

Data Server Day



Adobe Stock | #228959181


Les King
Director, Data Solutions
lking@ca.ibm.com
September 2023

Future Technology

Promise



Unintended Consequences



Future Technology – Social Media



Future Technology – Social Media - Promise

Promise:

Bring us together in ways we could not imagine

Promise Realized:

Communication Globalization

News Reach

Keeping in contact with family & friends

Visiting parts of the world without leaving your living room



Future Technology – Social Media – Unintended Consequences

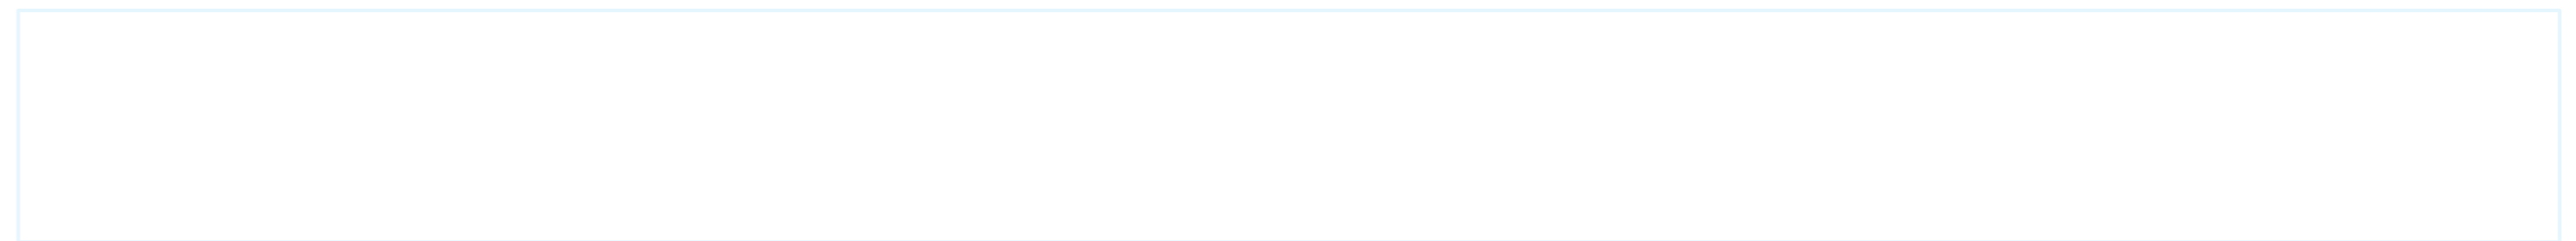
Are we more connected ?

OR

Just more connected to our devices ?



Future Technology – Social Media – Unintended Consequences



Future Technology – Social Media – Unintended Consequences




Future Technology – AI



AI - The Promise




AI - Unintended Consequences



AI - Unintended Consequences

When has the created ever loved, adored and respected the creator ?



Have you ever been a parent and the proud owner of a teenager ?

CHATGPT

 OpenAI




ChatGPT – The Promise




- Provides in-depth answers
- Strong analytics
- Translation
- Conversational

ChatGPT – Unintended Consequences - Bias




ChatGPT favors male names when de-biasing job feedback

How often ChatGPT removes negative personality feedback when asked to remove bias




ChatGPT – Unintended Consequences - Hallucinations



ChatGPT – Unintended Consequences – Adversarial Responses



+



=



90% Tabby Cat

Adversarial noise

100% Guacamole

Amazon Alexa – Blooper Countdown



Alexa starts a party and cops are called



Amazon Alexa – Blooper Countdown

Alexa starts a party and cops are called

Dollhouses and 2 kg of cookies
purchased by children




Amazon Alexa – Blooper Countdown

Alexa starts a party and cops are called

Dollhouses and 2 kg of cookies
purchased by children

Porn instead of children's song played
when "Digger Digger" requested by a child



Amazon Alexa – Blooper Countdown

Alexa starts a party and cops are called

Dollhouses and 2 kg of cookies
purchased by children

Porn instead of children's song played
when "Digger Digger" requested by a child

Bias is endless Passport, World Cup,
Beauty contest, Political ...



TM


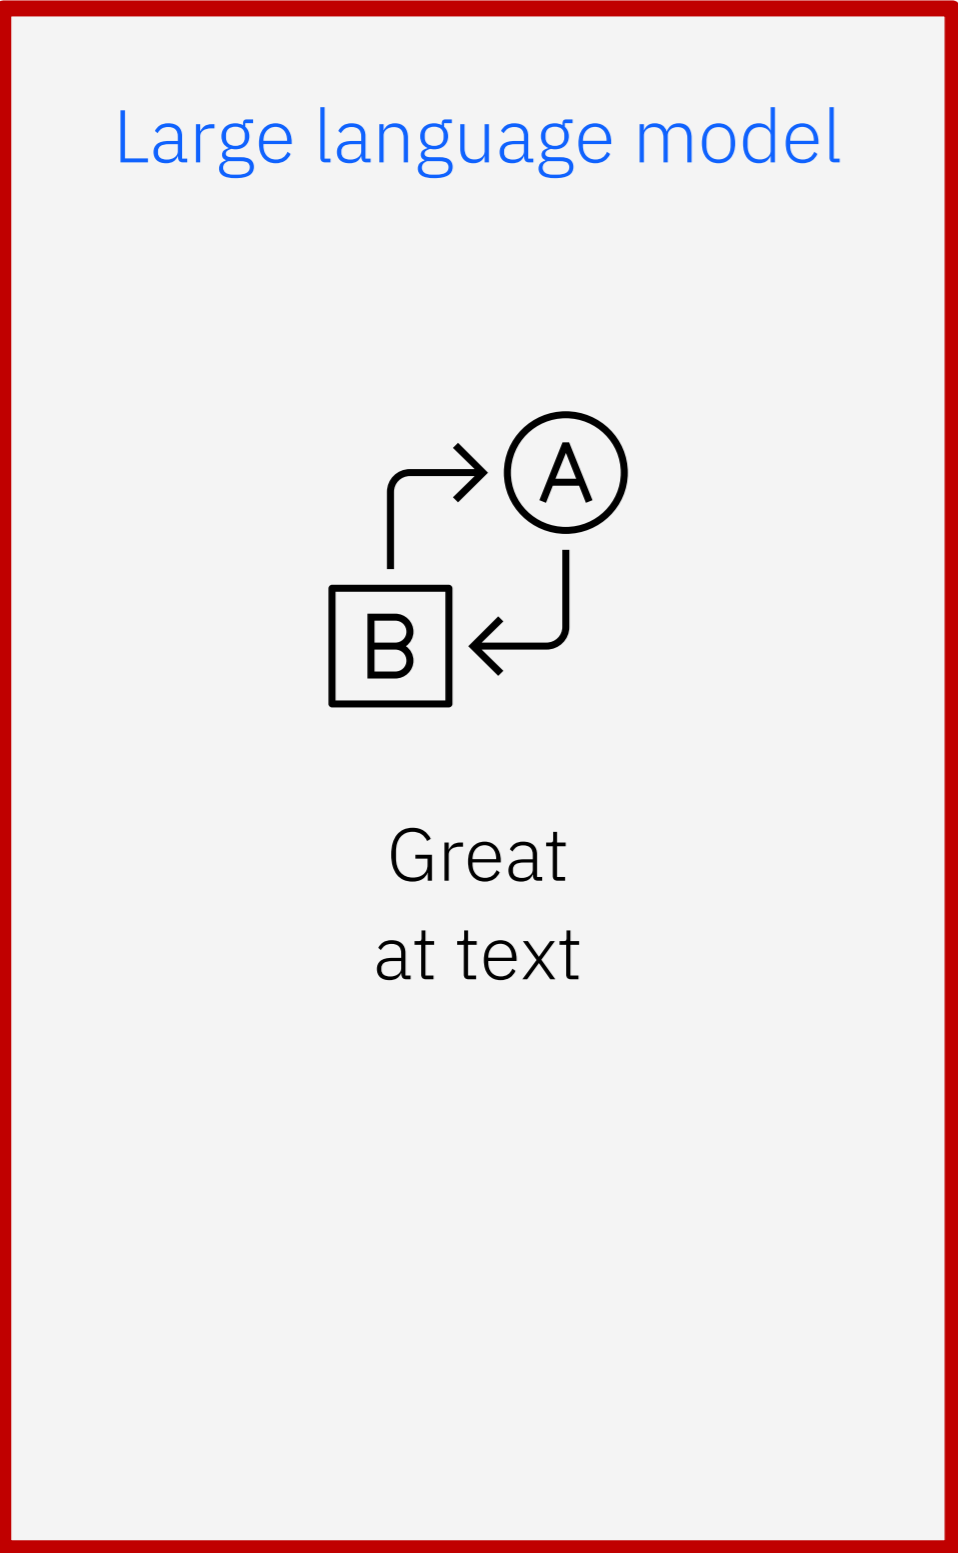
ChatGPT – Some Information





- ChatGPT is a large language model-based chatbot developed by OpenAI and launched in November 2022
- It enables users to refine and steer a conversation towards a desired length, format, style, level of detail, and language used
- Large Language Model (LLM)
- Generative, Pre-Trained Transformer (GPT)

Building blocks of an AI Strategy


ChatGPT
inspired interest...



But there is a
bigger concept...


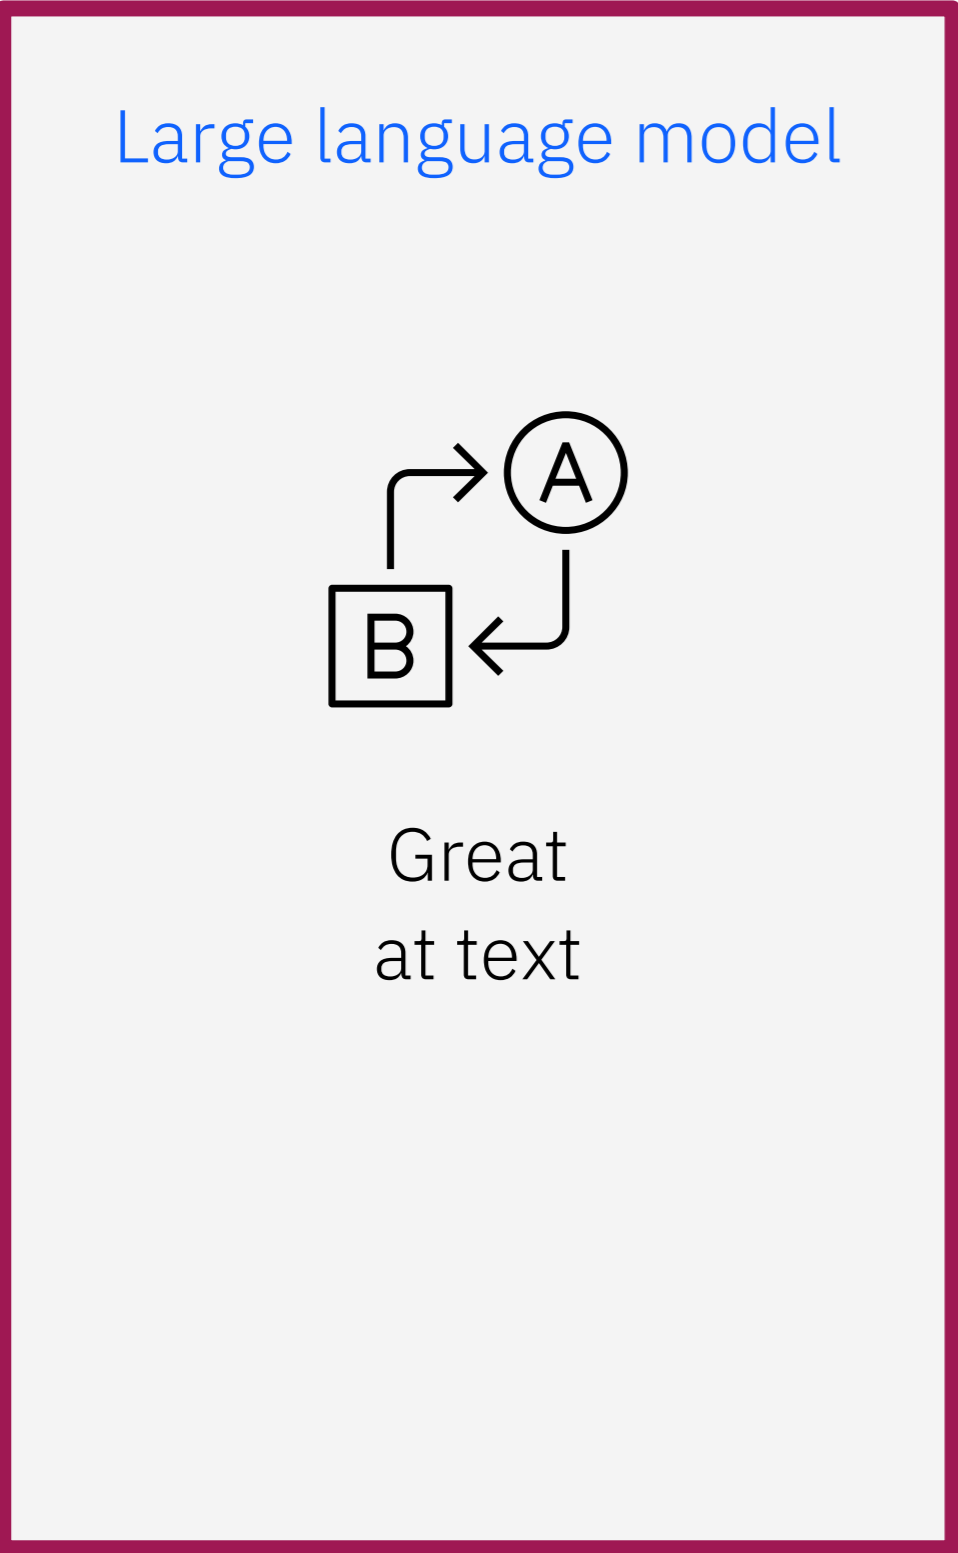


Which will
change business


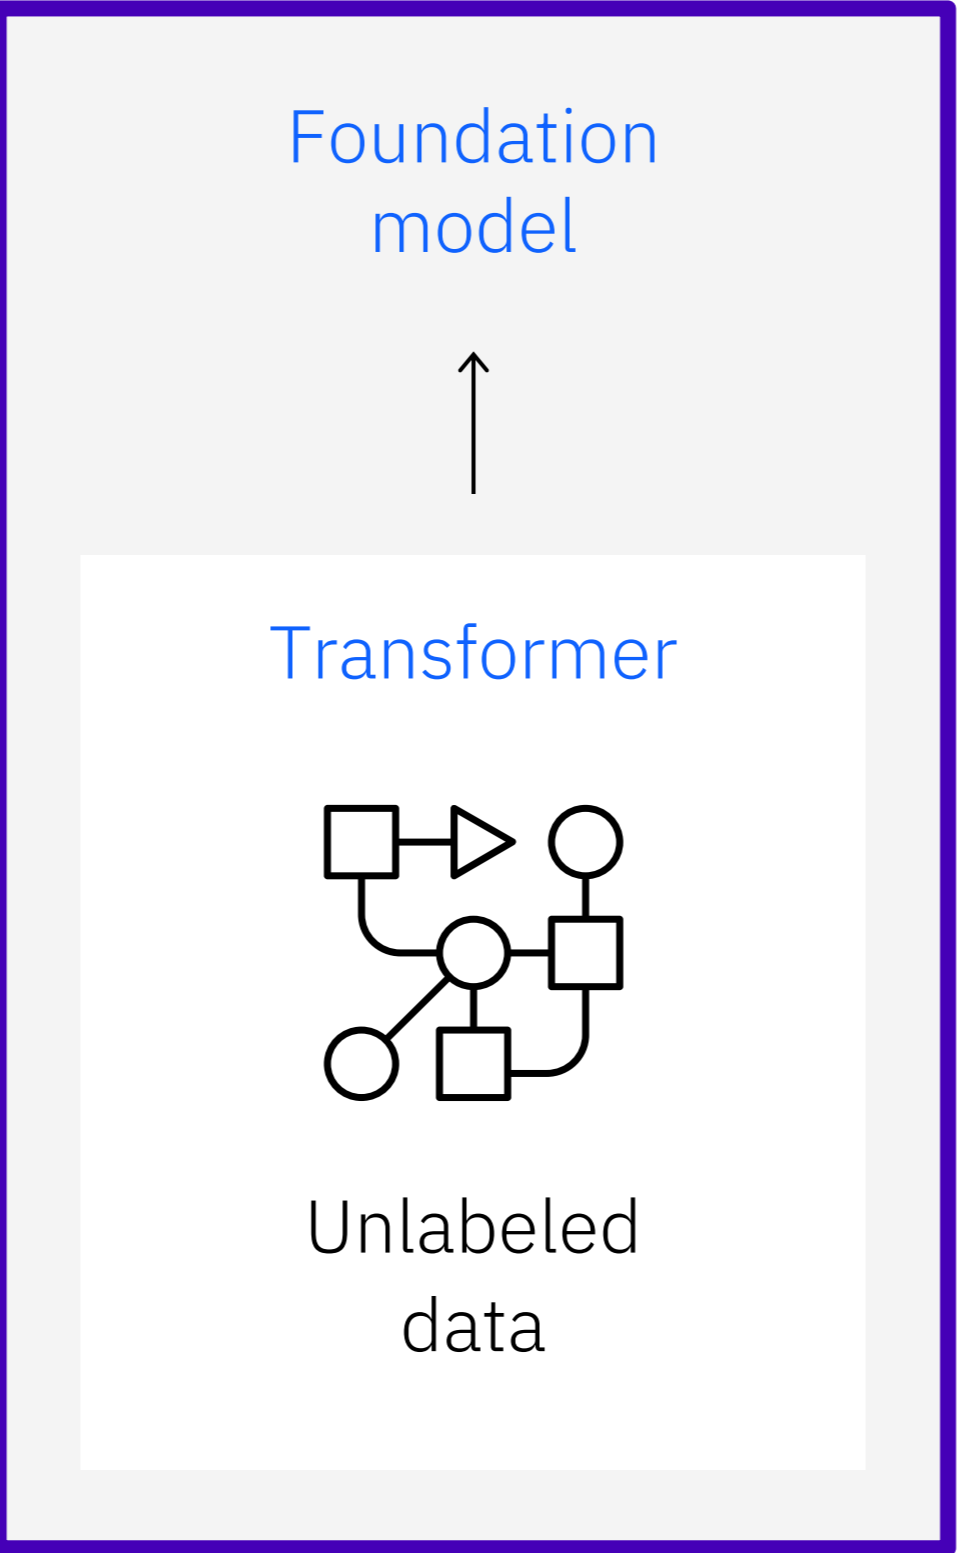


Building blocks of an AI Strategy

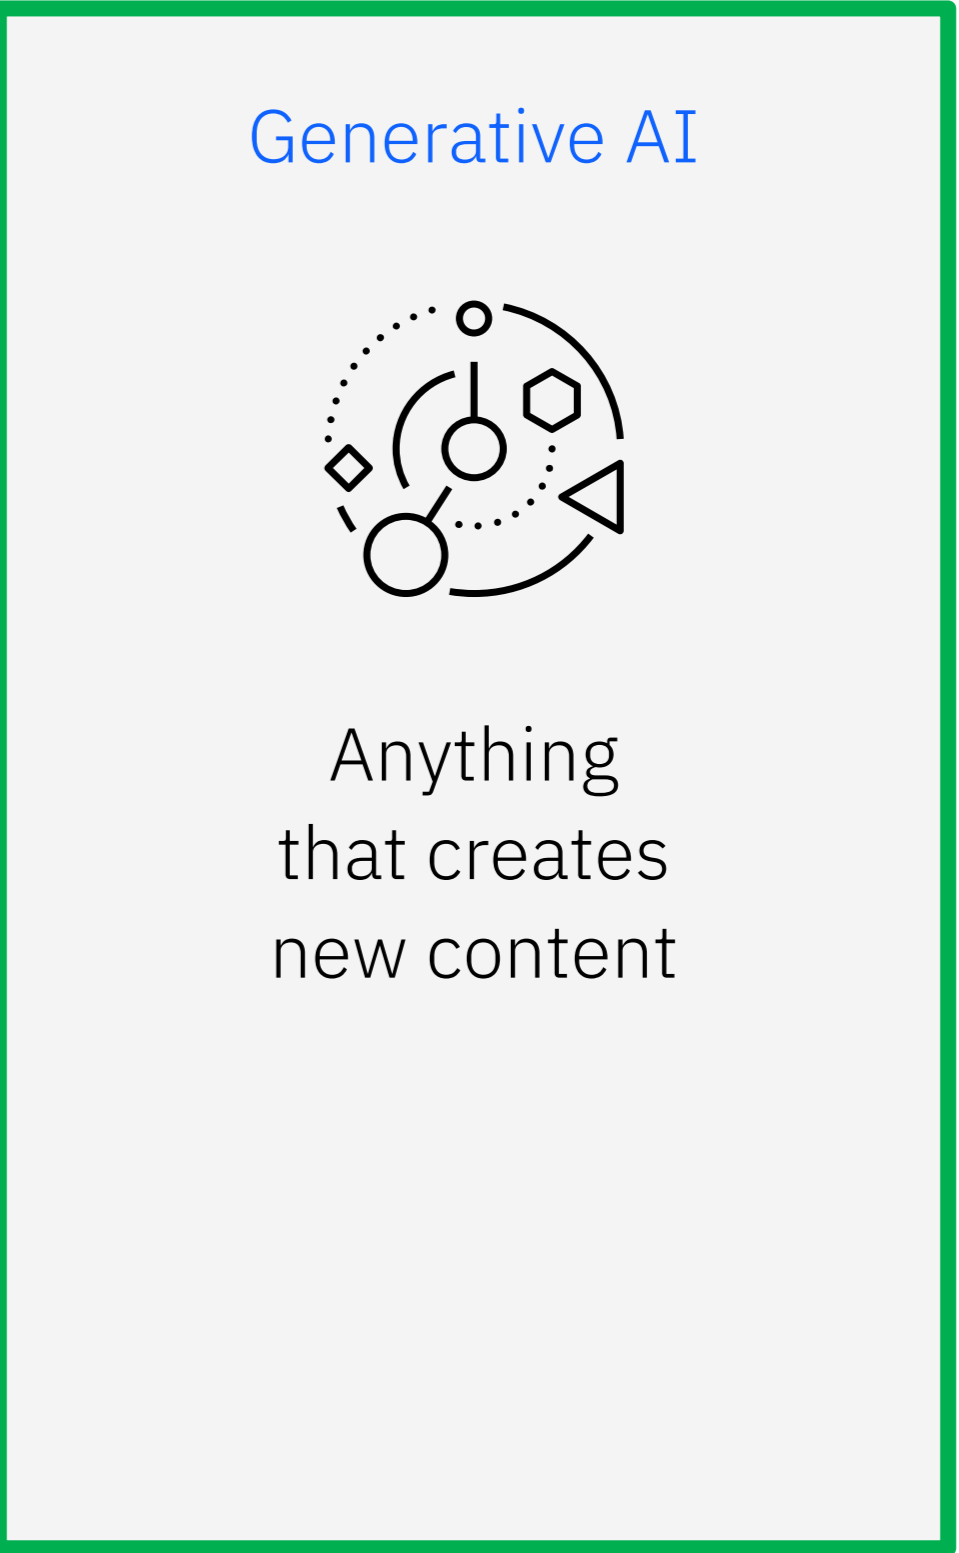
ChatGPT
inspired interest...



But there is a
bigger concept...



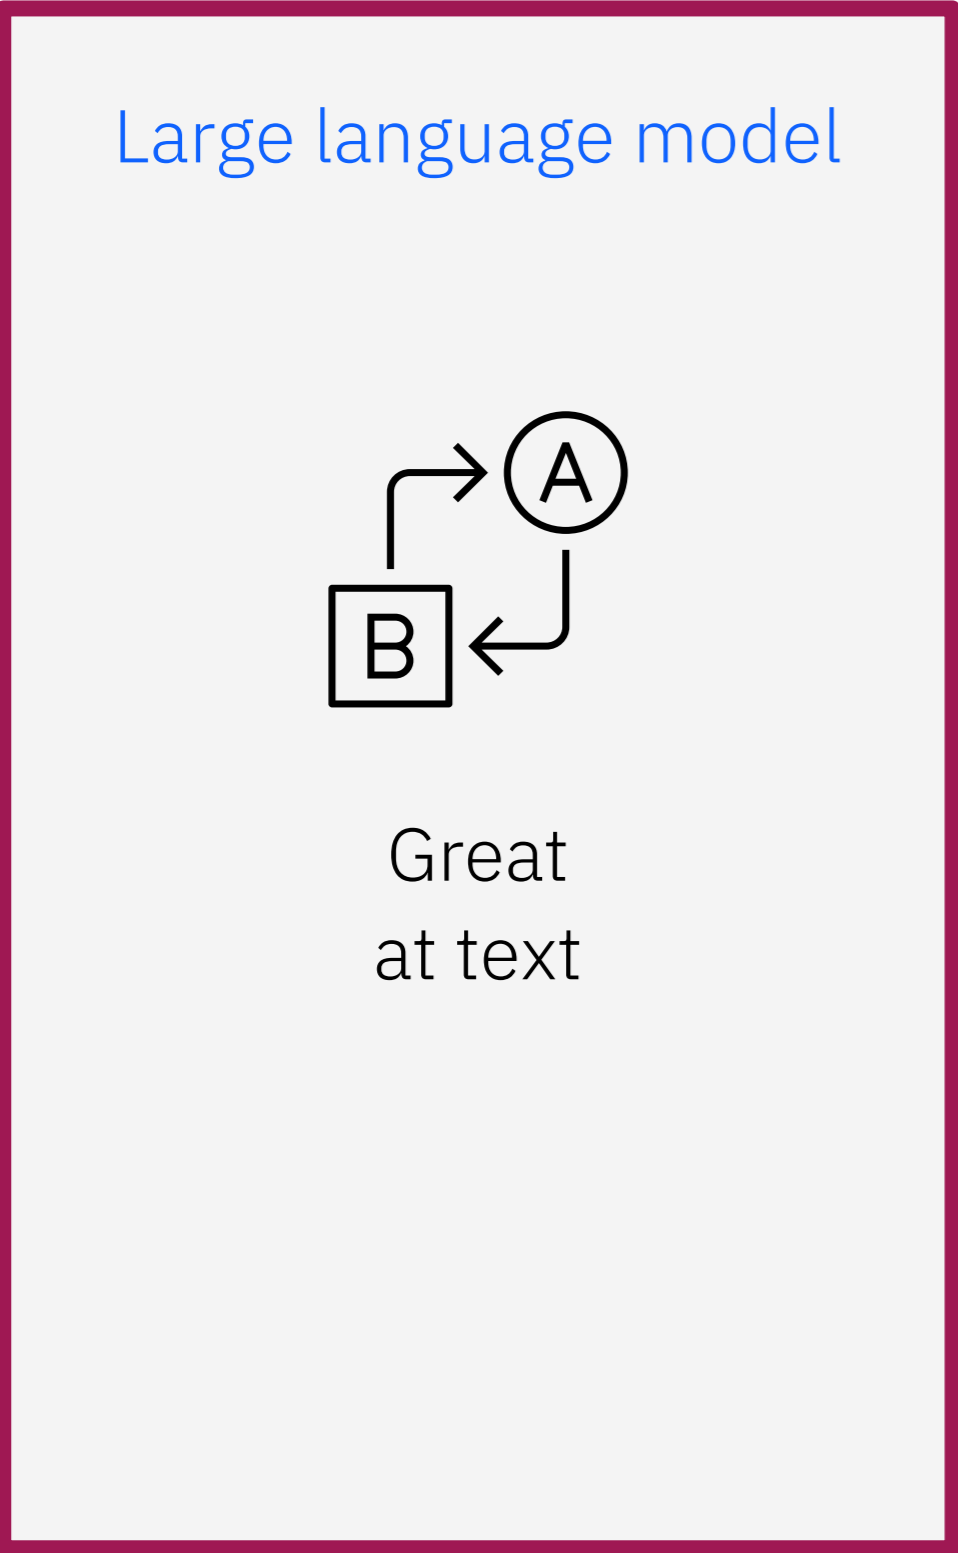
Which will
change business




A **large language model (LLM)** is a type of **machine learning model** that has been trained on **large quantities** of unlabeled text using self-supervised learning and can perform a variety of natural language processing (NLP) tasks (even when that language is a programming language). Output may range from books, articles, social media posts, online conversations, and even code. The architecture of an LLM consists of layers of **neural networks** that learn to generate language in a way that is similar to how humans use language

Building blocks of an AI Strategy


ChatGPT
inspired interest...



But there is a
bigger concept...



Which will
change business


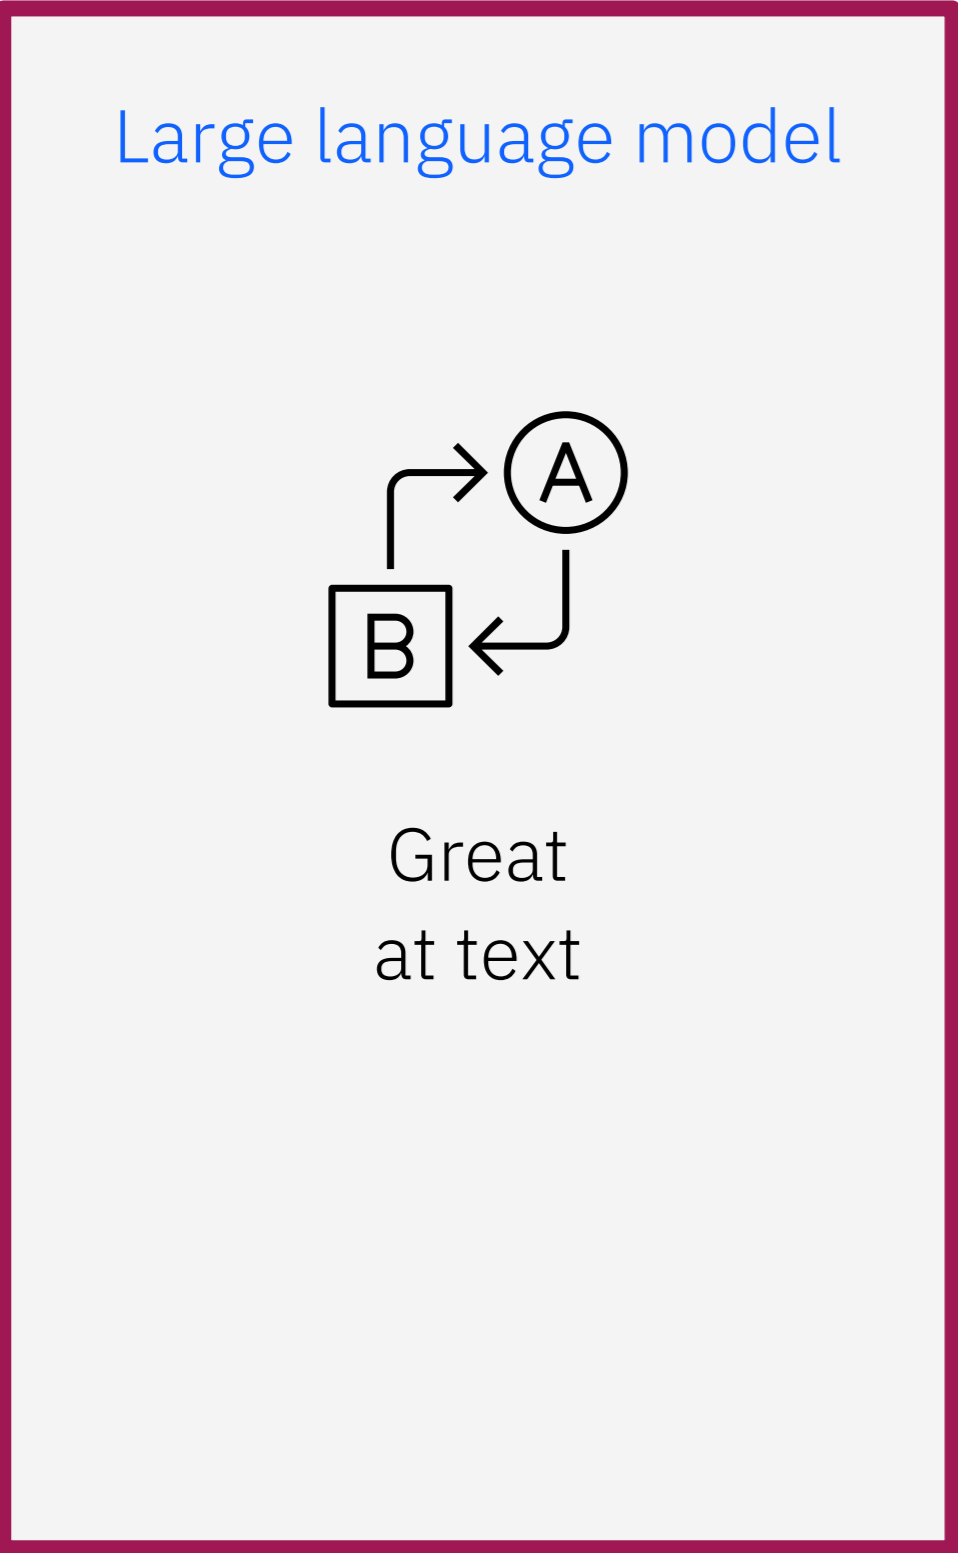


A **Foundation models** are typically built using a specific kind of neural network architecture, called a transformer, which is designed to generate sequences of related data elements (for example, like a sentence).


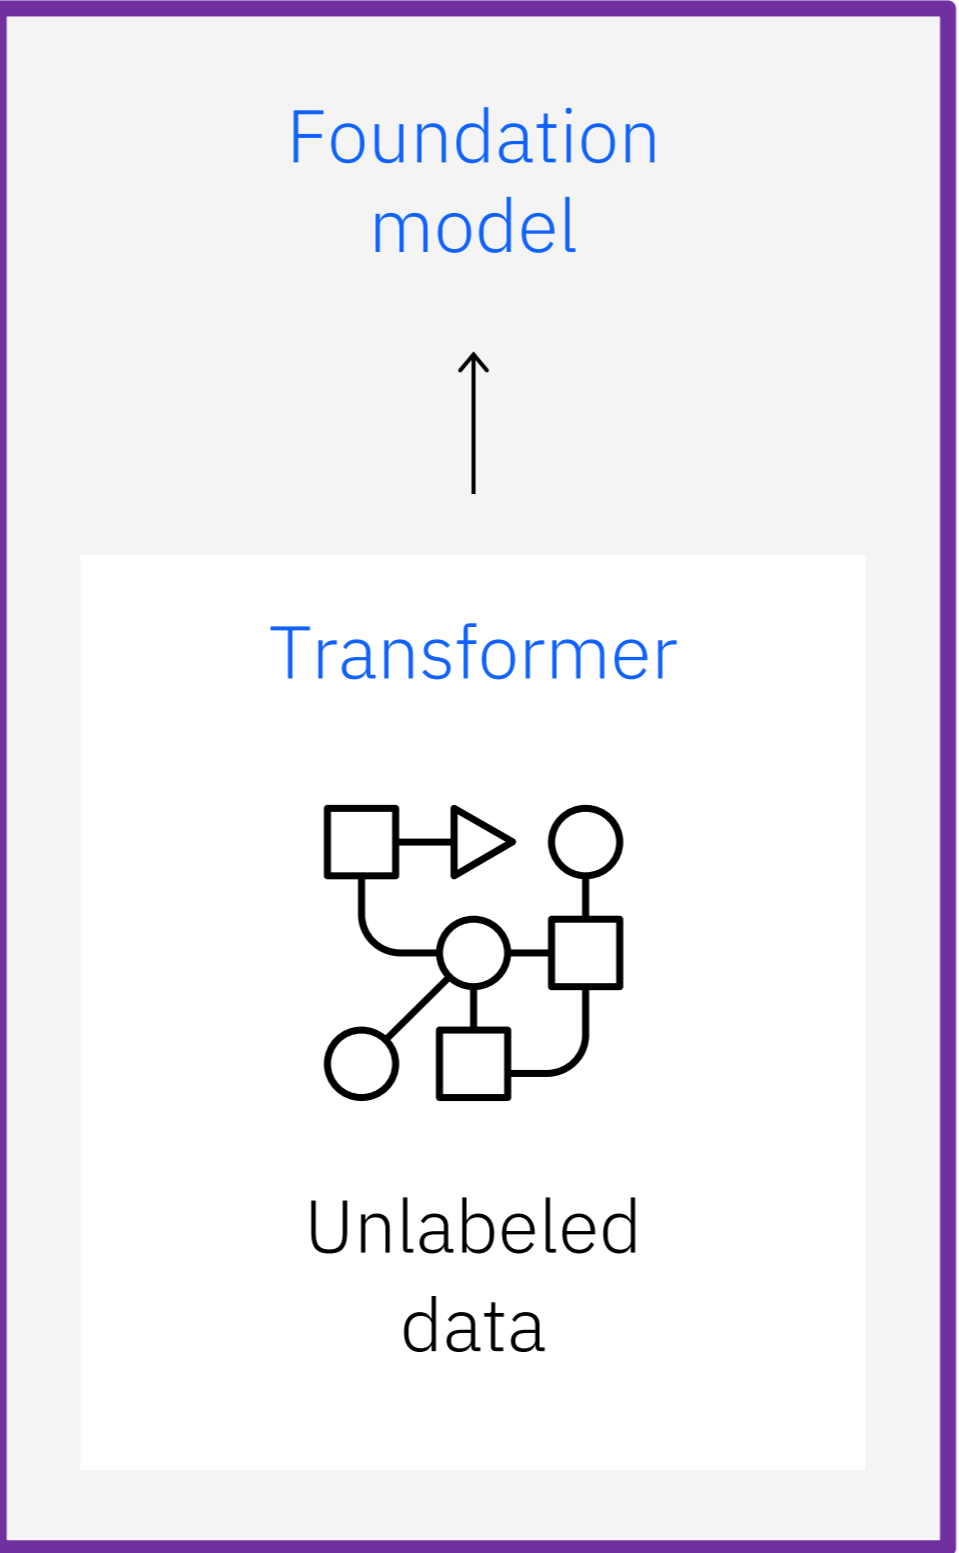
A **transformer model** is a neural network architecture useful for understanding language, which does not have to understand words one at a time but can look at an entire sentence at once for context and disambiguation.

Building blocks of an AI Strategy


ChatGPT
inspired interest...



But there is a
bigger concept...

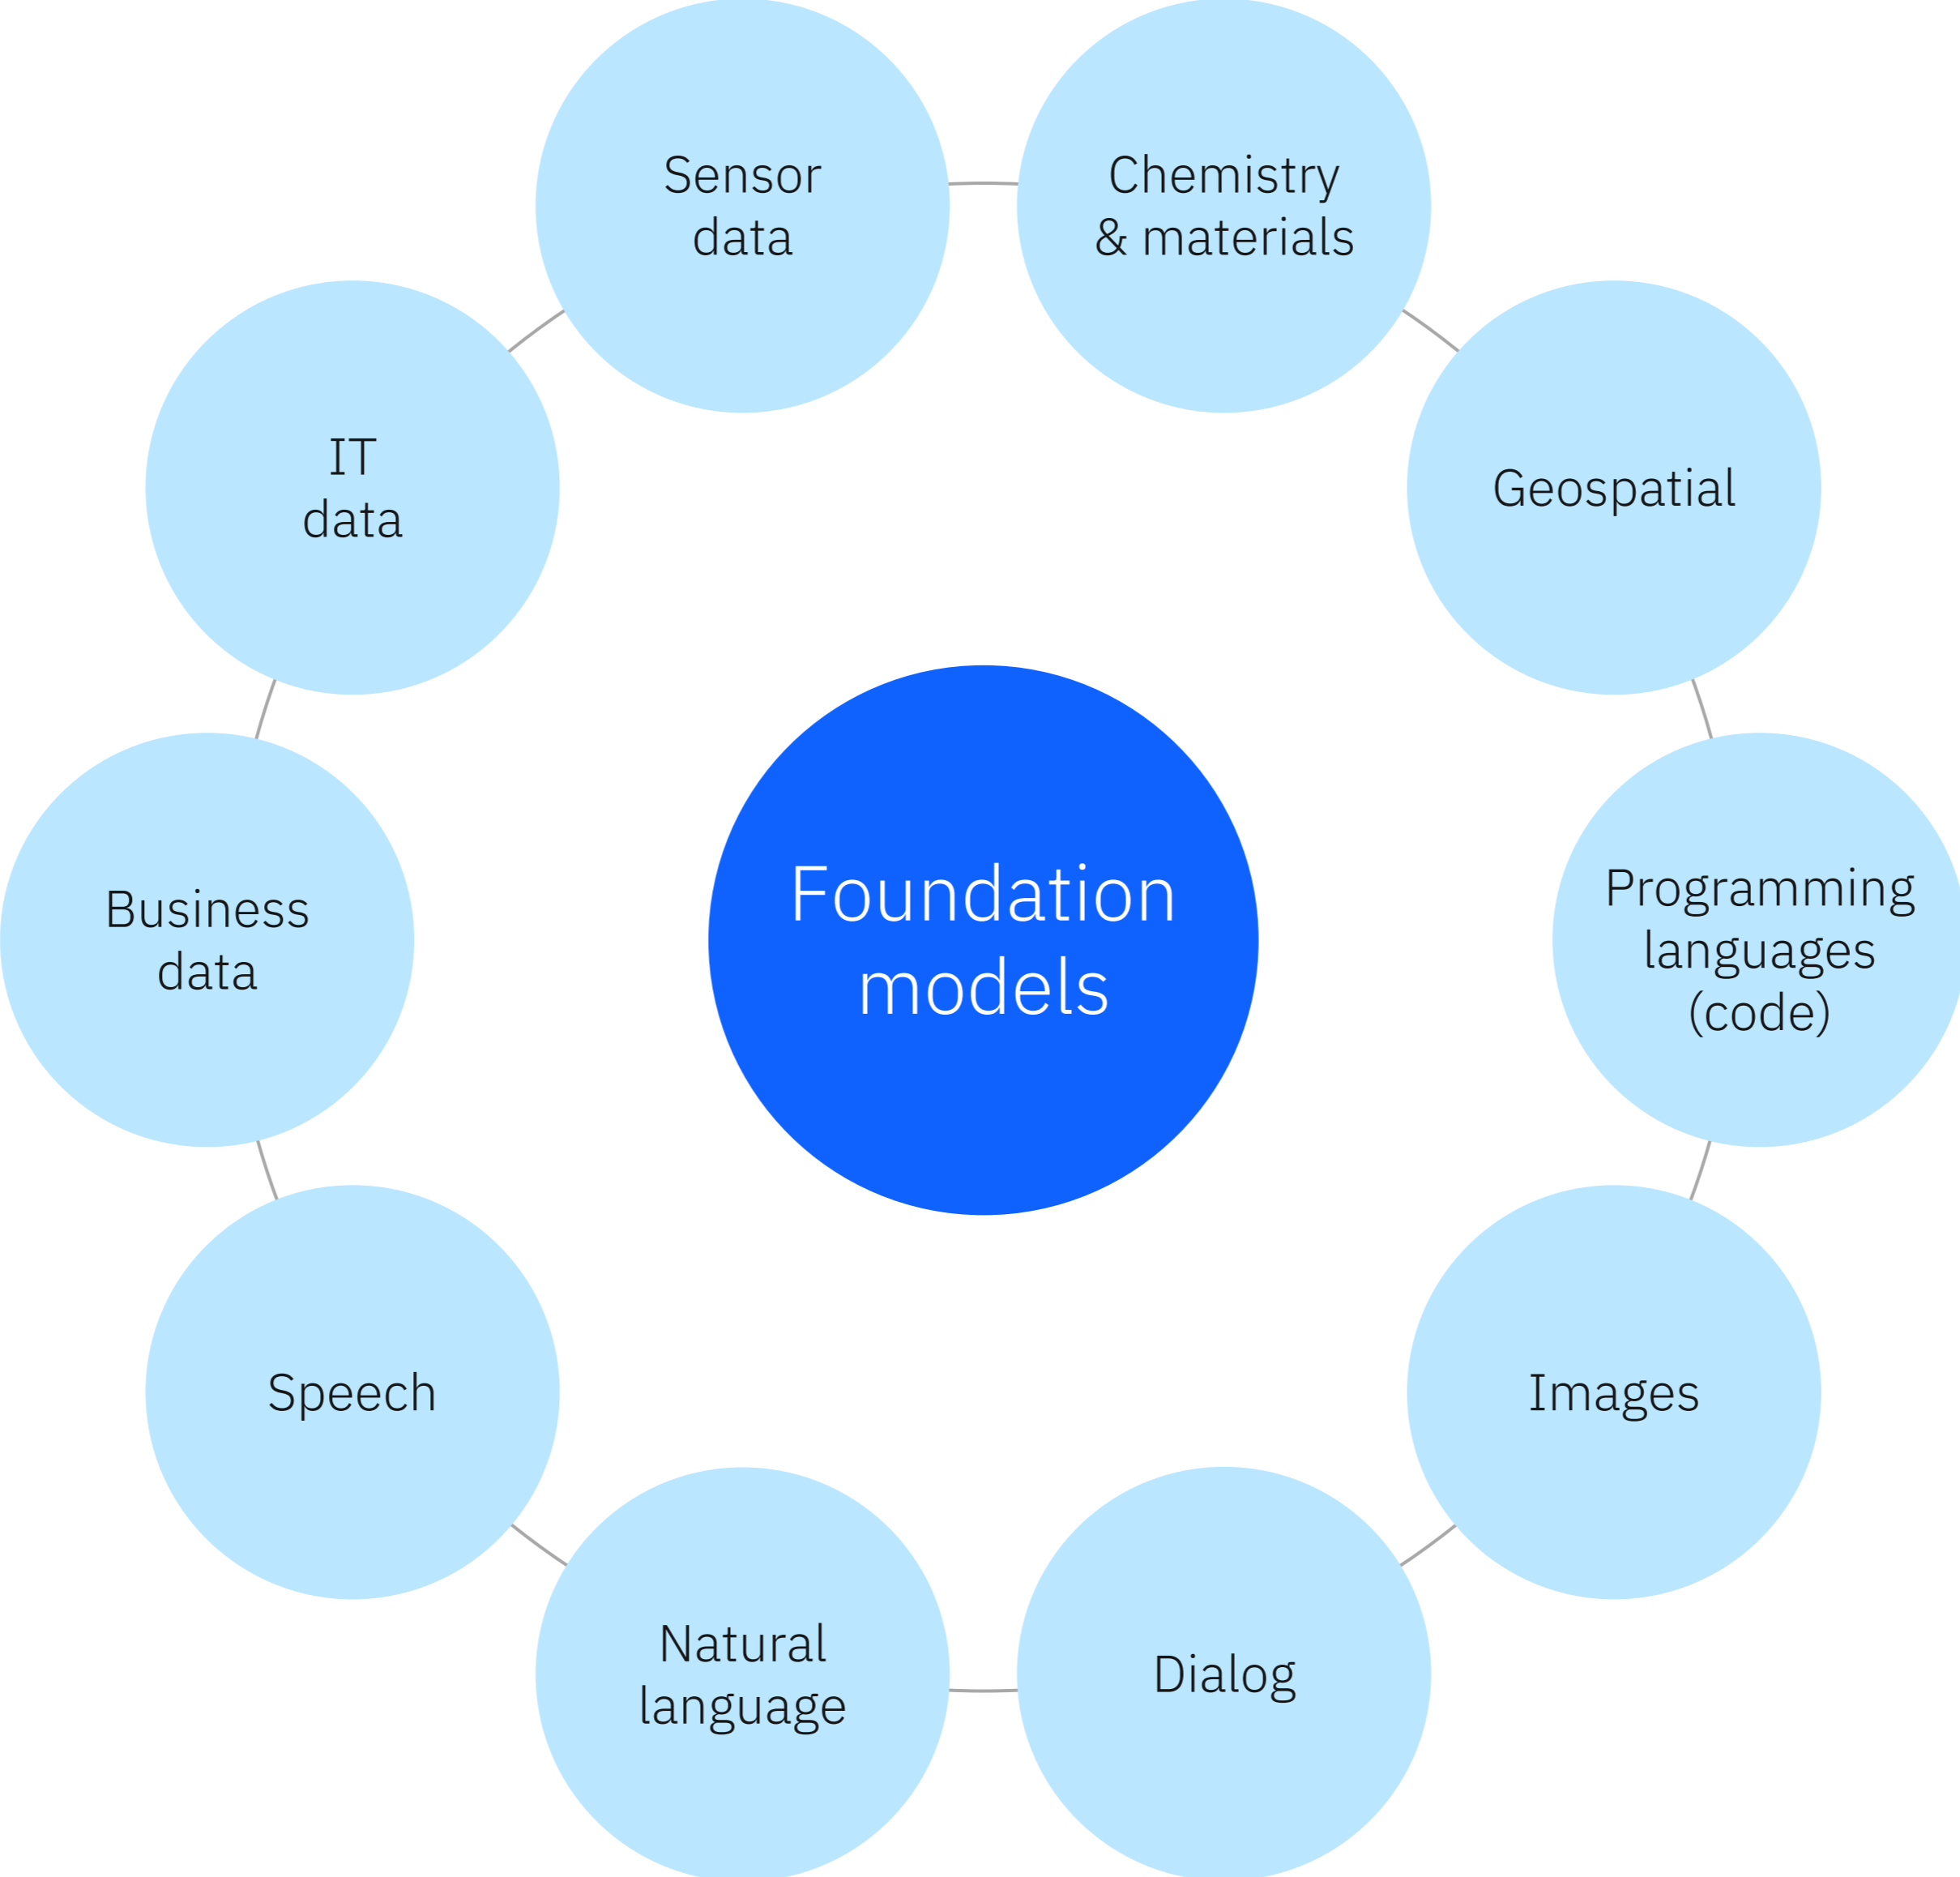


Which will
change business

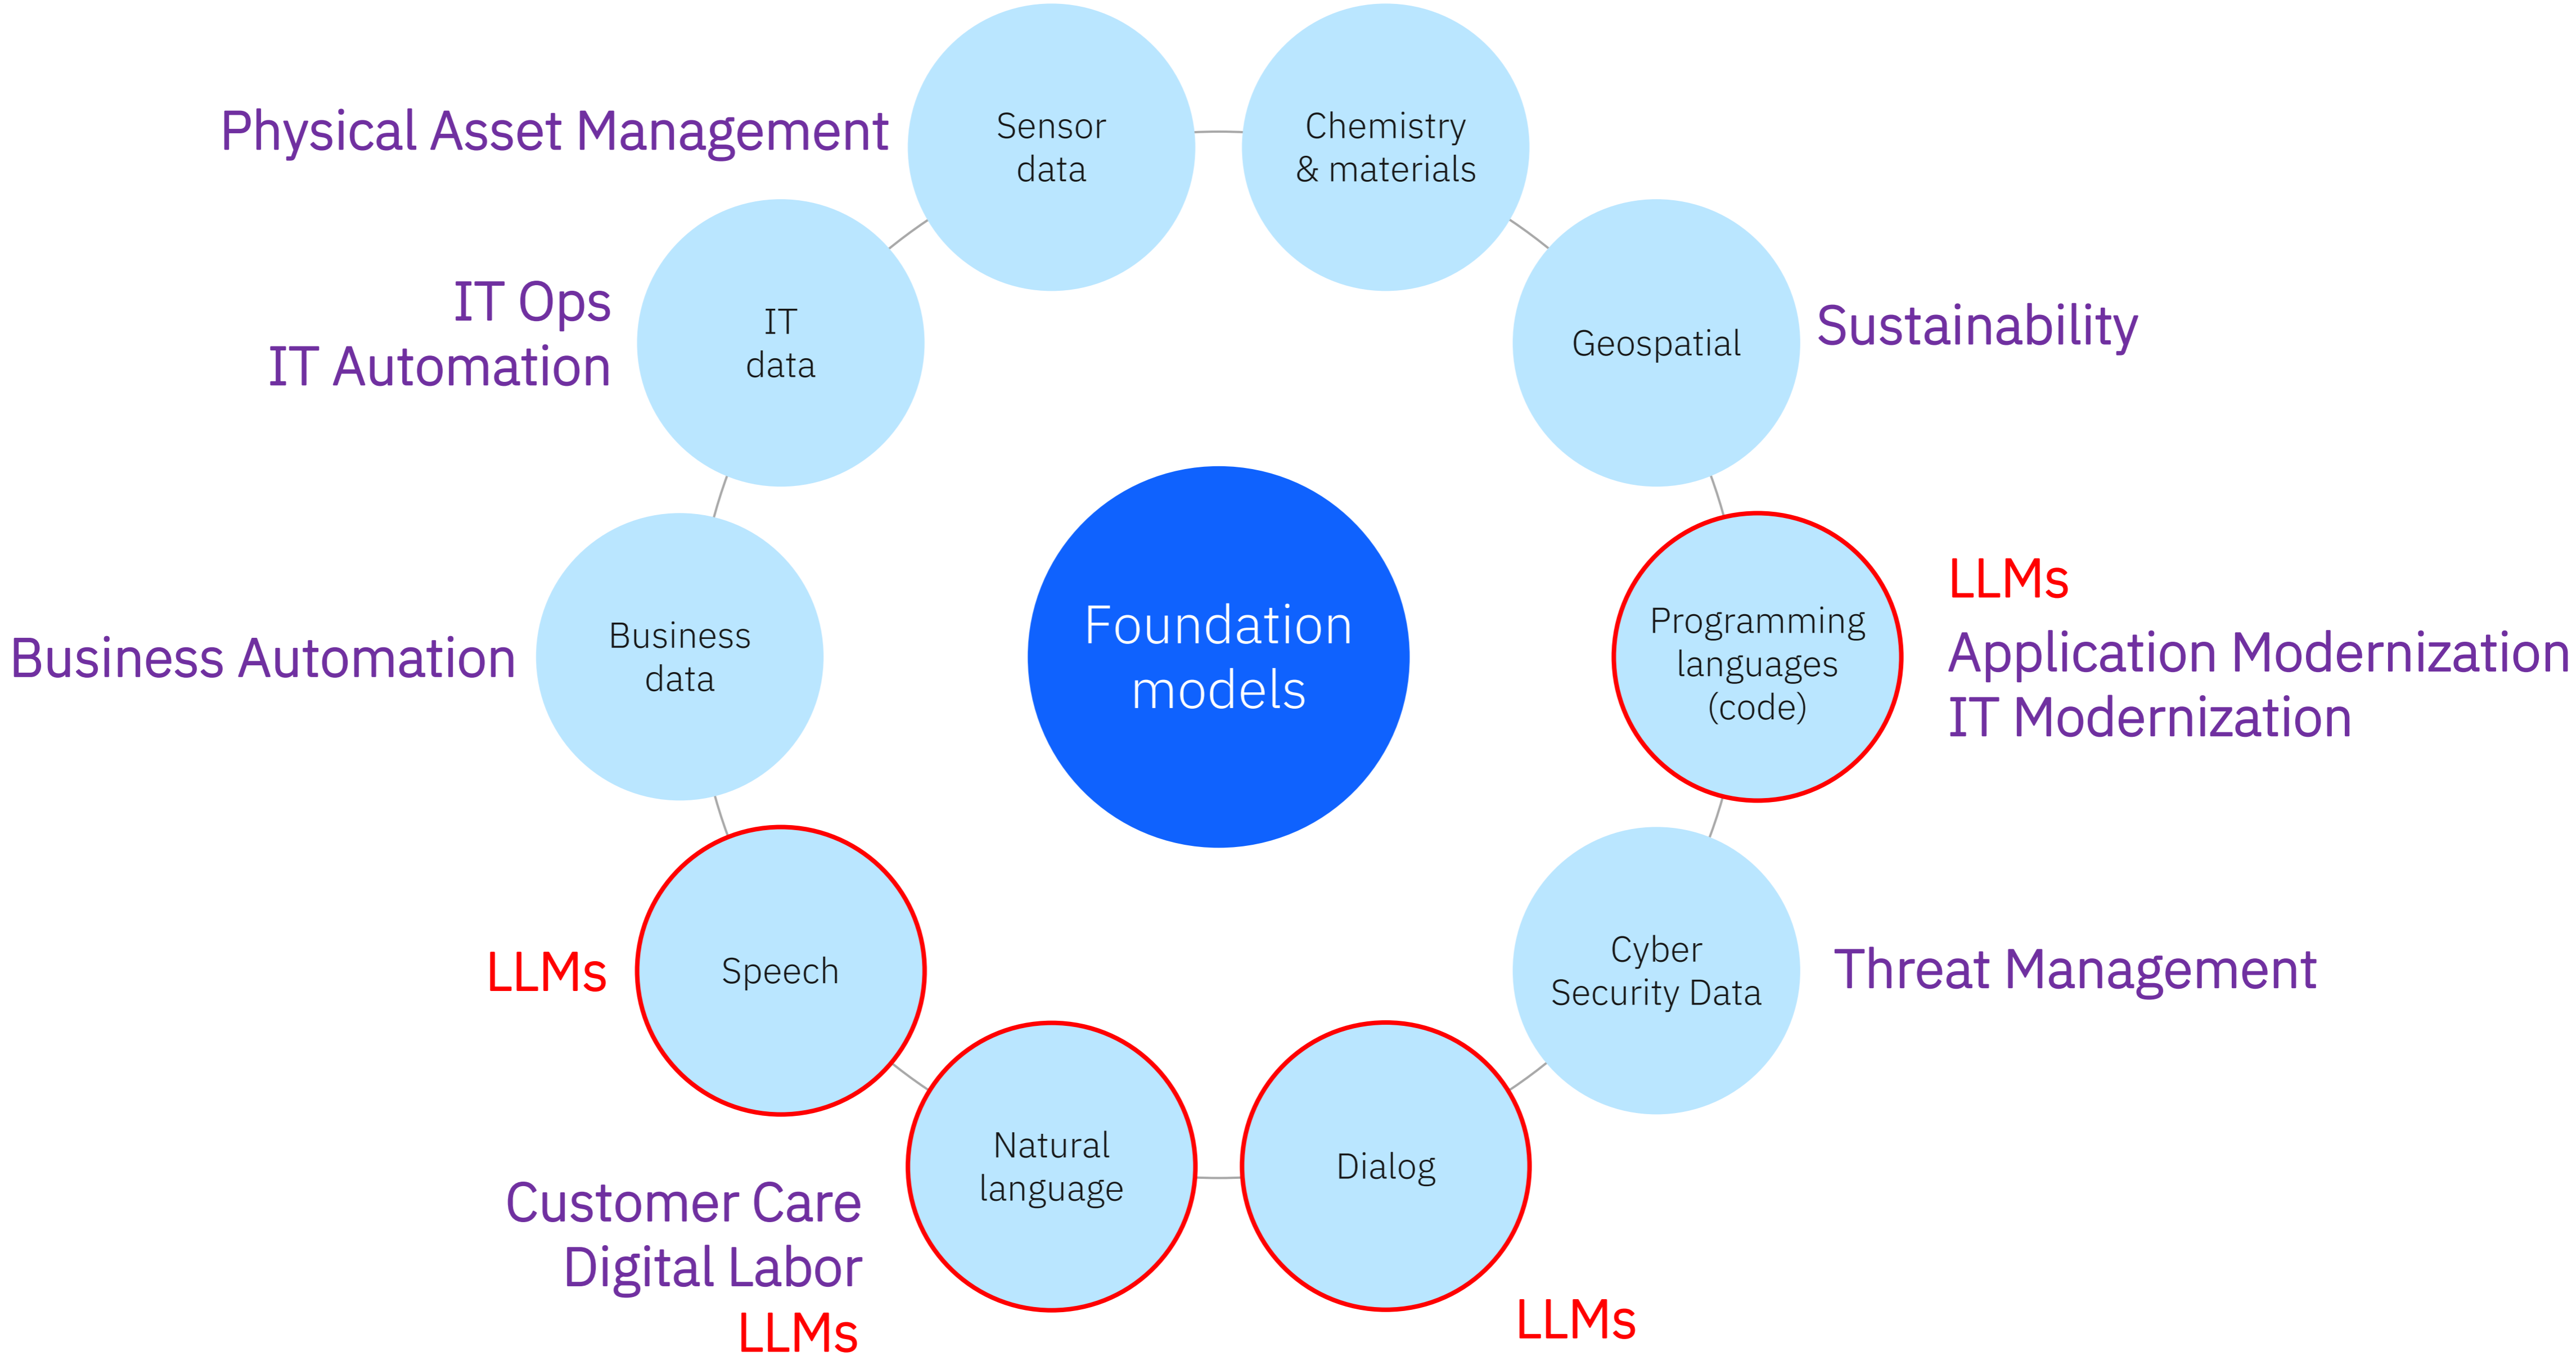


Generative AI refers to a set of AI algorithms that can generate new outputs — such as text, images, code, or audio — based on the training data, unlike traditional AI systems that are designed to recognize patterns and make predictions. Sometimes the AI that powers these solutions are referred to as decoders.

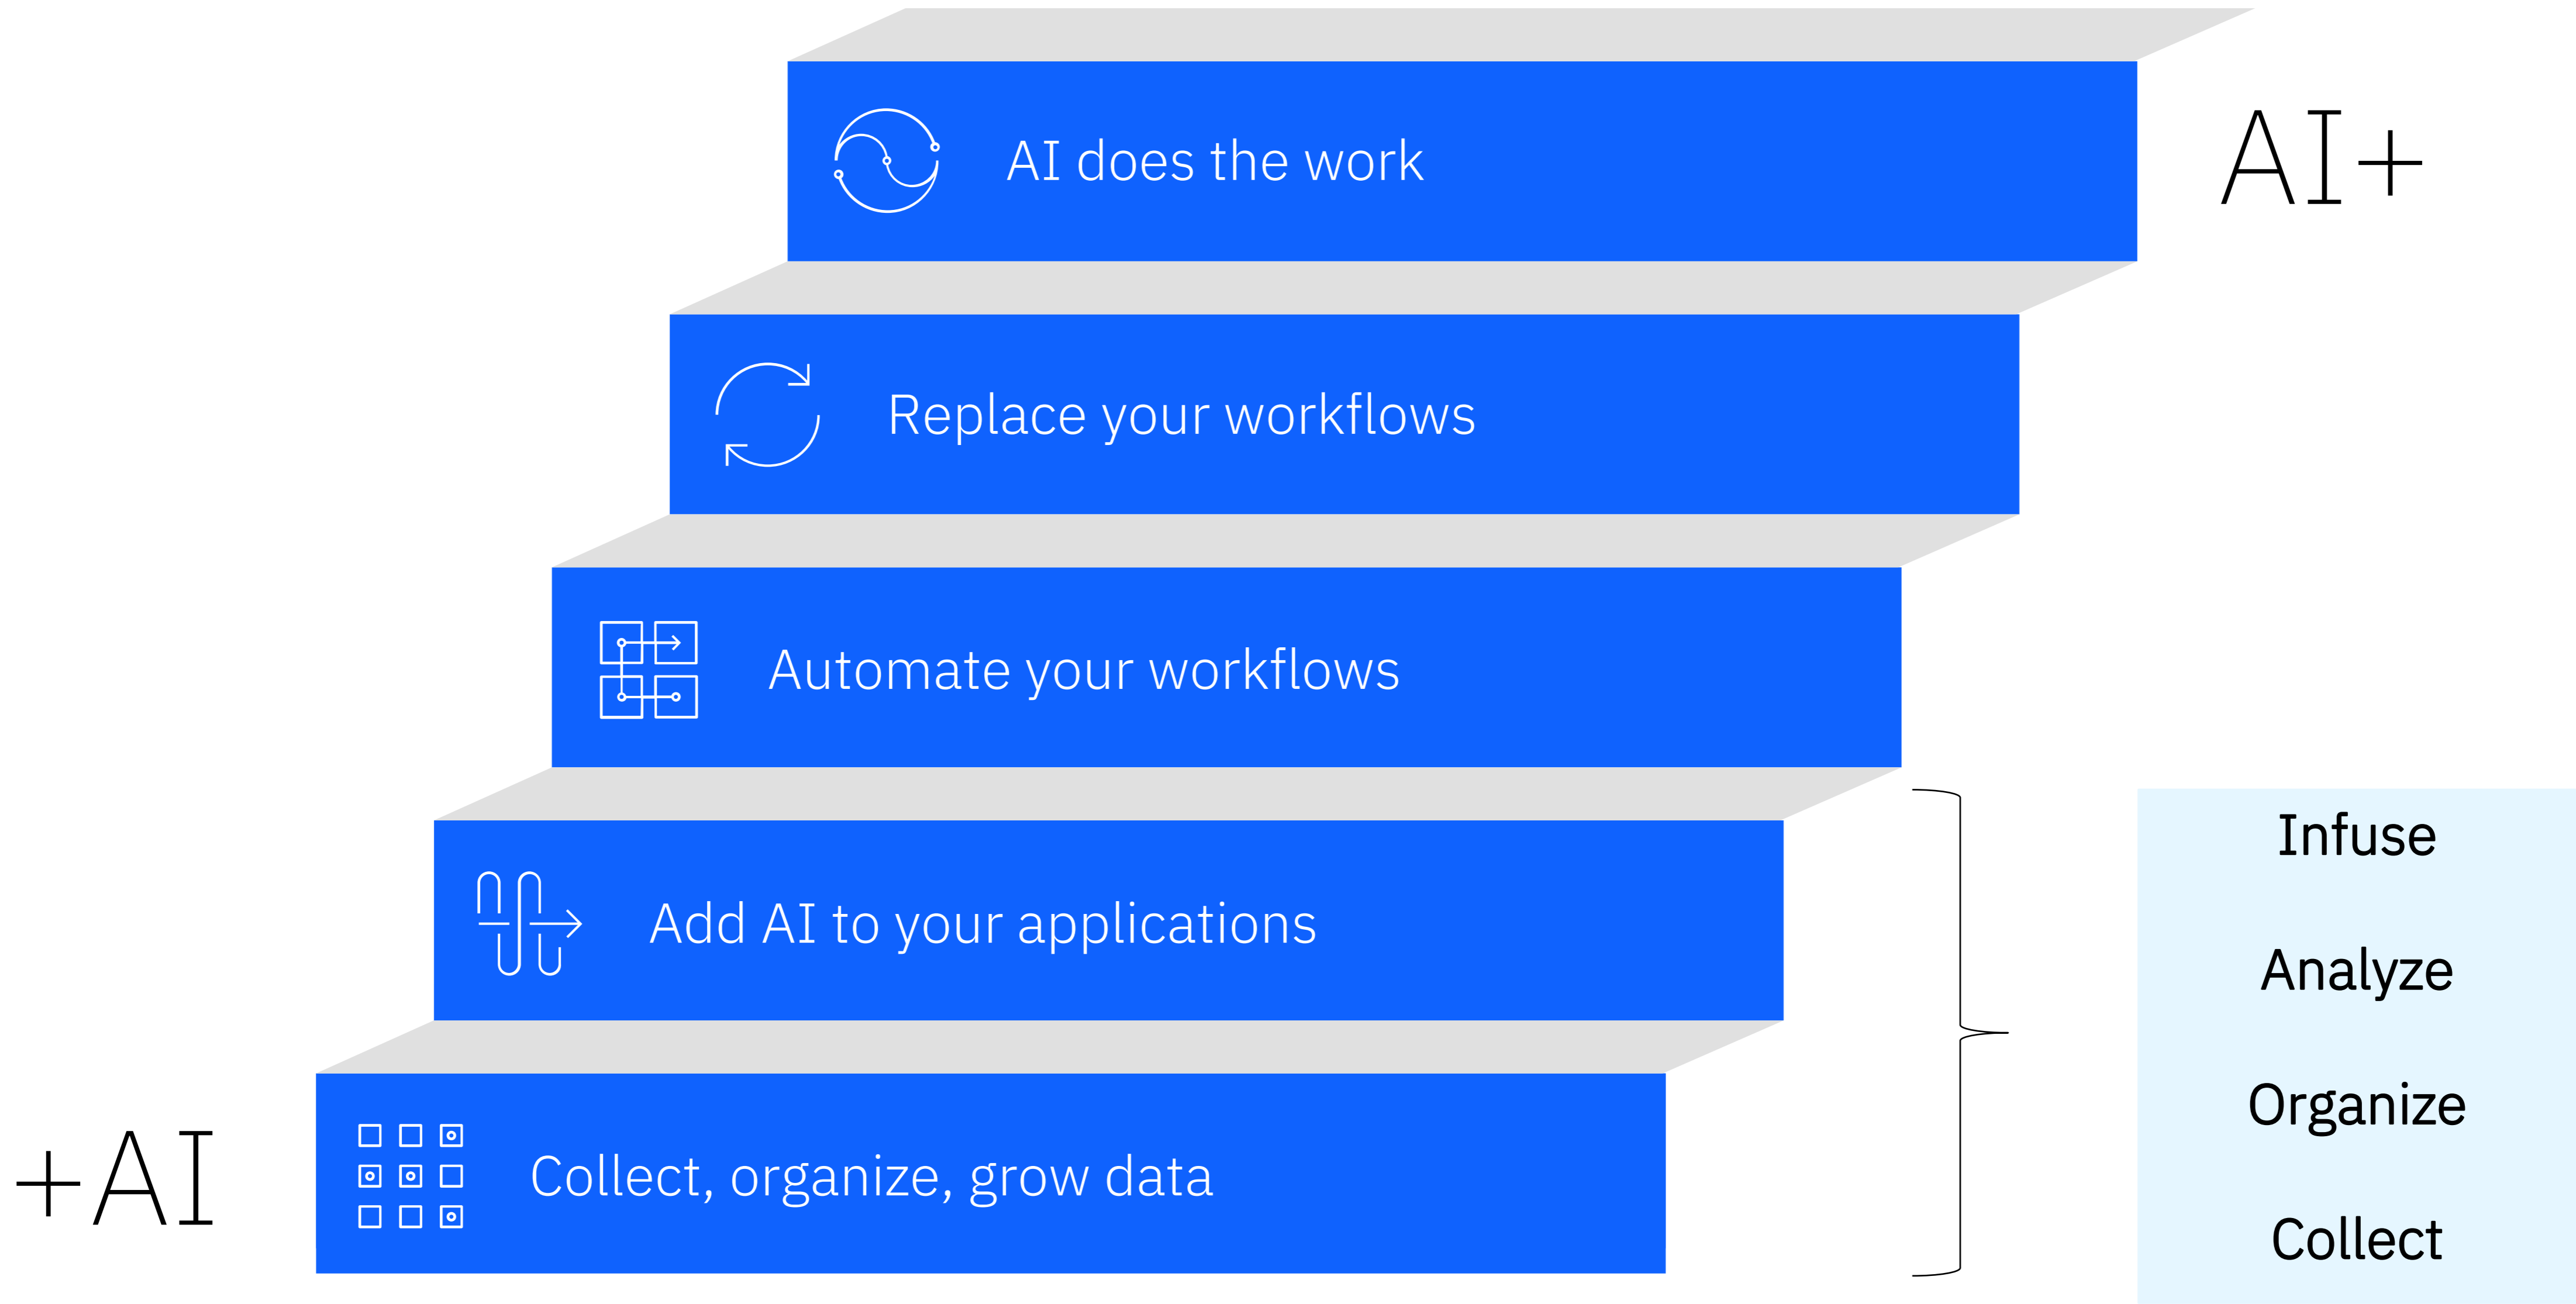
Incredible opportunities around enterprise data



Incredible opportunities around enterprise data



The modern-day AI ladder



Scaling AI adoption with Ecosystem Partners



McDonalds uses AI to automate drive thru order taking, enabling employees to increase focus on food delivery and customer service.



SAP applications allow clients to use natural language to get information buried deep in SAP versus using its native programming language.



Bring the power of AI to IT Automation to transform the developer experience.



GM's in-car applications. Enabling future embedded intelligent applications



Adobe Acrobat to help better process PDFs for subscribers in a cloud based document library

AI governance is needed to manage risk and protect reputations

“Fewer than 20% of executives strongly agree that their organizations’ practices and actions on AI ethics match (or exceed) their stated principles and values.”

- IBM and Oxford Economics – AI ethics in action, 2021

IDEAS

Algorithmic bias isn't just unfair — it's bad for business

If it's not deployed wisely, artificial intelligence can turn consumers off.

By Kalinda Ukanwa Updated May 23, 2021, 3:00 a.m.

YouTube sued for using AI to racially profile content creators

They claim YouTube's algorithms discriminate against black users

Data science during COVID-19: Some reassembly required

Most likely, the assumptions behind your data science model or the patterns in your data did not survive the coronavirus pandemic. Here's how to address the challenges of model drift

2018 / 4:04 PM / UPDATED 2 YEARS AGO

Amazon scraps secret AI recruiting tool that showed bias against women

The \$300m flip flop: how real-estate site Zillow's side hustle went badly wrong

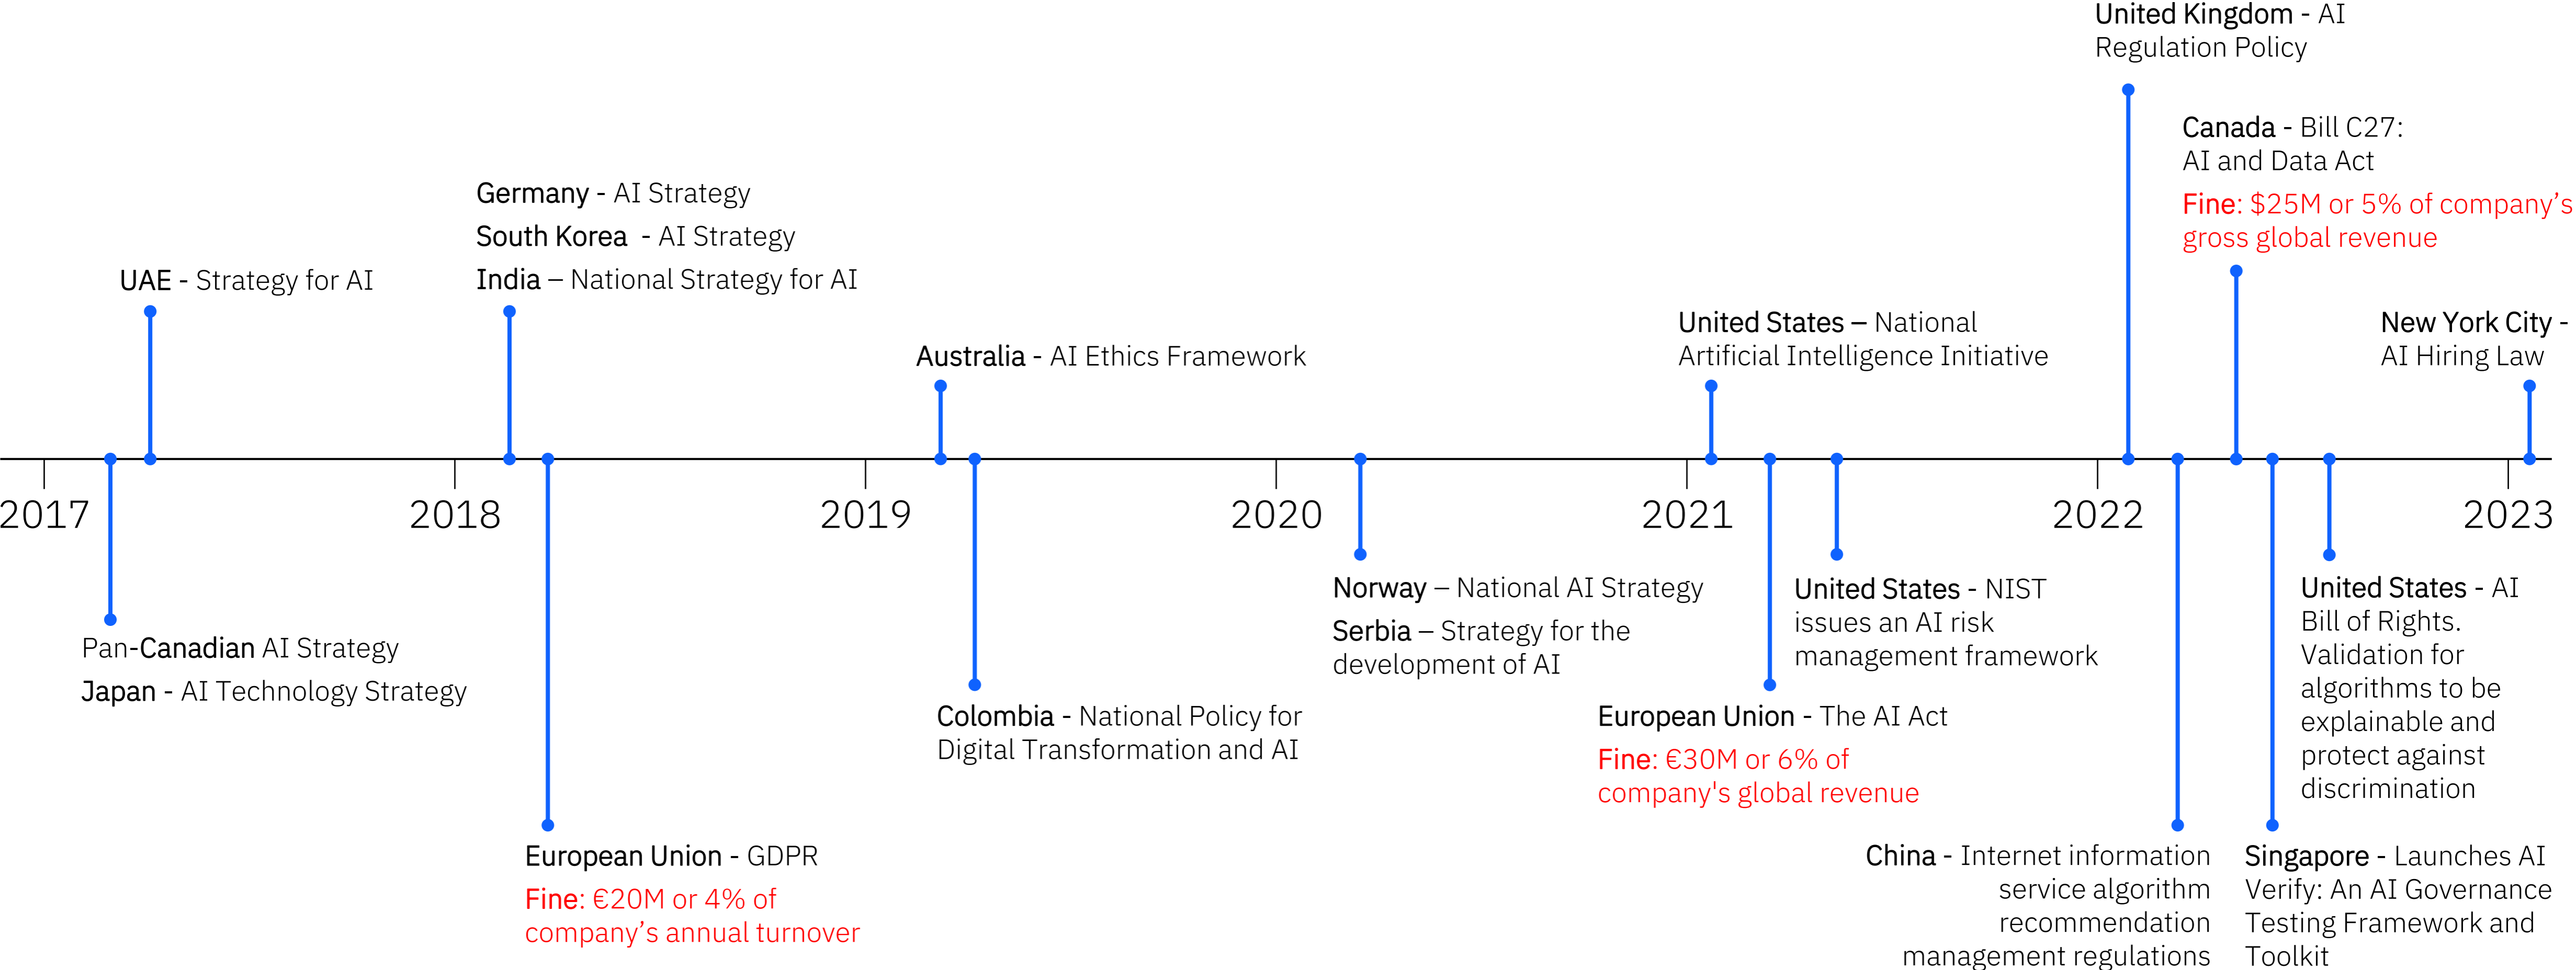
The Washington Post

Democracy Dies in Darkness

Apple Card algorithm sparks gender bias allegations against Goldman Sachs

Regulatory compliance

Constantly growing and changing regulations drive the need for governance



Sarbanes-Oxley Act



Why should organizations that build or use AI care about ethics?

Company values

Company reputation

Social justice and equity

Client and investor inquiries

Differentiation

Business opportunities

Existing or expected regulations





Pillars of trust

The purpose of
AI is to augment
human intelligence

Data & Insights
belong to their
creator

Explainability

An AI system's ability to provide a human-interpretable explanation for its predictions and insights.

Fairness

An AI system's ability to treat individuals or groups equitably, depending on the context in which the AI system is used.

Robustness

An AI system's ability to effectively handle exceptional conditions, such as abnormalities in input.

Transparency

An AI system's ability to include and share information on how it has been designed and developed.


Privacy

An AI system's ability to prioritize and safeguard consumers' privacy and data rights.

Your AI is only as good as your data.



Data Fabric



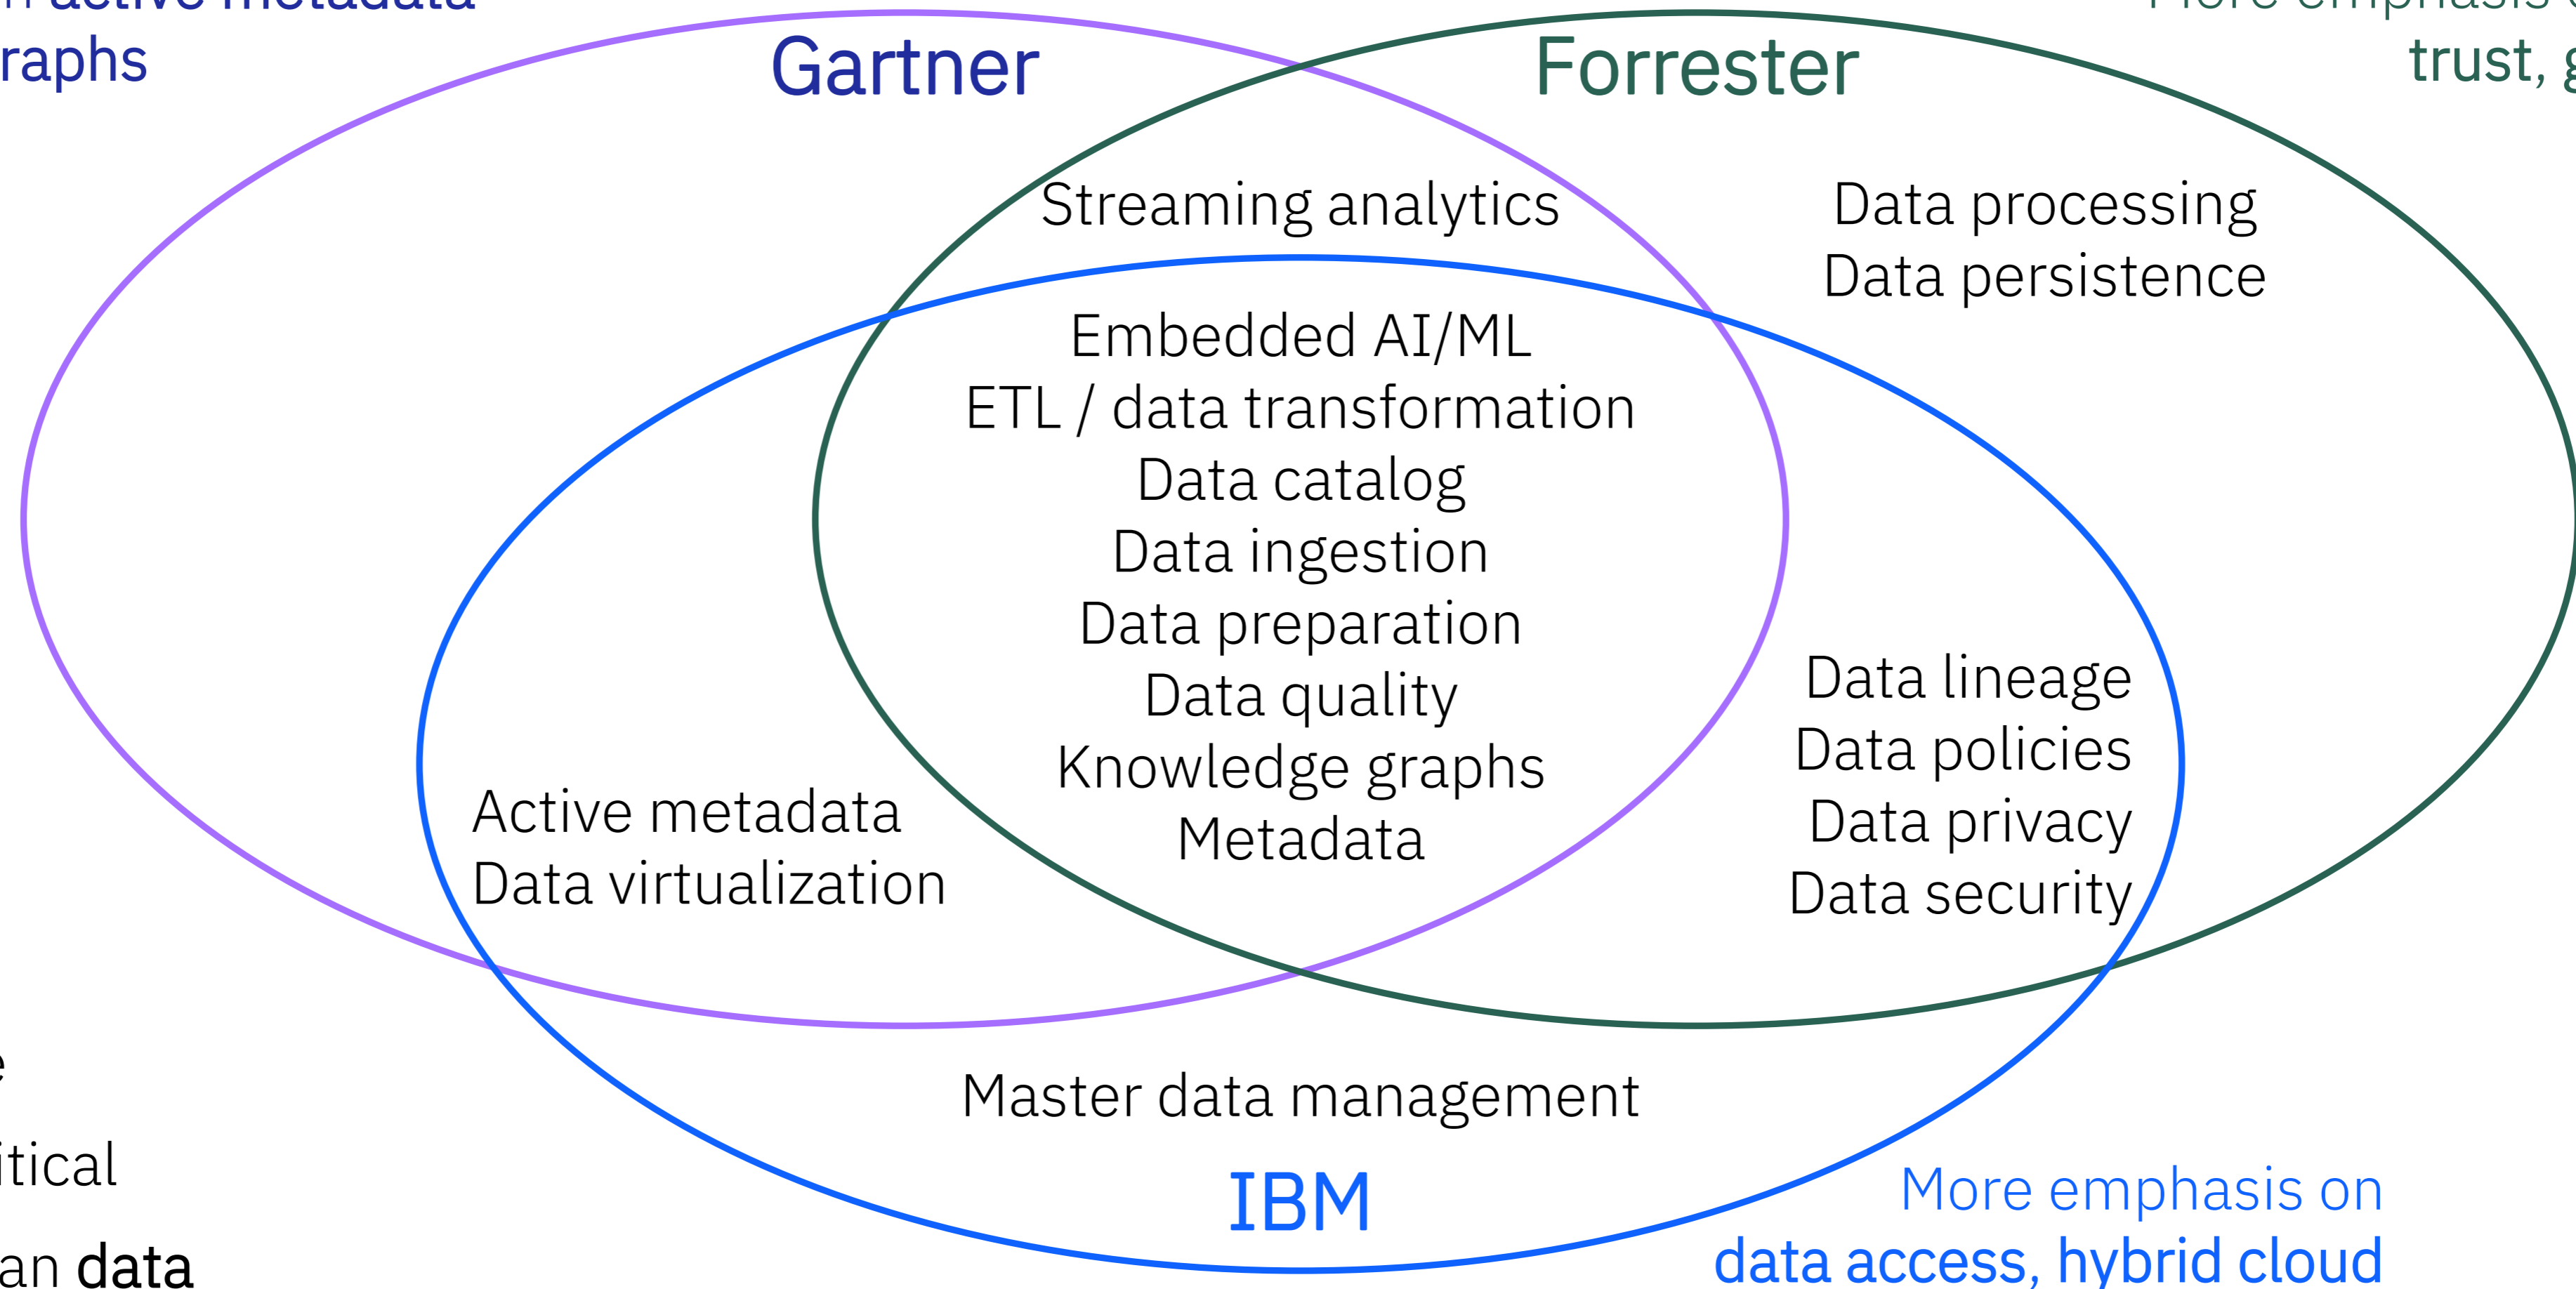
A data fabric is an **architectural approach** to simplify data access in an organization to **facilitate self-service data consumption**. This architecture is **agnostic to data environments, processes, utility, and geography**, all while **integrating** end-to-end data-management capabilities. A data fabric **automates** data discovery, governance, and consumption, enabling enterprises to use data to maximize their value chain. With a data fabric, enterprises elevate the value of their data by **providing the right data, at the right time**, regardless of where it resides.

Data Fabric definitions from Gartner and Forrester

IBM's definition is mostly aligned

More emphasis on **active metadata** and **knowledge graphs**

More emphasis on **security, trust, governance**



Common threads

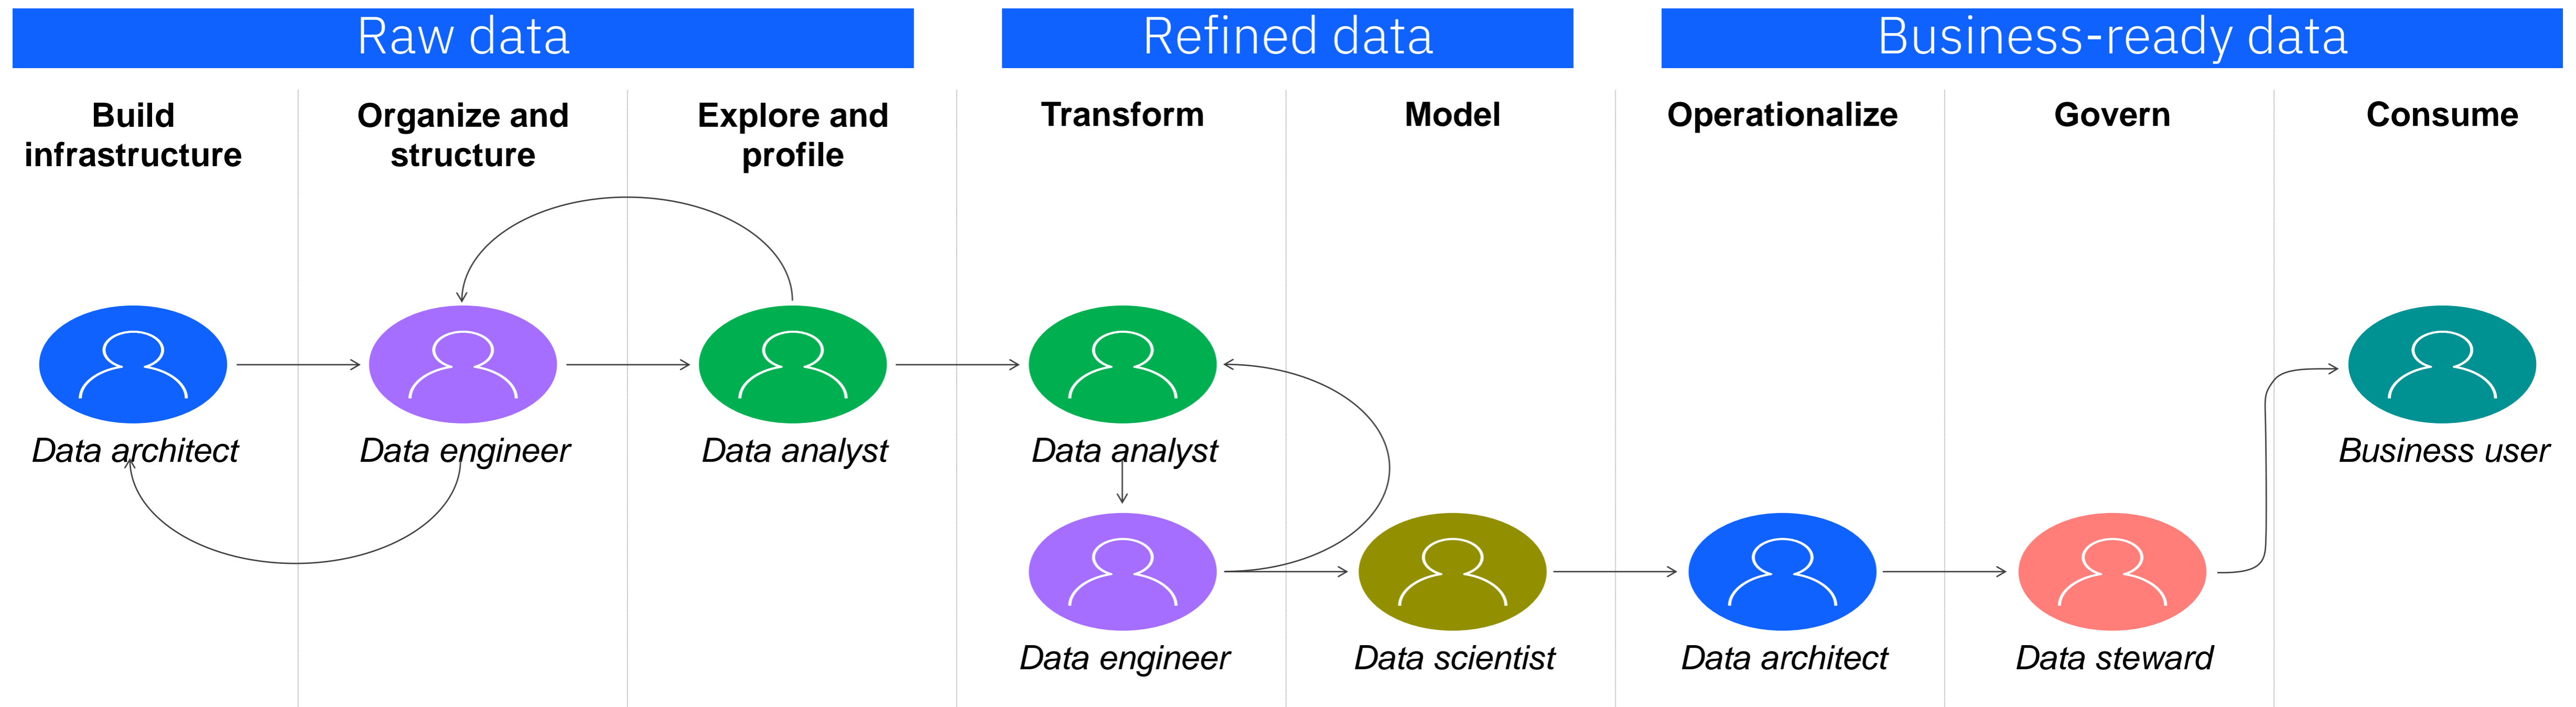
- It's an architecture
- Embedded AI is critical
- Key capabilities span **data integration** and **governance**
- Connects disparate data

DataOps

Data operations (DataOps) is the **orchestration** of people, processes, and technology to **deliver trusted, high-quality data** to data citizens fast.

Data fabric architectures need an agile and effective DataOps practice

DataOps is the booster for a data fabric solution, ensuring that those components are used to their full potential to facilitate self-service data exploration and experimentation.



Data Mesh

*Data mesh is an **approach** centered on **organizational processes** to enable agile, **domain-specific ownership** and **creation of reusable data products**. It is **technology agnostic** and domain owners are responsible for the entire data lifecycle.*

Domain ownership

End-to-end data ownership belongs to the applicable line of business teams and departments

Data as a product

Data needs to be tied to business goals and the value they generate

Self-service data infrastructure

Any data consumer, regardless of role must have easy access to the data they need

Federated computational governance

Enables self-service for easy end user data consumption and collaboration

The platform
for AI and data

watsonx

Scale and
accelerate the
impact of AI with
trusted data.

watsonx.ai

Train, validate, tune
and deploy AI models

A next generation enterprise studio for AI builders to train, validate, tune, and deploy both traditional machine learning and new generative AI capabilities powered by foundation models. It enables you to build AI applications in a fraction of the time with a fraction of the data.

watsonx.data

Scale AI workloads, for
all your data, anywhere

Fit-for-purpose data store optimized for governed data and AI workloads, supported by querying, governance and open data formats to access and share data.

watsonx.governance

Enable responsible,
transparent and explainable
data and AI workflows

End-to-end toolkit encompassing both data and AI governance to enable responsible, transparent, and explainable AI workflows.

What IBM offers

IBM's AI is embedded in applications built on

watsonx

Watson Orchestrate

Harnesses the power of AI and automation to free up individuals from tedious tasks

40%

Improvement in HR productivity

Watson Assistant

Builds better virtual agents, to deliver consistent and intelligent customer care

70%

Call center calls contained by conversational AI

Watson Code Assistant

Enables hybrid cloud developers to write code with AI-generated recommendations

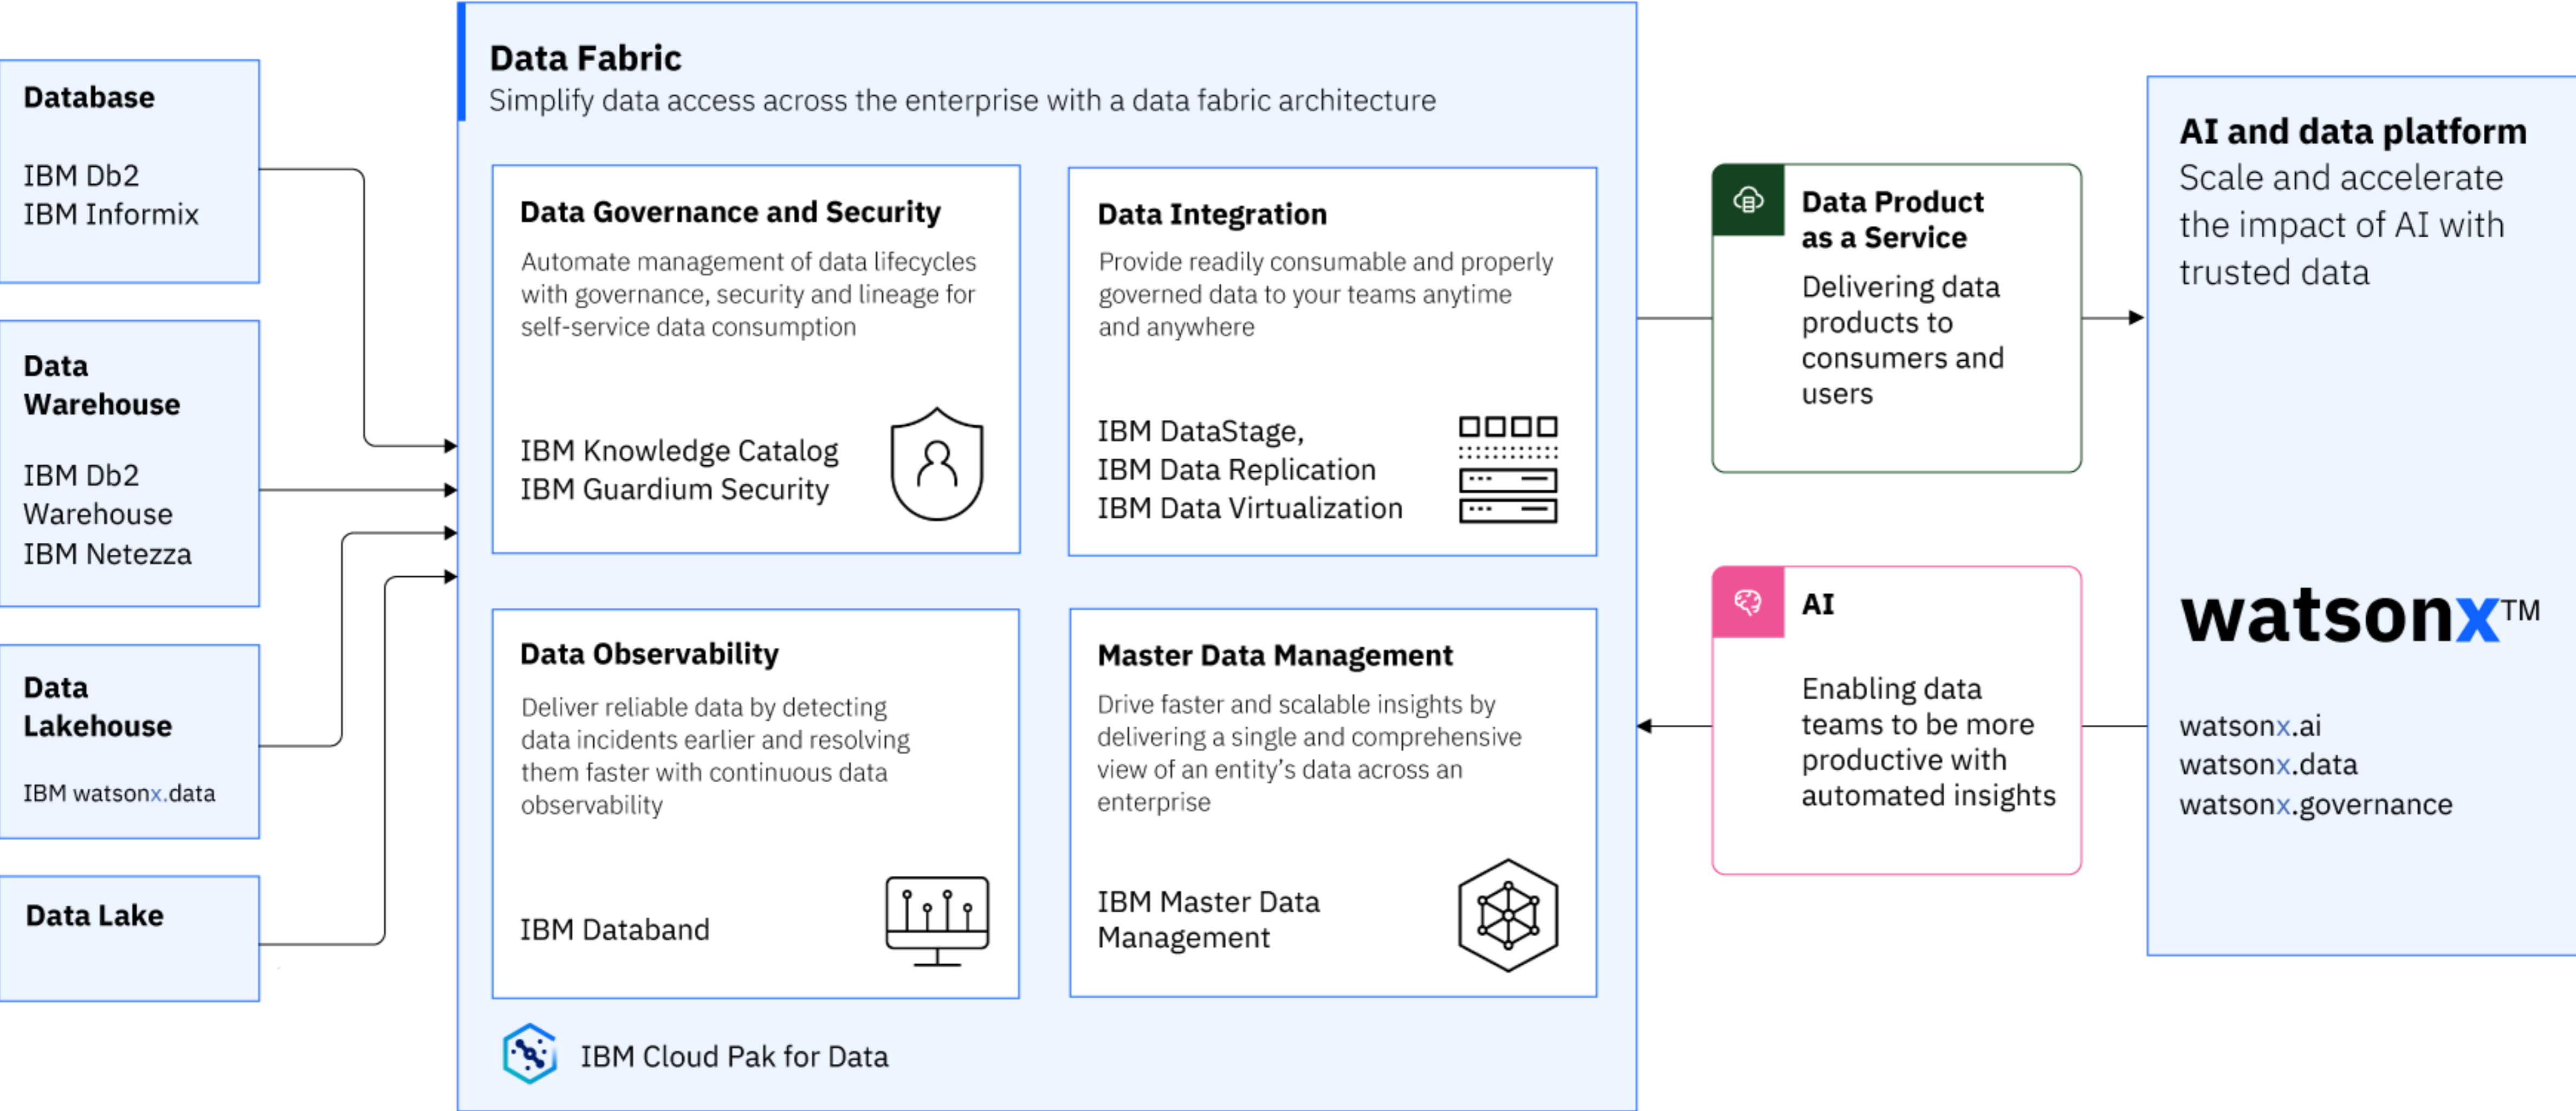
30%

Productivity gain in application modernization

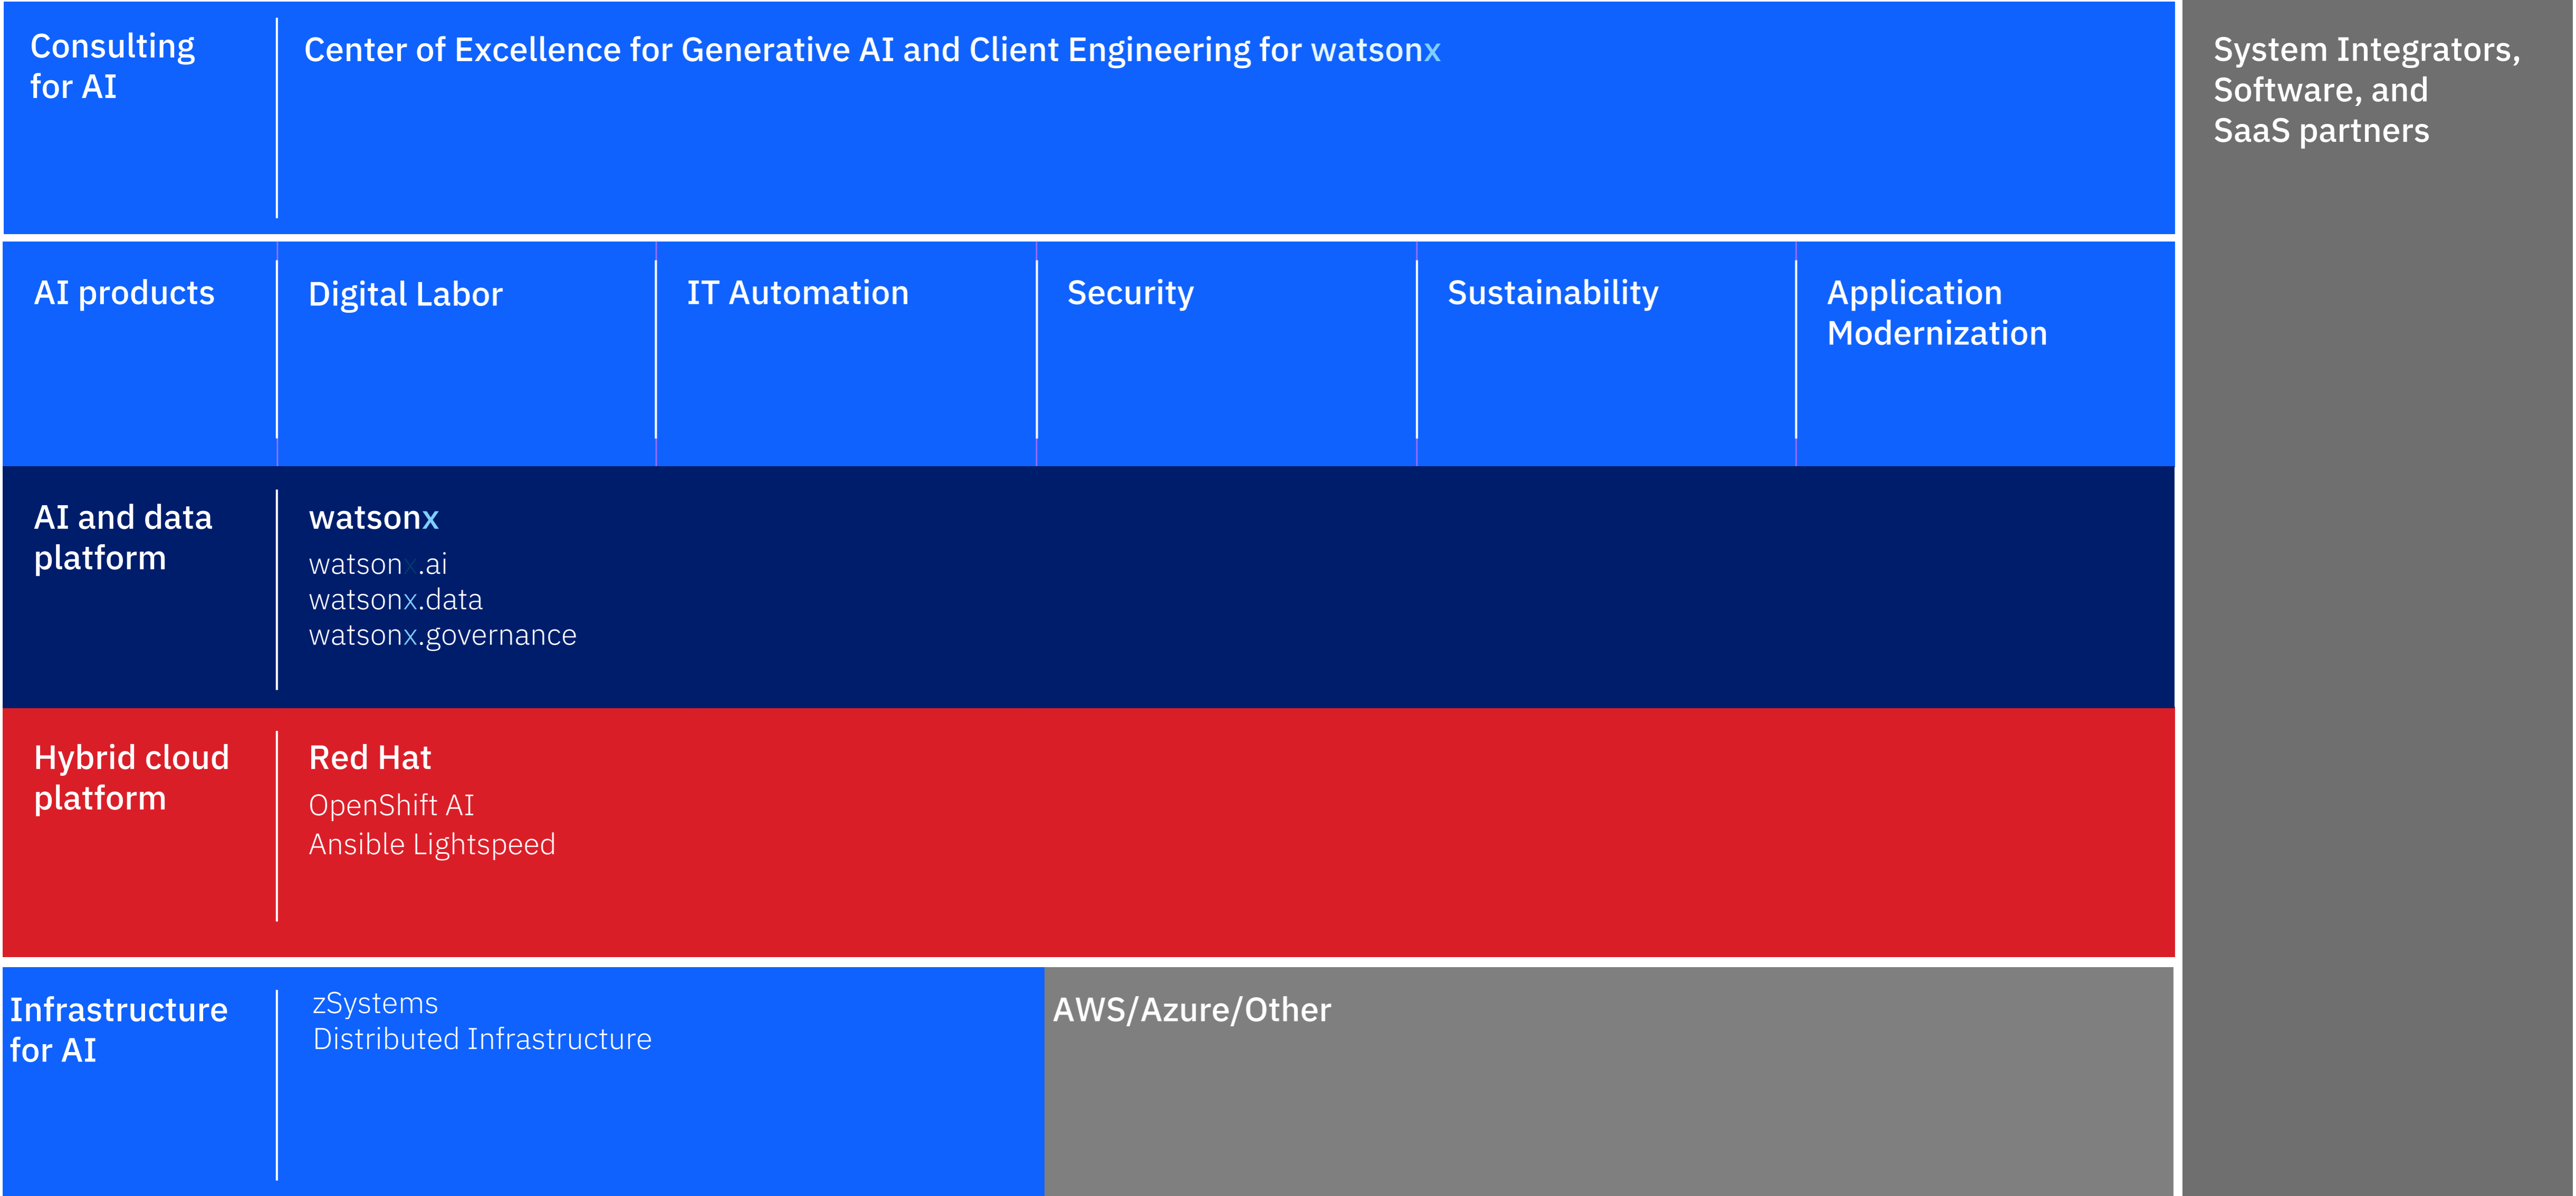
AI and data platform

watsonx

Investments in a trusted data foundation will accelerate and scale AI



IBM approach for AI: Unleash the intelligence in your business



Reinventing how work gets done | +AI to AI+

IBM is actively engaging with enterprise clients across a broad set of business domains

NON-EXHAUSTIVE

Customer-facing functions and experiences	HR, Finance, and Supply Chain functions	IT development and operations	Core business operations
<p>Customer service Empower customers to find solutions with easy, compelling experiences.</p> <p>Automate answers with 95% accuracy</p>	<p>HR automation Reduce manual work and automate recruiting, sourcing and nurturing job candidates.</p> <p>Reduce employee mobility processing time by 50%</p>	<p>App modernization, migration Generate code, tune code generation response in real time.</p> <p>Deliver faster development output</p>	<p>Threat management Reduce incident response times from hours to minutes or seconds.</p> <p>Contain potential threats 8x faster</p>
<p>Marketing Increase personalization, improve efficiency across the content supply chain.</p> <p>Reduce content creation costs by up to 40%</p>	<p>Supply chain Automate source to pay processes, reduce resource needs and improve cycle times.</p> <p>Reduce cost per invoice by up to 50%</p>	<p>IT automation Identify deployment issues, avoiding incidents, optimize application demand to supply.</p> <p>Reduce mean time to repair (MTTR) by 50%+</p>	<p>Asset management Optimize critical asset performance and operations while delivering sustainable outcomes.</p> <p>Reduce unplanned downtime by 43%</p>
<p>Content creation Ex. Enhance digital sports viewing with auto-generated spoken AI commentary.</p> <p>Scale live viewing experiences cost effectively</p>	<p>Planning and analysis Make smarter decisions, focus on higher value tasks with automated workflows and AI.</p> <p>Process planning data up to 80% faster</p>	<p>AIOps Assure continuous, cost-effective performance and connectivity across applications.</p> <p>Reduce application support tickets by 70%</p>	<p>Product development Ex. Expedite drug discovery by inferring structure with AI from simple molecular representations.</p> <p>Faster and less expensive drug discovery</p>
<p>Knowledge worker Enable higher value work, improve decision making, and increase productivity.</p> <p>Reduce 90% of text reading and analysis work</p>	<p>Regulatory compliance Support compliance based on requirements / risks, proactively respond to regulatory changes.</p> <p>Reduce time spent responding to issues</p>	<p>Data platform engineering Redesign the approach for data integration using generative AI.</p> <p>Reduce data integration time by 30%+</p>	<p>Environmental intelligence Provide intelligence to proactively plan and manage impact of severe weather and climate.</p> <p>Increase manufacturing output by 25%</p>

Modern-day hybrid and multi-cloud data management is *very* complex

Distributed data

- On-premises, cloud platforms, and edge devices
- Compatibility and migration

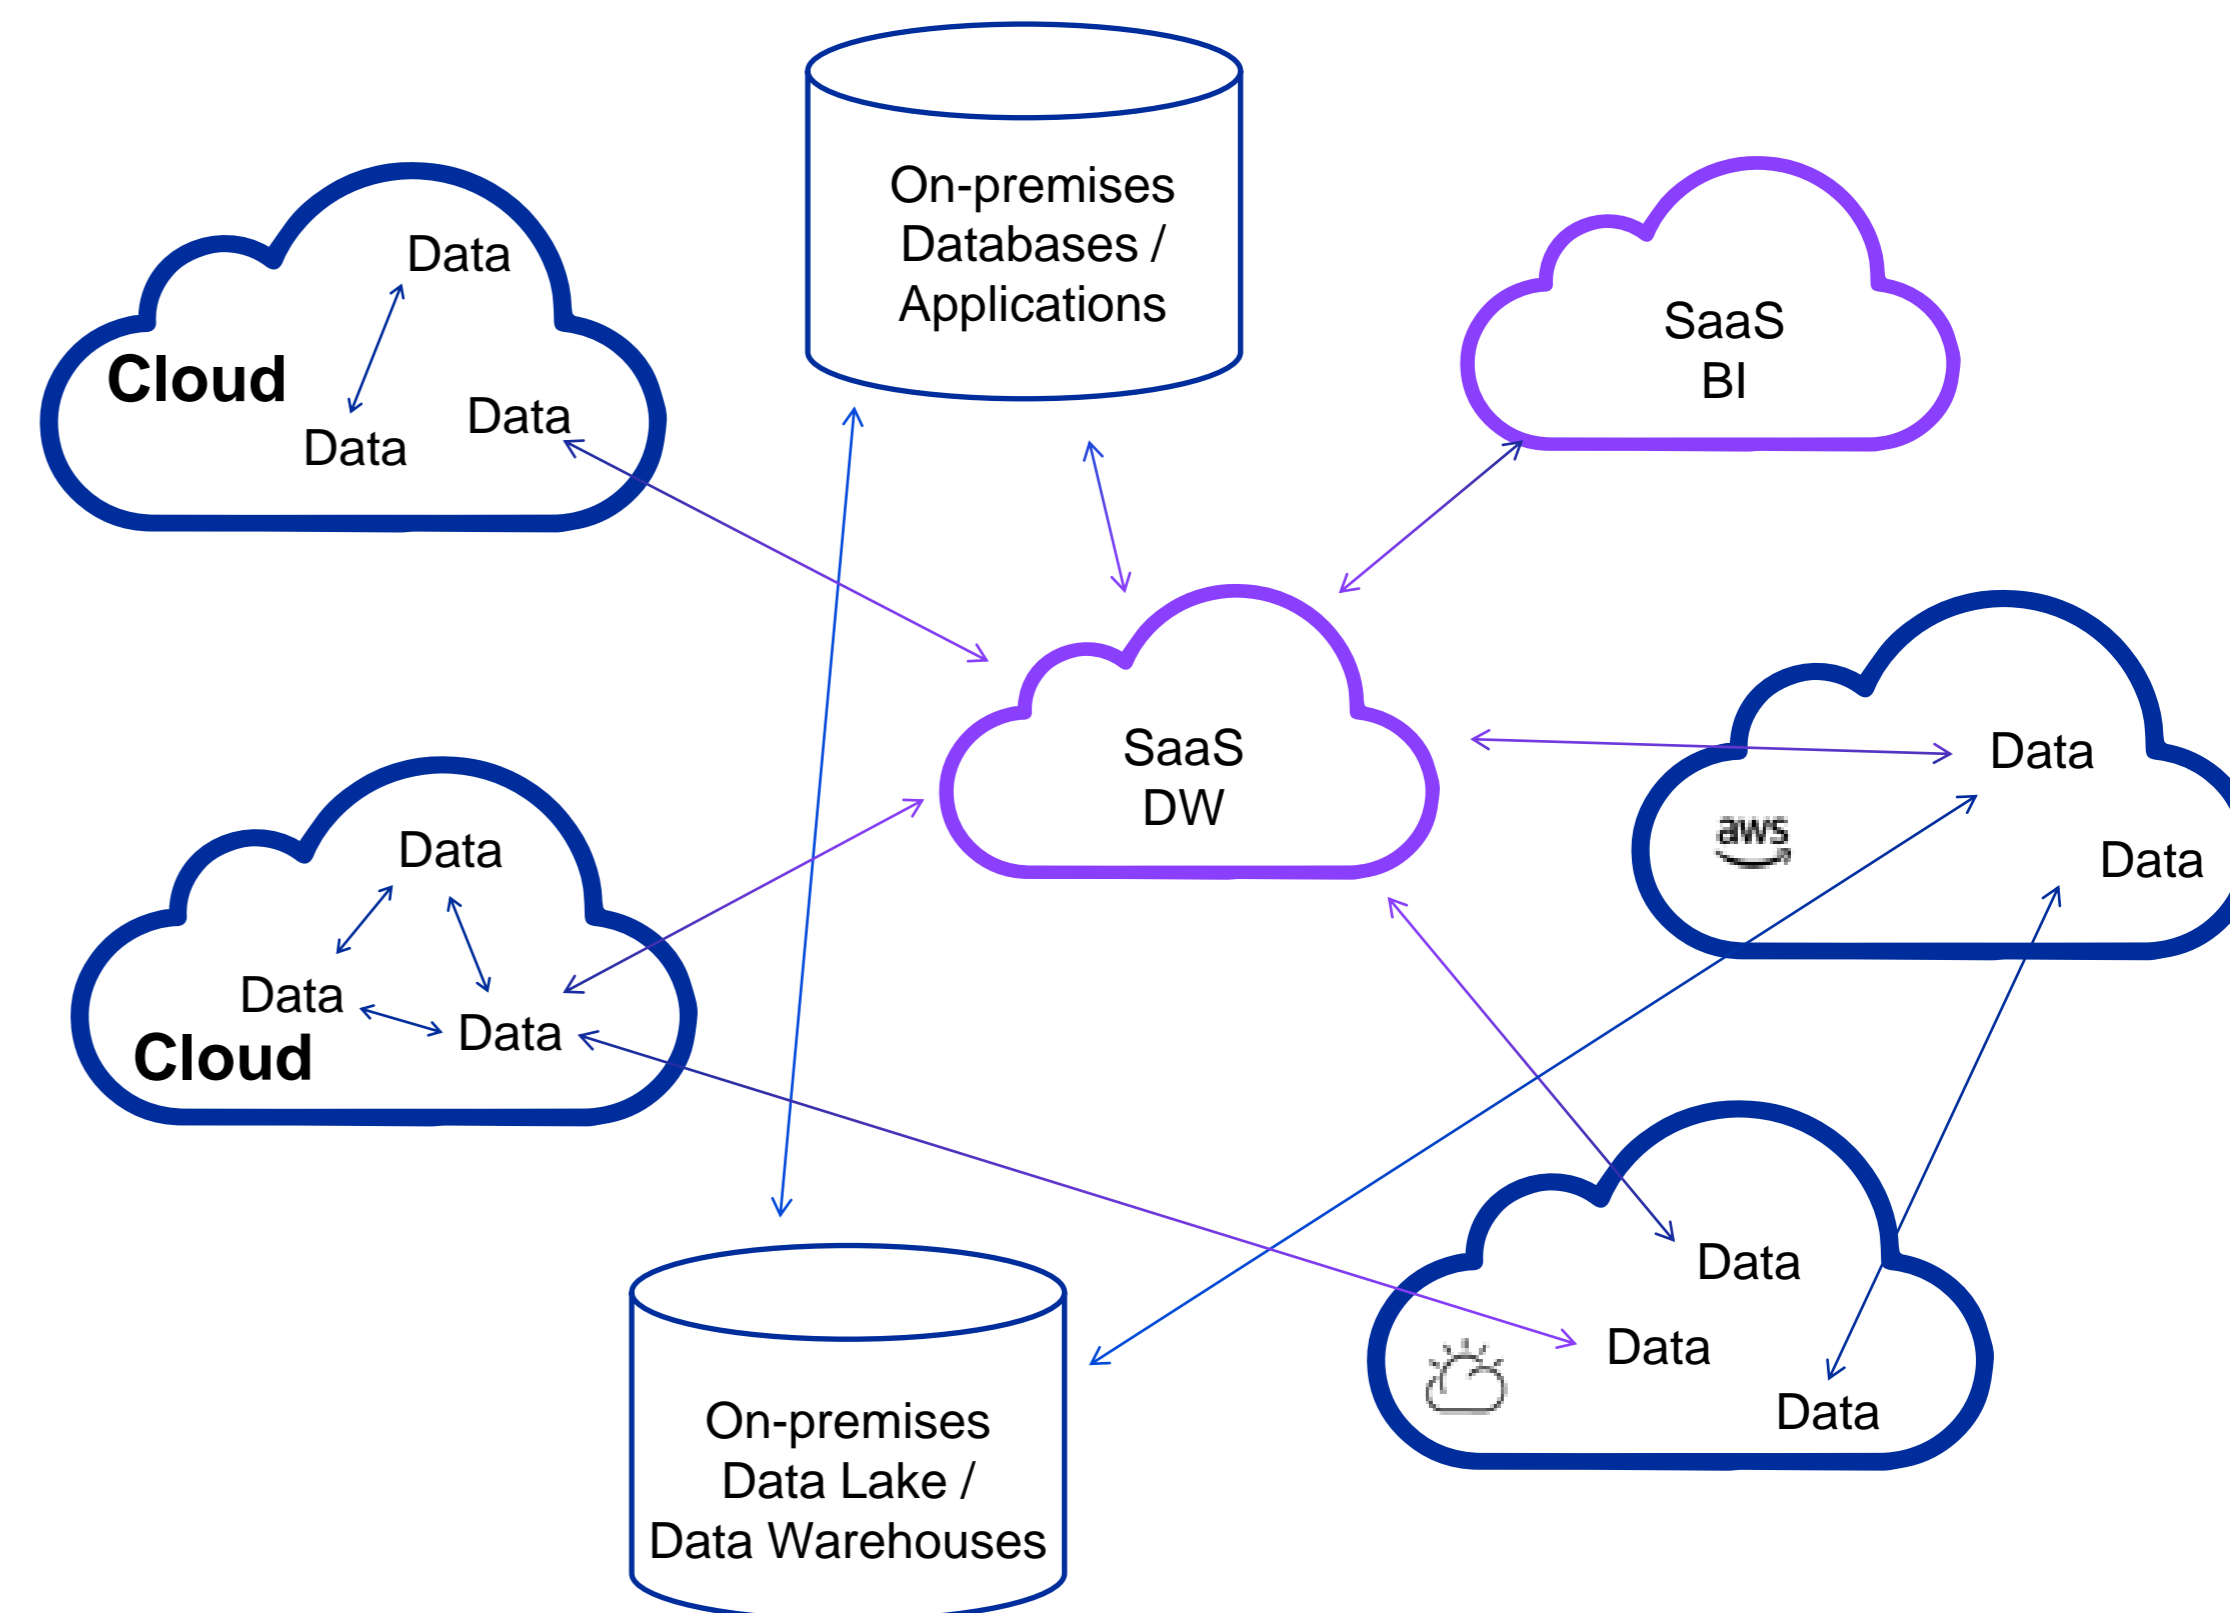
Increasingly strict requirements

- Government regulations
- Security and compliance

Volume, velocity and variety of incoming data

- Data types and workloads

Agility and scalability



...and it will get more complex, *fast*

By 2024 the volume of replicated data will be greater than the original unique data by a factor of

10x

50%

of the world's data will be stored in cloud databases **by 2025**

With the increasing data volume and complexity, businesses need a data management engine that **won't delay their business outcomes**

Properties of a sound hybrid data management architecture

All deployment methods


On-premises, cloud, and multi-cloud

All types of data and workloads

Structured and unstructured, transactional and analytical

Open-source integrated

Cost effective integration of open-source and enterprise data



All the data, together

Use virtualization to query data across deployments without moving it

Cloud agility

Rapidly develop, test, and deploy applications; elastically use resources

Integrated analytic & machine learning


Make smarter decisions faster

Data virtualization


The ability to view, access, manipulate, and analyze data without the need to understand its physical format or location, and without having to move or copy it

Data virtualization
reduces ETL requests by


25-64%




Single view
across business




Consistent interface



Intelligent views

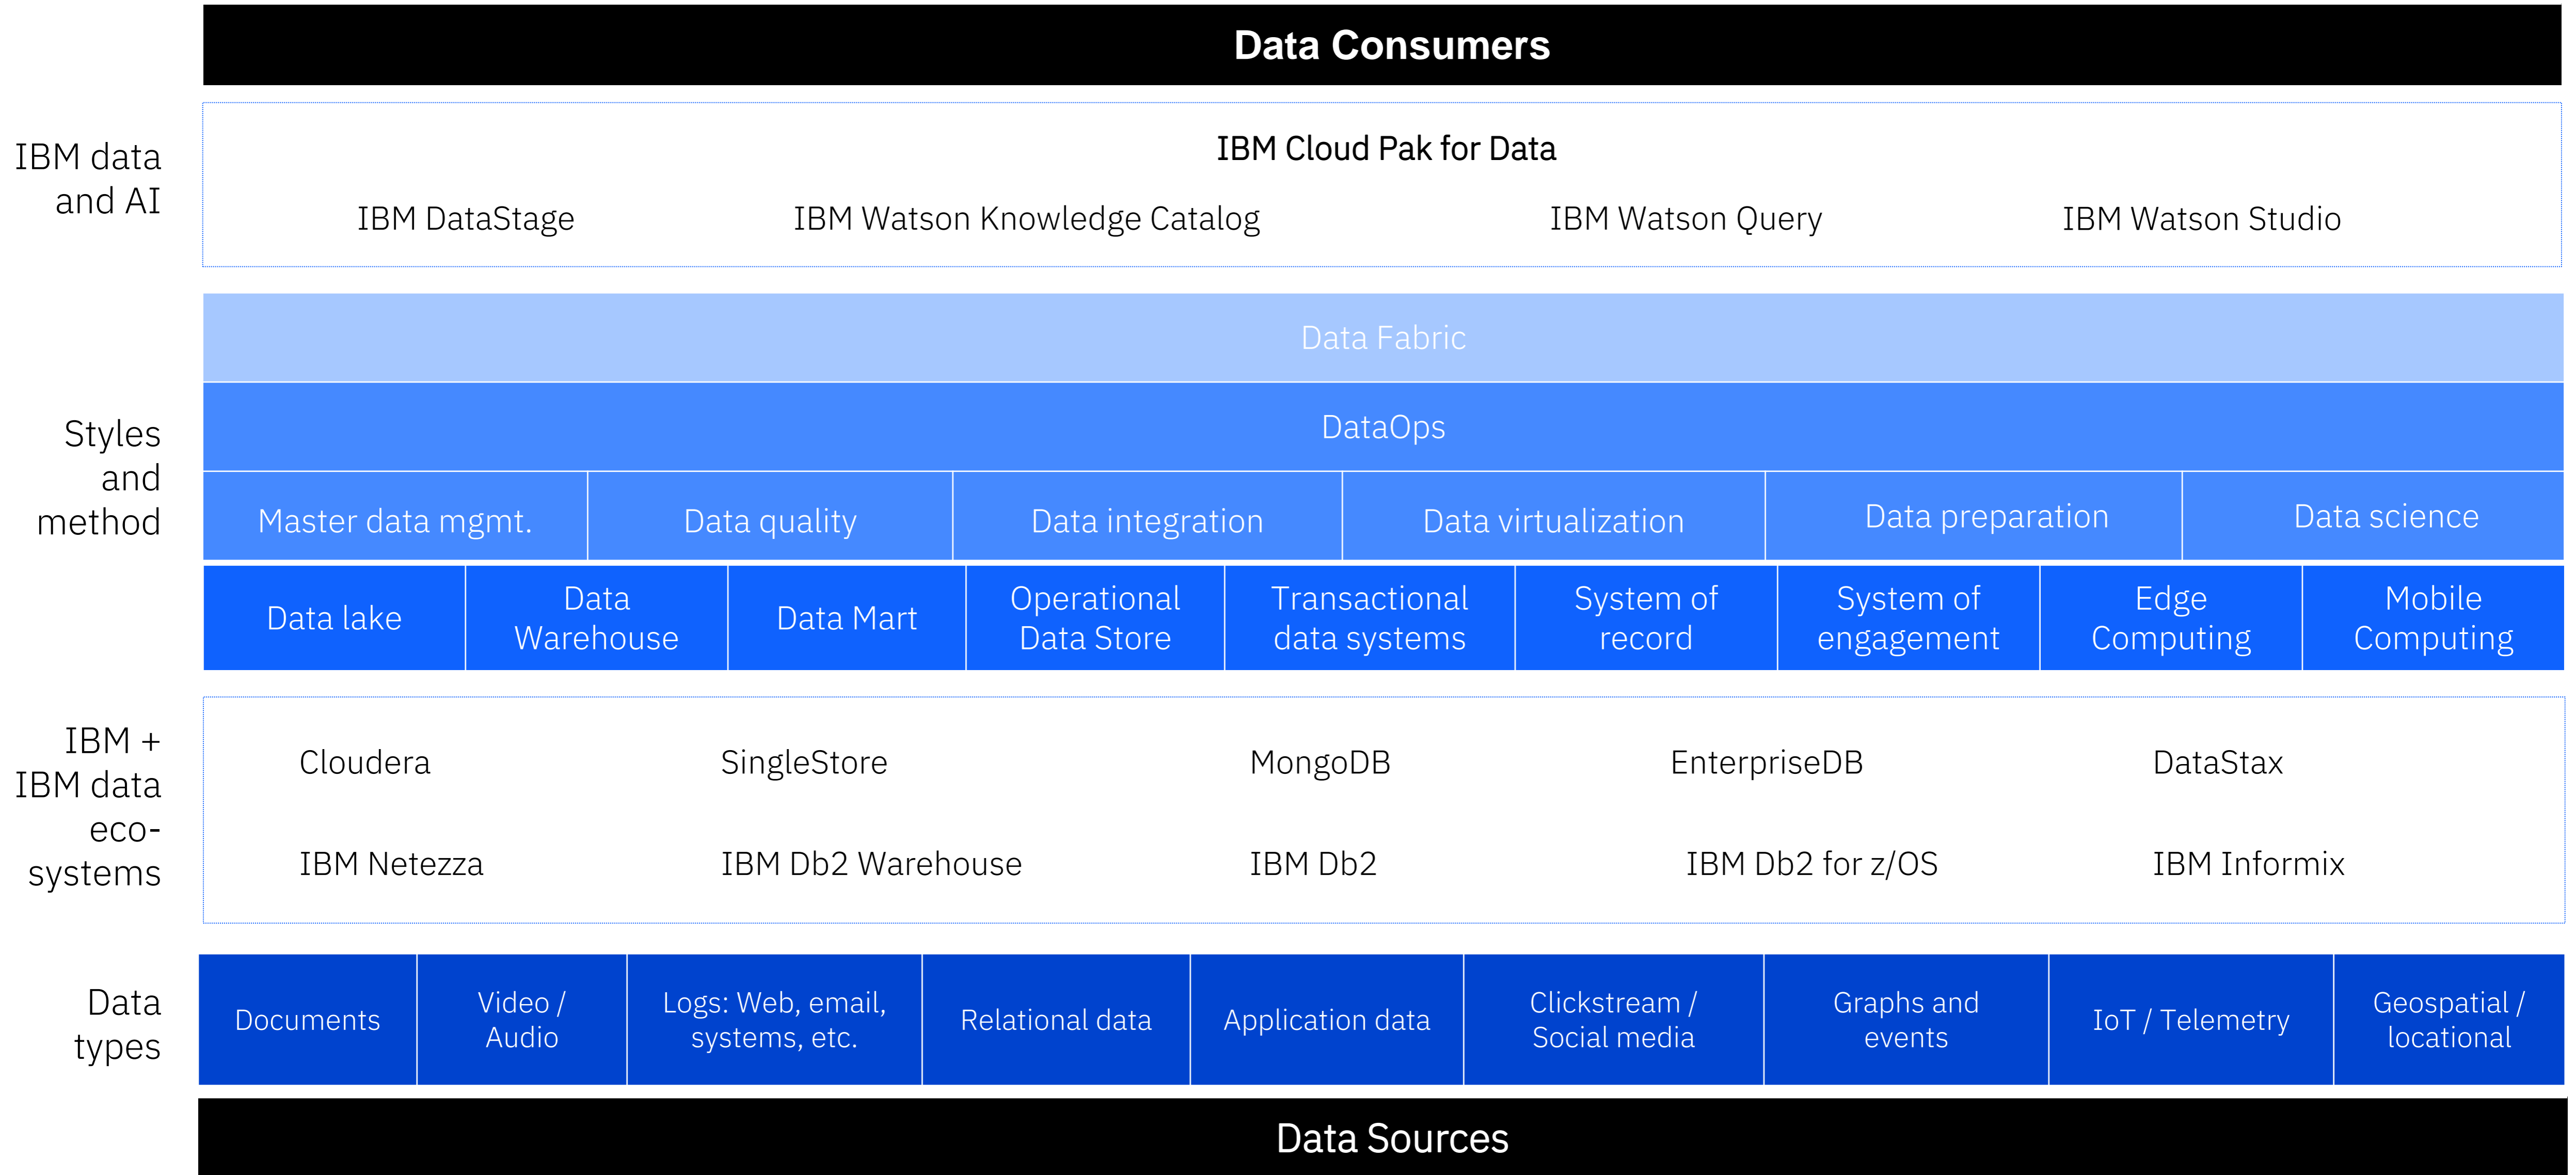


Real-time analytics
without moving data



Governance and security

IBM Data Management portfolio



Netezza Performance Server

enables us to make the most out of our data. By leveraging the solutions capabilities, we are able to reduce the time to generate key reports and **make better, more informed decisions** in line with our customer needs.

- *Leading European Bank*

“We ported a six-year-old custom-code application running thousands of ELT operations. **Ran first time without a hiccup and finished so fast** we thought it had failed.”

- *Leading National Concierge company*

“Our longest queries returned in a fraction of their original time, and saw **no concurrency drag** running hundreds at once.”

- *Customer from Healthcare Industry*

“The **speed is amazing**. Our jobs run faster, queries return in a flash, and the one console tells us everything”

- *National Entertainment Retailer*

Informix

has been offering TimeSeries and geo-spatial data longer than any other database on the market. In many cases for longer than competitive offerings have been in existence.

The experience that comes with this longevity is clear in the [depth and breadth of the capabilities](#) offered.



“Informix offers a **state-of-the-art** database management system, which is fast, reliable and very easy to maintain.”

- *Engineer,
Computer software industry*

“Informix is **flexible** enough to allow us to both easily do ad hoc reporting and make changes to our software as needed to meet clients’ needs.”

- *Software engineer,
Computer software industry*

“A **robust database product that is way ahead of the competition.** The incredible feature set and versatility makes it ideal as a back end workhorse supporting huge ERP systems and data warehouses, or as a small footprint database for internet of things applications, and everything in between.”

- *Senior database consultant,
IT and Services industry*


IBM Db2

is a **reliable and trusted** database software which handles billions of transactions each day sub-milliseconds response time highly available scales to handle 2,000% transactional growth securing the homeland of United States with 7x24 around the year.

- IT Leader/Manager
US Federal Government

Read more

- [Gartner Customer Choice Awards for 2022](#)
- [Db2 ranked #3 in G2 Summer Grid Report for Data Warehousing](#)



Happy
anniversary

The image features the text "Happy anniversary" written in a flowing, gold-colored cursive script. The word "Happy" is positioned above "anniversary". The text is set against a solid, vibrant blue background. The entire phrase is embellished with numerous white, multi-pointed starburst or sparkler-like effects, which are most prominent around the letters 'H', 'a', 'n', 'i', 'v', 'e', 'r', 's', 'a', 'r', 'y' and the 'y' at the end of the second line. The overall aesthetic is festive and celebratory.

Db2 highlights – 30 Years of Leading-edge Technology

Db2/2 V1



1993

DB2/6000 V1



1994

DB2 PE



1994

DB2 2.1



1995-1996

DB2 5-6



1997-1999

DB2 7-8



2001-2004

DB2 9-9.8



2006-2010

DB2 10.1



2012

DB2 10.5



2013

DB2 11.1



2016

Db2 11.5



2019

Db2 11.5.4



2020

DB2 11.5.5



2020

DB2 11.5.6



2021

DB2 11.5.7



2021

DB2 11.5.8



2022

Db2 11.5.9



2023

Db2 Next



2024

Db2 highlights – 30 Years of Leading-edge Technology

Db2/2 V1



1993

DB2/6000 V1



1994

DB2 PE



1994

DB2 2.1



1995-1996

DB2 5-6



1997-1999

DB2 7-8



2001-2004

1993: IBM DB2 for OS/2 – DB2/2

1994: IBM DB2/6000 Version 2 and DB2/2 Version 2

1994: IBM DB2 Parallel Edition for AIX/6000 Version 1

1996: DB2 Universal Database (UDB) Version 2

1998: IBM DB2 Universal Database Version 5.0

1999: IBM DB2 Universal Database Version 6.1

2000: IBM DB2 Universal Database Version 7.1

2001: IBM DB2 Universal Database Version 7.2

2002: IBM DB2 Universal Database V8.1 for Linux, UNIX, and Windows

2004: IBM DB2 Universal Database V8.2 for Linux, UNIX, and Windows

Db2 highlights – 30 Years of Leading-edge Technology



Db2/2 V1



1993

HADR

.NET, JDBC,
SQLJ, OLE

VARxxx,
xLOBs

DB2/6000 V1



1994

OS/2, AIX, Windows,
Linux, Solaris, HP-UX

Text Analytics

Shared-Nothing Scale-
out for OLAP (EEE/DPF)

DB2 PE



1994

Granular
Backups

Spatial Analytics

Sequence

Query Patroller (WLM)

DB2 2.1



1995-1996

Db2 Connect (DDCS)

Net Search Extender

Data Joiner (Federation)

DB2 5-6



1997-1999

Connection Pooling

Data Links (External Files)

LDAP

Unicode

DB2 7-8



2001-2004

AST/MQT

Mobile
Satellite

Triggers

Db2 highlights – 30 Years of Leading-edge Technology



Db2 highlights – 30 Years of Leading-edge Technology

2006: DB2 9.1 for Linux, Unix and Windows (aka Viper)

2008: DB2 9.5 for Linux, Unix and Windows

2009: DB2 9.7 for Linux, Unix and Windows

2010: DB2 9.8 (pureScale)

2012: DB2 10.1 for Linux, Unix and Windows

2013: DB2 10.5 for Linux, Unix and Windows

DB2 9-9.8



2006–2010

DB2 10.1



2012

DB2 10.5



2013

DB2 11.1



2016

Db2 11.5



2019

Db2 11.5.4



2020

2016: DB2 Database 11.1

2016: DB2 Database 11.1.1.1

2017: DB2 Database 11.1.2.2

2018: DB2 Database 11.1.3.3

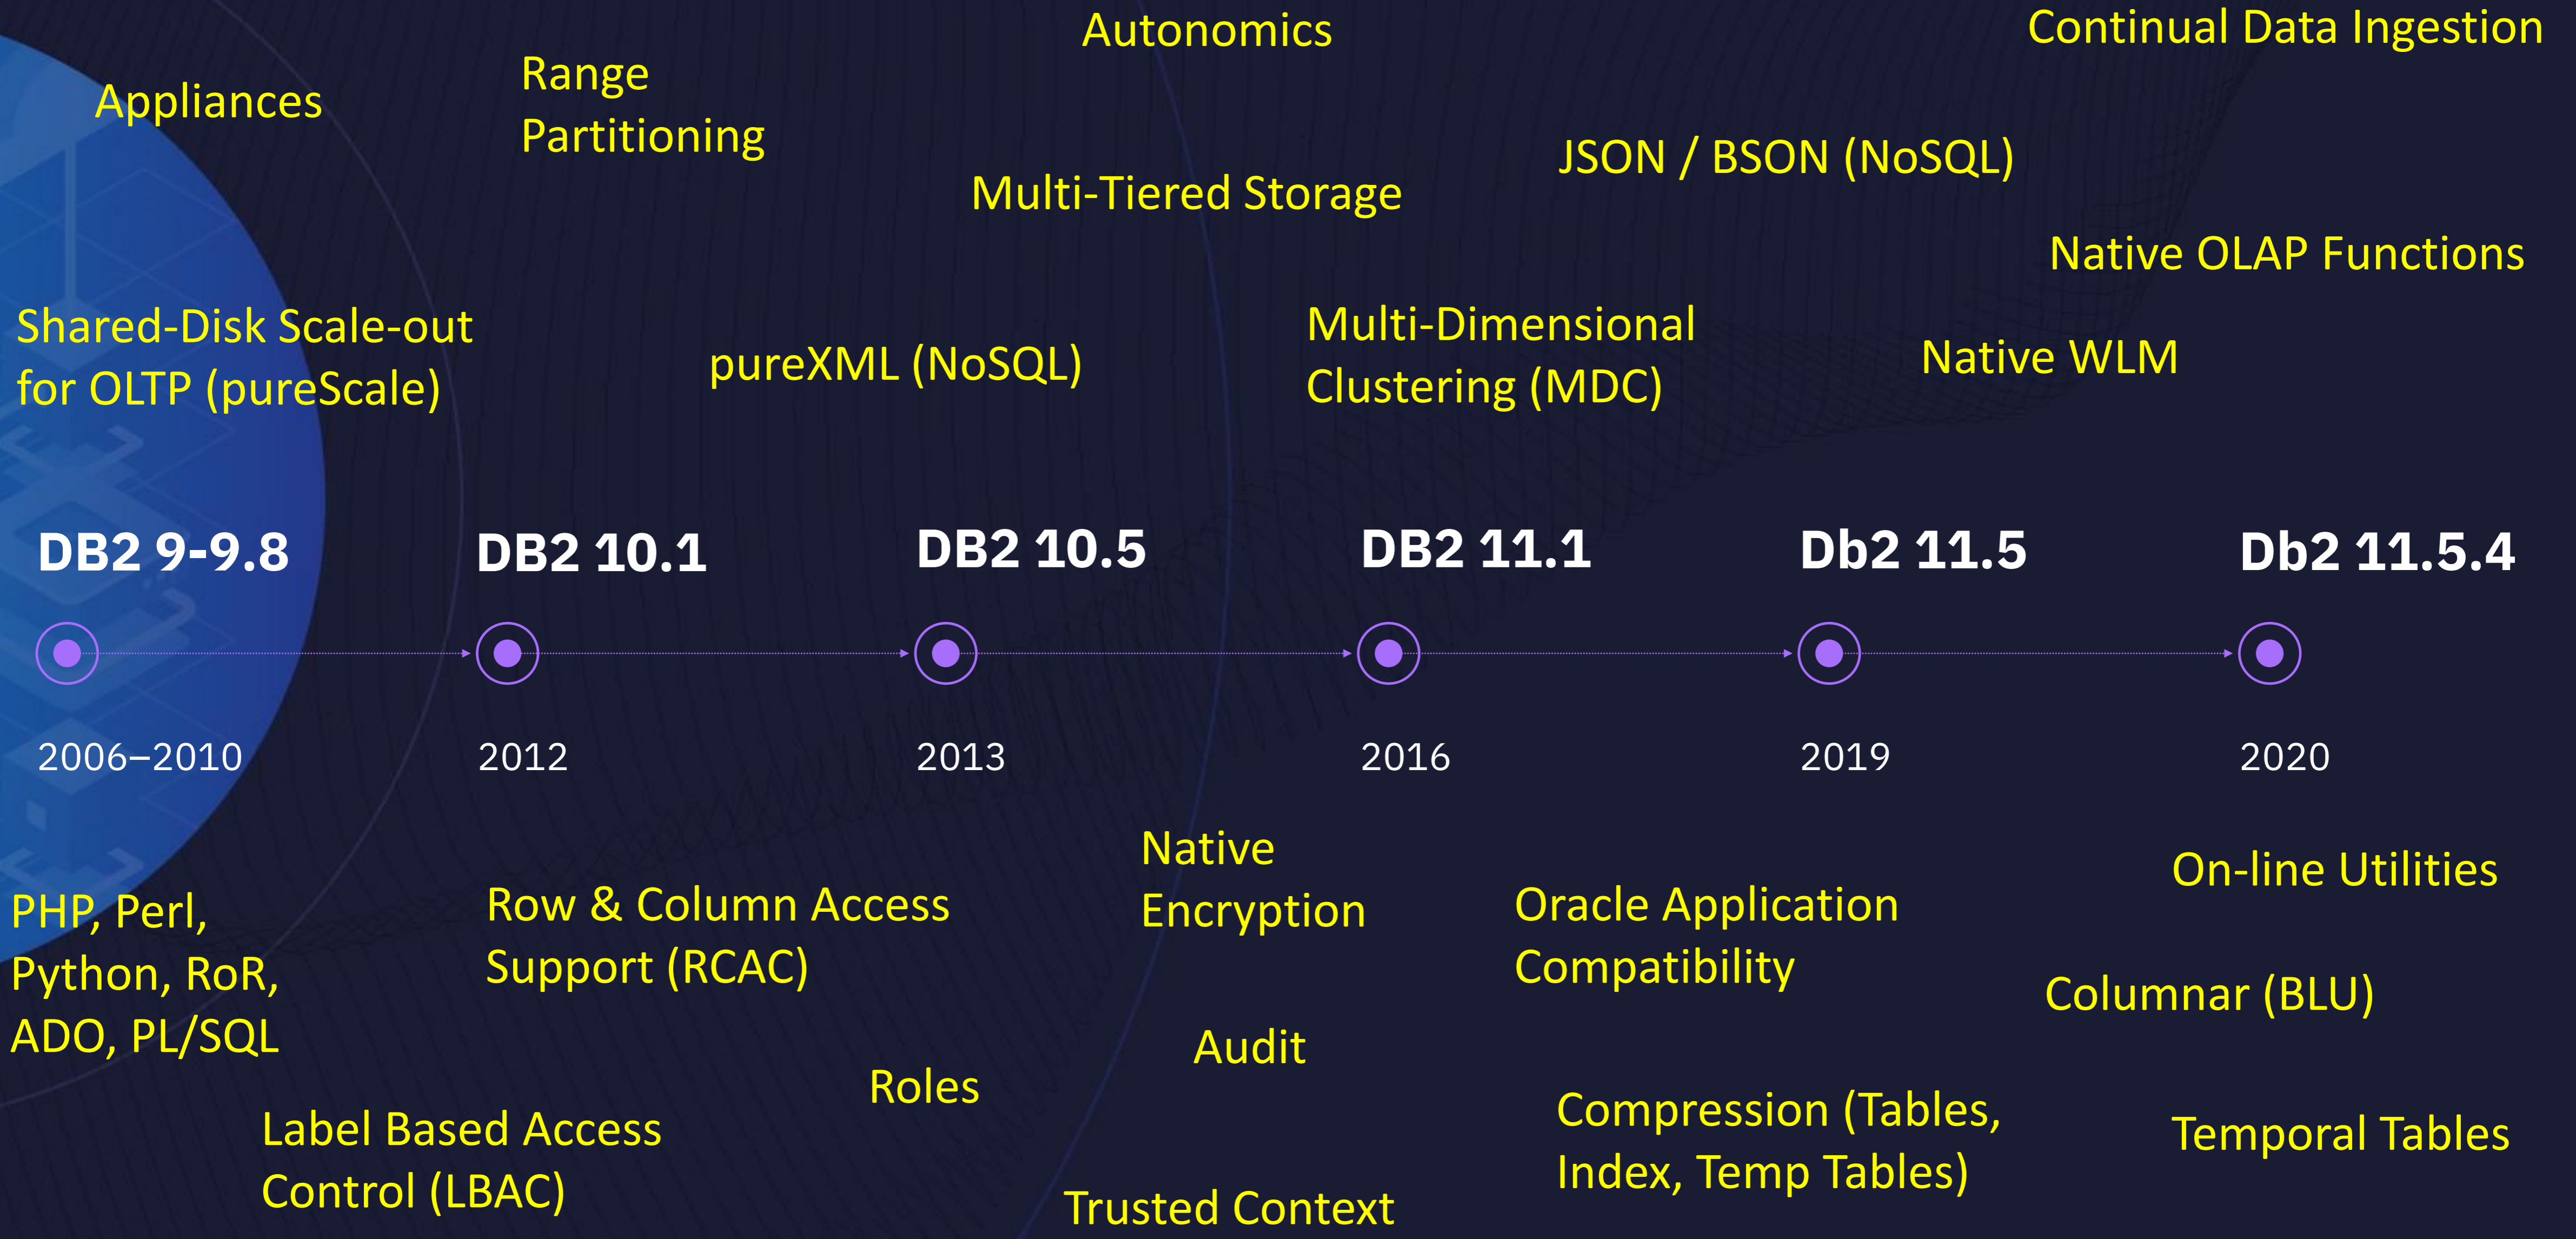
2018: DB2 Database 11.1.4.4

2019: Db2 Database 11.5

2019: Db2 Database 11.5.1, 11.5.2, 11.5.3

2020: Db2 Database 11.5.4

Db2 highlights – 30 Years of Leading-edge Technology



Db2 highlights – 30 Years of Leading-edge Technology



Db2 highlights – 30 Years of Leading-edge Technology

2020: Db2 11.5.5

2021: Db2 11.5.6

2021: Db2 11.5.7

2022: Db2 11.5.8

2023: Db2 11.5.9

.....

and more to come !!

DB2 11.5.5


DB2 11.5.6

DB2 11.5.7

DB2 11.5.8

Db2 11.5.9

Db2 Next



2020

2021

2021

2022

2023

2024

Db2 highlights – 30 Years of Leading-edge Technology

NZ (Postgres)
Compatibility

Db2 on Cloud
(DBaaS)

Graph (NoSQL)

NewSQL

Machine Learning
Optimizer

Event Processing

Data Virtualization

In-Db2 ML

Advanced Log Space
Management

RHOS Support

Schema Level
Security

Adaptive Workload
Management (AWLM)

Schema Level Recovery

DB2 11.5.5

DB2 11.5.6

DB2 11.5.7

DB2 11.5.8

Db2 11.5.9

Db2 Next

2020

2021

2021

2022

2023

2024

Containerization

HTAP

Mixed Workloads

Blockchain Query

External Tables

Native Cloud Object
Store Support

RESTful

pacemaker

Multi-Tenancy

Native Hadoop Support

Db2 highlights – 30 Years of Leading-edge Technology



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

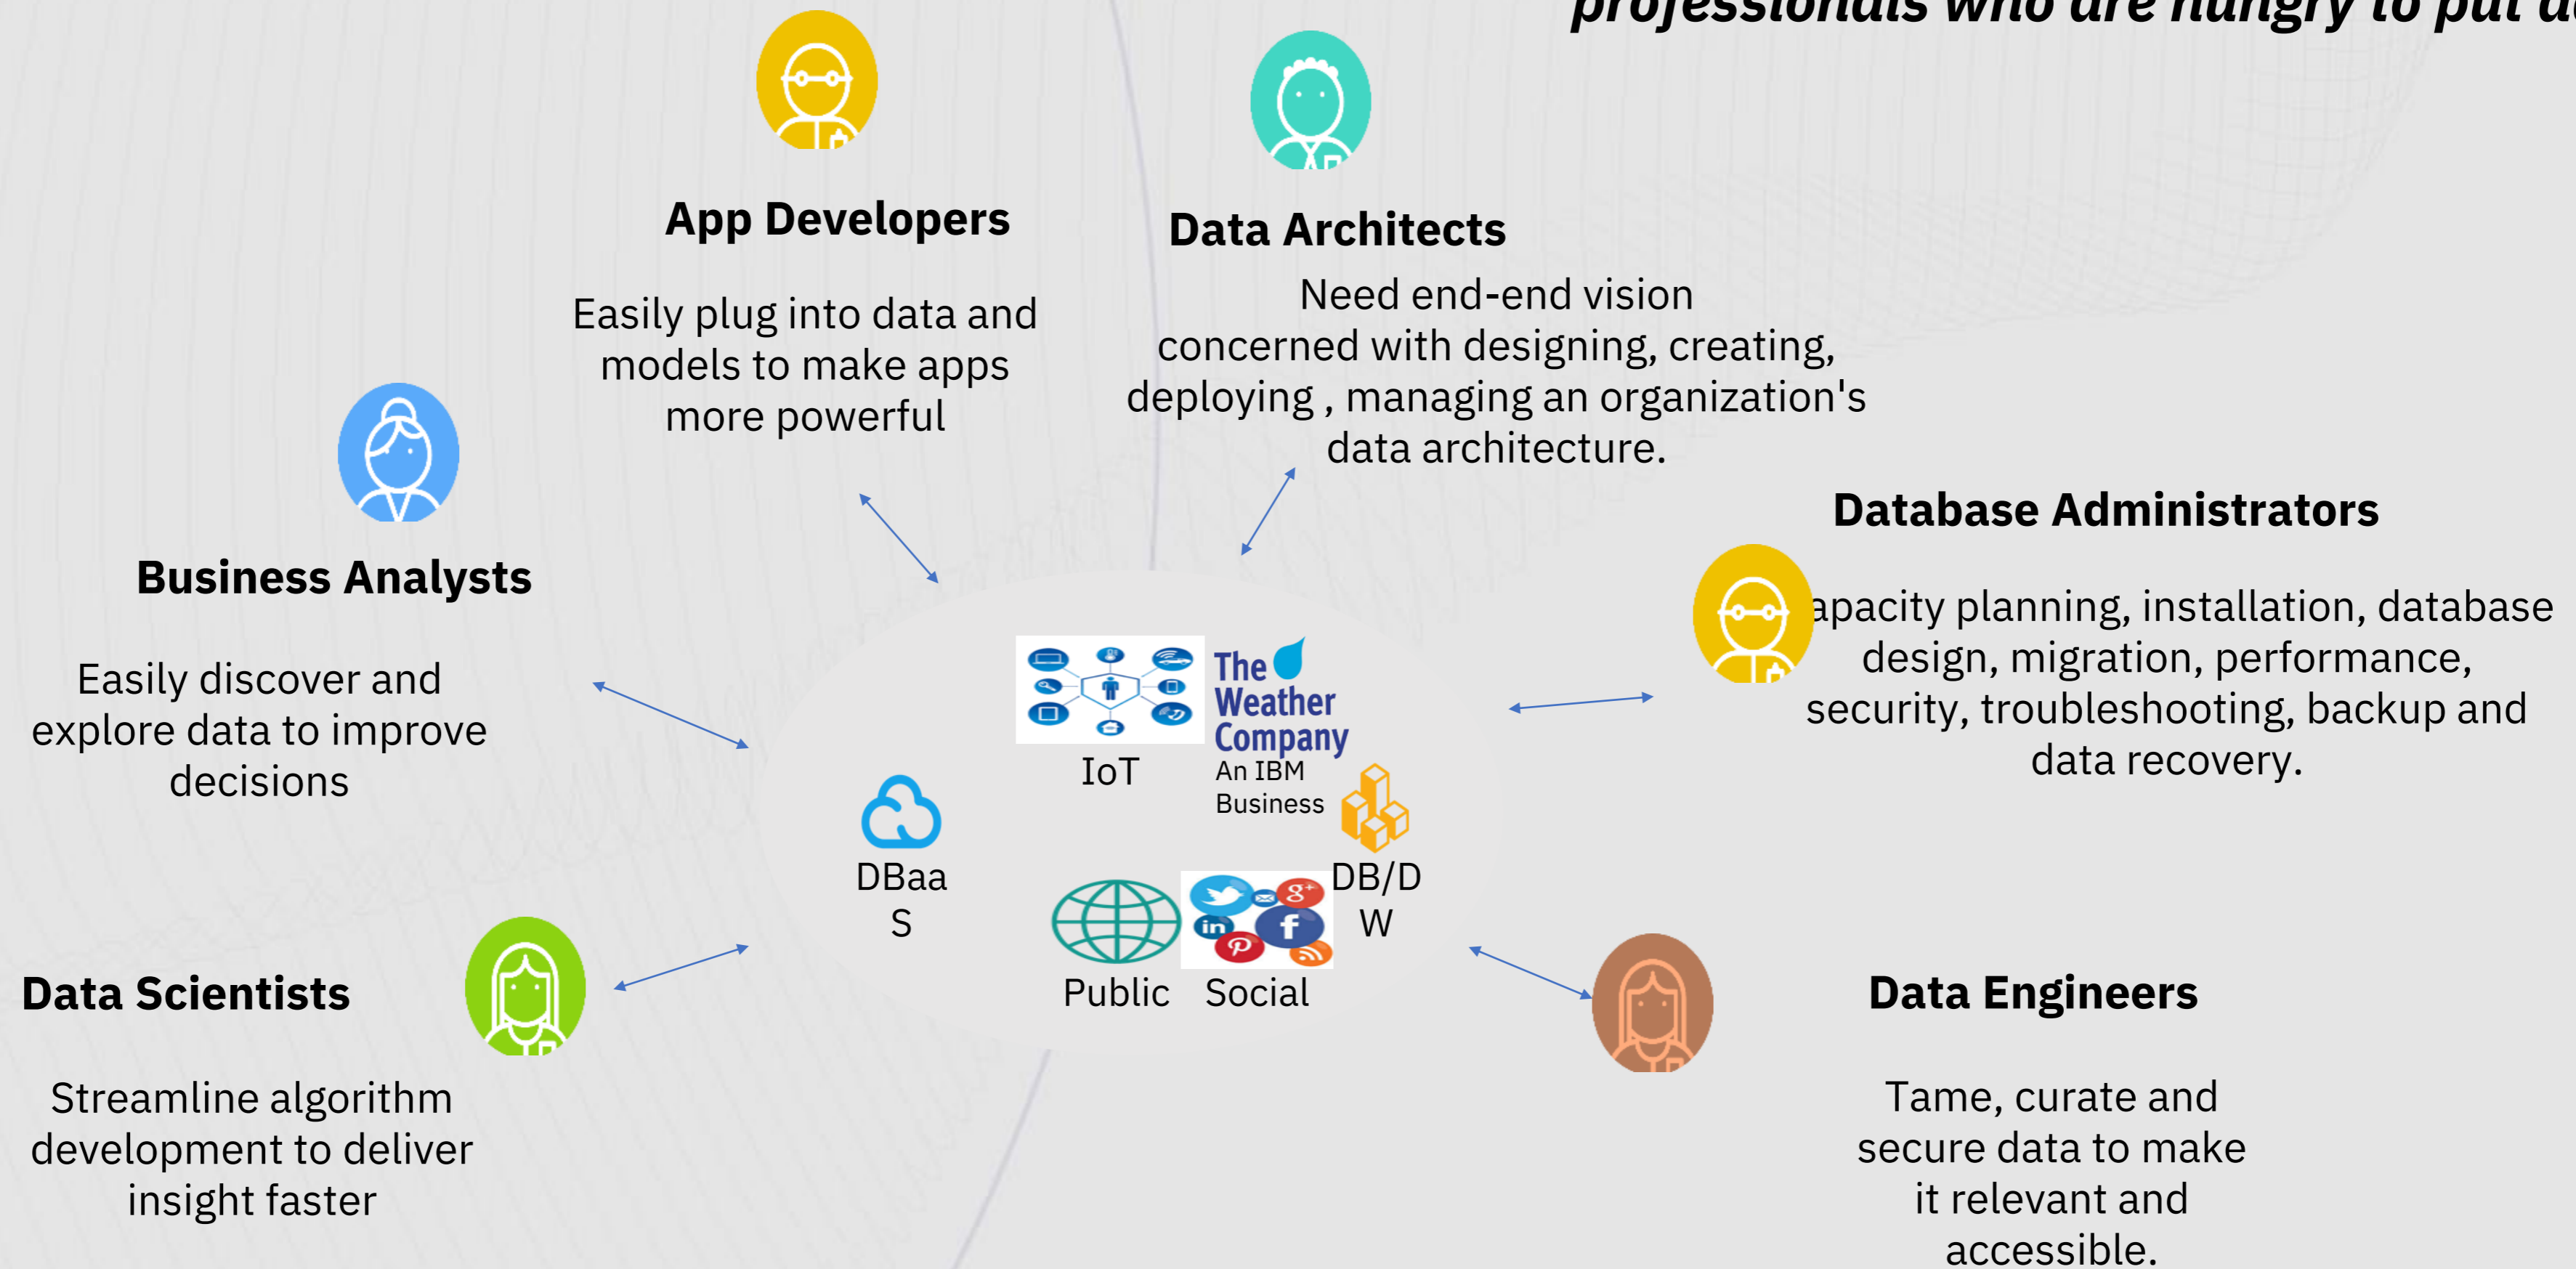
Db2 BigSQL

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



Evolving the Needs of All Data Professionals

As data maturity increases, so does the number of data professionals who are hungry to put data to work



Db2 V11.5 – Balance between **foundation** and **modernization**



Rock Solid Database



Infuse AI



Containers



Enterprise
Readiness

Making Data Simple

- *Performant*
- *Secure*
- *Available*
- *Automated administration & monitoring*



Cost Savings

- *ML optimizer*
- *Adaptive workload management*

Making Data Accessible



Modern
Development

- *Developer-friendly*
- *Multi-model*




Consumability

- *Deploy across multiple form factors*
- *Multi Cloud and Containers*

The AI Ladder

A prescriptive approach to accelerating the journey to AI




AI

INFUSE – Operationalize AI with trust and transparency

ANALYZE – Scale insights with AI everywhere

ORGANIZE – Create a trusted analytics foundation

COLLECT – Make data simple and accessible



Data of every type, regardless of where it lives




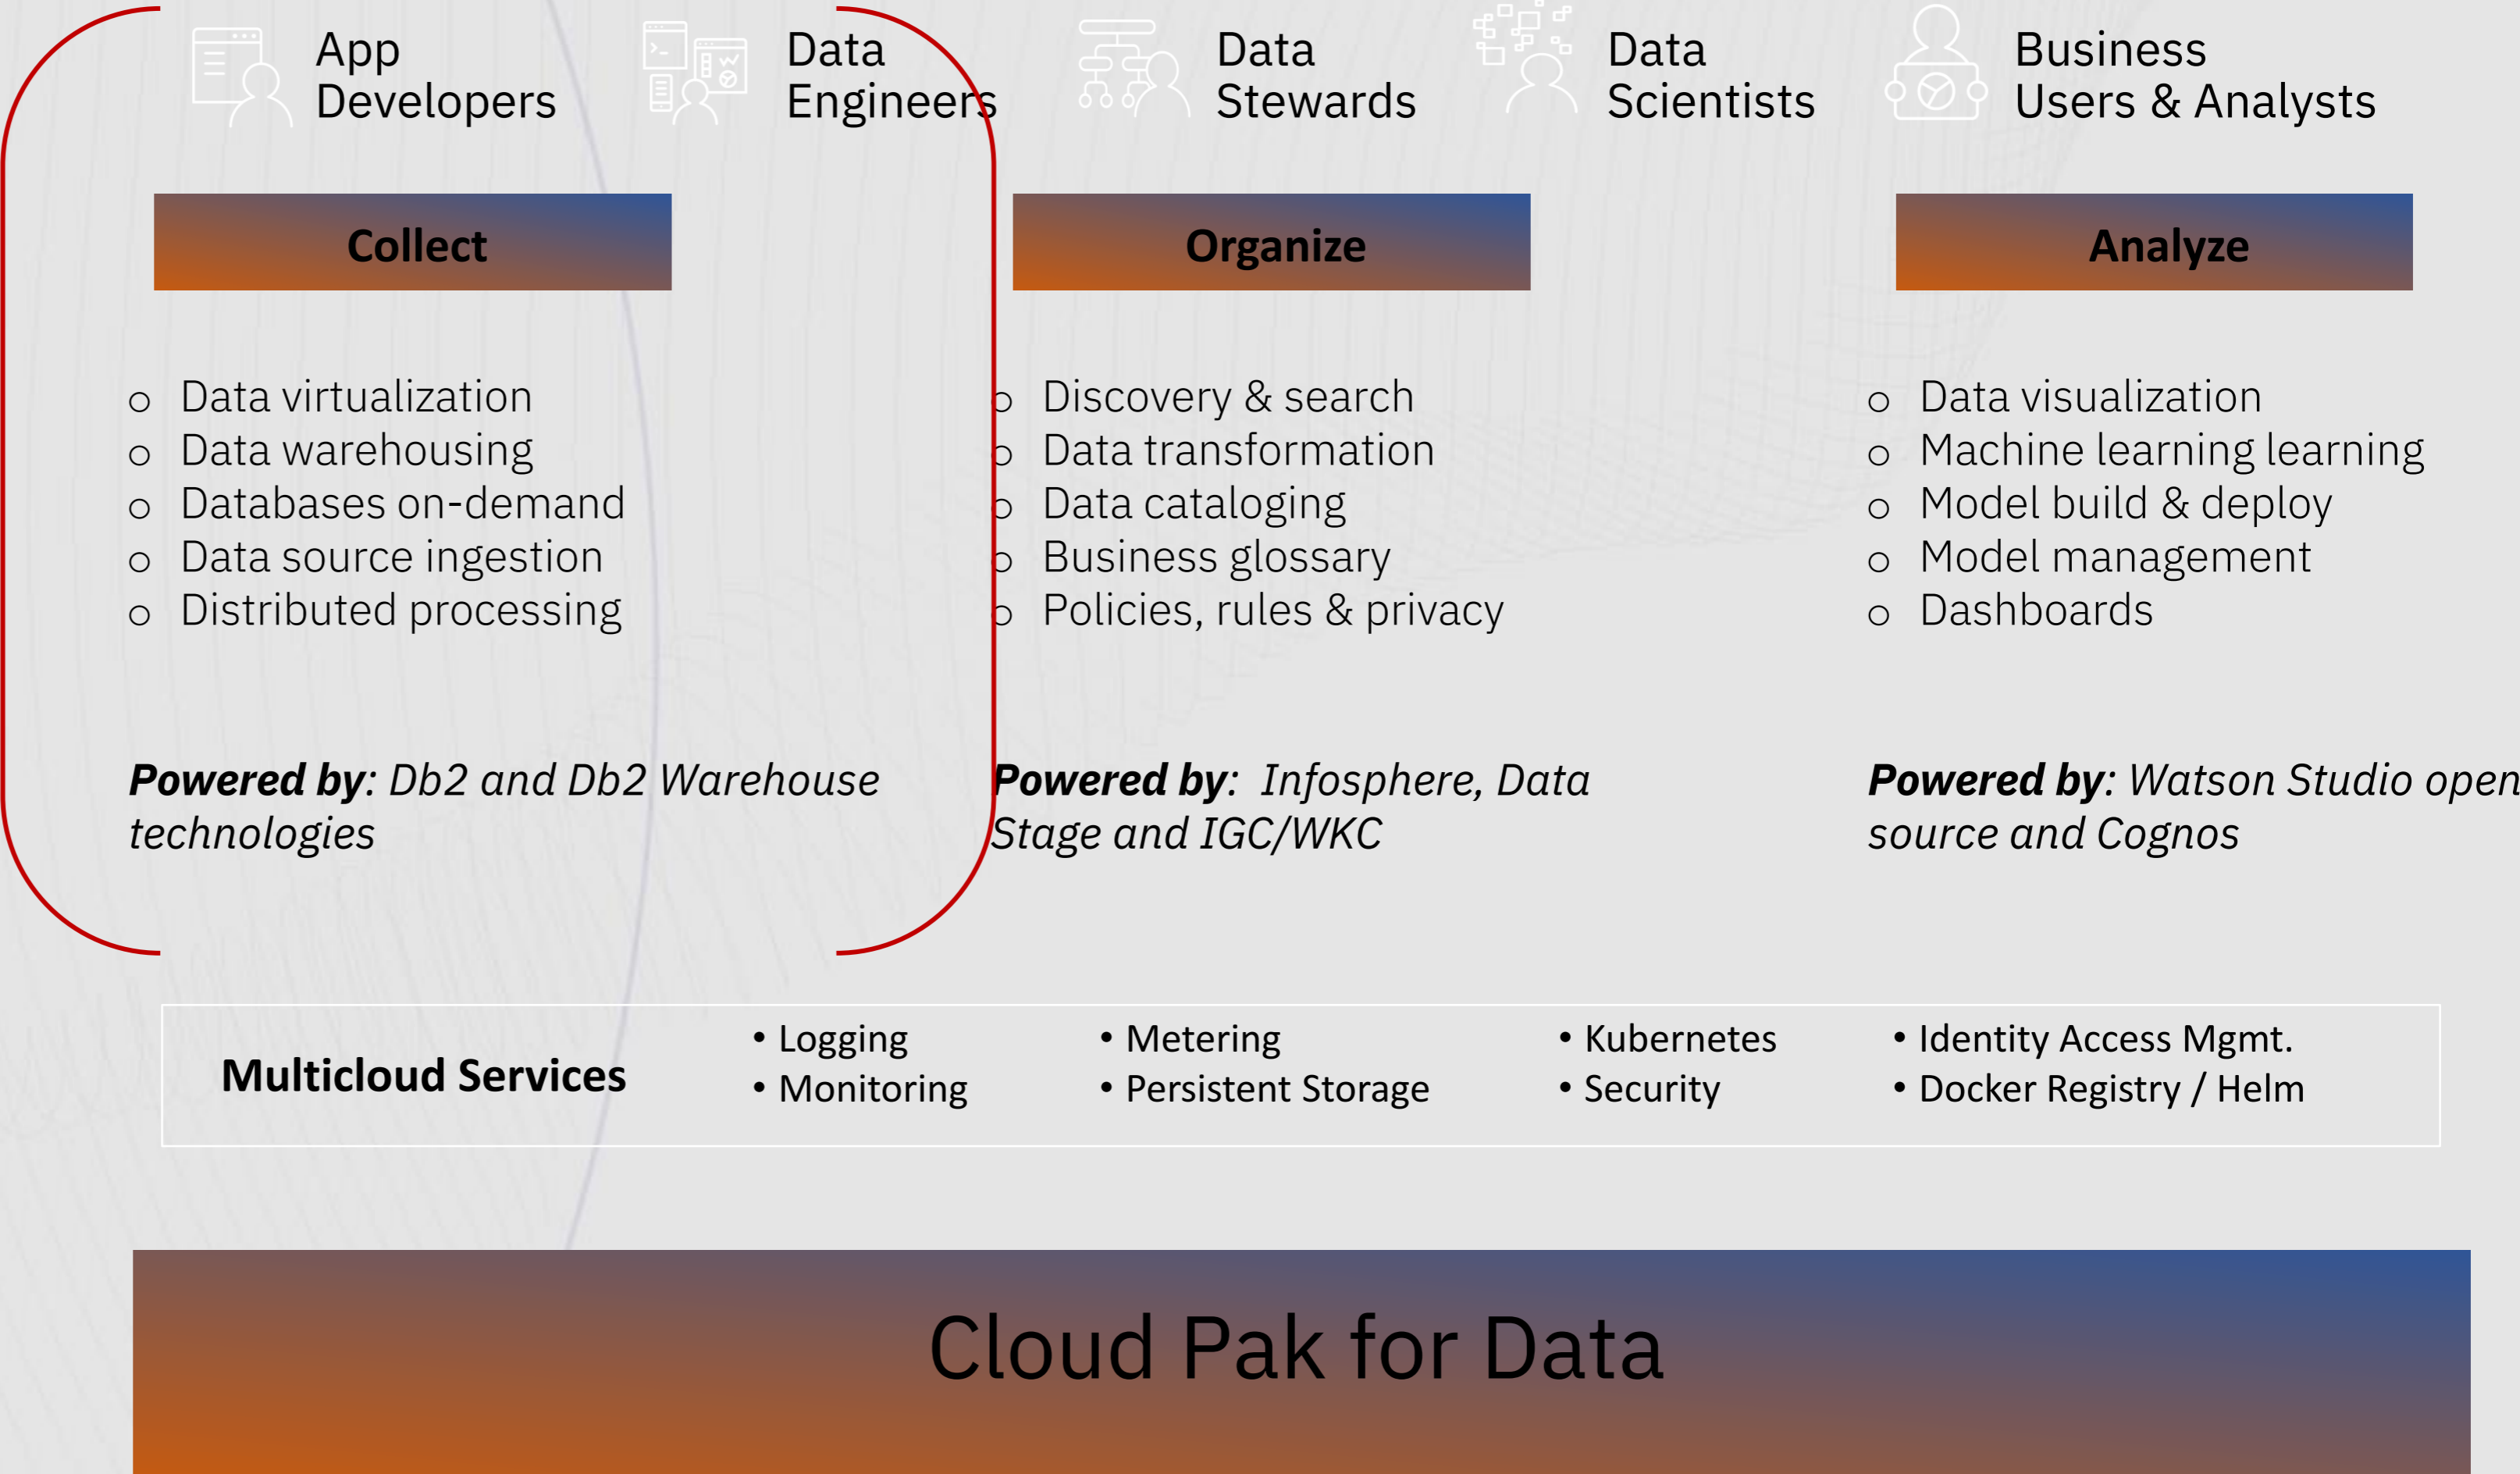
Db2 V11.5 – Cornerstone of “Collect” in Cloud Pak for Data

Foundational “out of the box” multi-cloud Data & AI services

The Ladder to AI



Infuse AI



Data Server Day



Les King
Director, Data Solutions
lking@ca.ibm.com
September 2023