylight*
Helm	Icinga 2*		Open vSwitch*
Istio	DataDog*		
Minikube			
Apache Aurora*			

* Pending input on priority
^ IBM offering

** Content and priority are subject to change*

Various sources of input: e.g. BlueMix, Github stats, feedback from: direct client input, IBM client reps, on going research

Open Source & ISV Linux SW Capability

Ported - verified
Work in progress (target quarter/half)
Work under Evaluation/not started

Big Data & Analytics	App development & DevOps	Web Application Dev/Perf & CMS	Front End
Elasticsearch (ELK)	Xerces-c	Apache jMeter	Kibana (ELK)
Apache Spark	XMLSec	Wordpress	D3
Apache Solr	protobuf	Ceilometer	Angular (MEAN)
Apache Storm	Doxygen	Apache Tomcat	Express (MEAN)
Anaconda	ANTLR	HAProxy	Beats
	Apache Maven	NGNIX	Grafana
	Jenkins	Apache HTTP(utils)	Ionic*
Notebooks	Apigility	JBoss	Graphene*
Apache Zeppelin	Mule	Drupal	Meteor*
Jupyter*	PM2	Joomla	
Ipython*	Htop	SugarCRM	
	IPTraf	Magento	
	Hibernate	Alfresco	
Distributed File System	SonarQube	OpenResty	
GlusterFS	Akka	TomEE*	
Heketi	Graphite*		
Ceph*			

* Pending input on priority
^ IBM offering

** Content and priority are subject to change*

Various sources of input: e.g. BlueMix, Github stats, feedback from: direct client input, IBM client reps, on going research