

ioblend

MAKE YOUR DATA ESTATE
STATE-OF-THE-ART

try now @ www.ioblend.com

WHAT IS IOBLEND?

IOBLEND IS A NEW, ENTERPRISE REAL-TIME DATA INTEGRATION SOLUTION WITH AUTOMATED DATAOPS CAPABILITY IN ITS DNA.

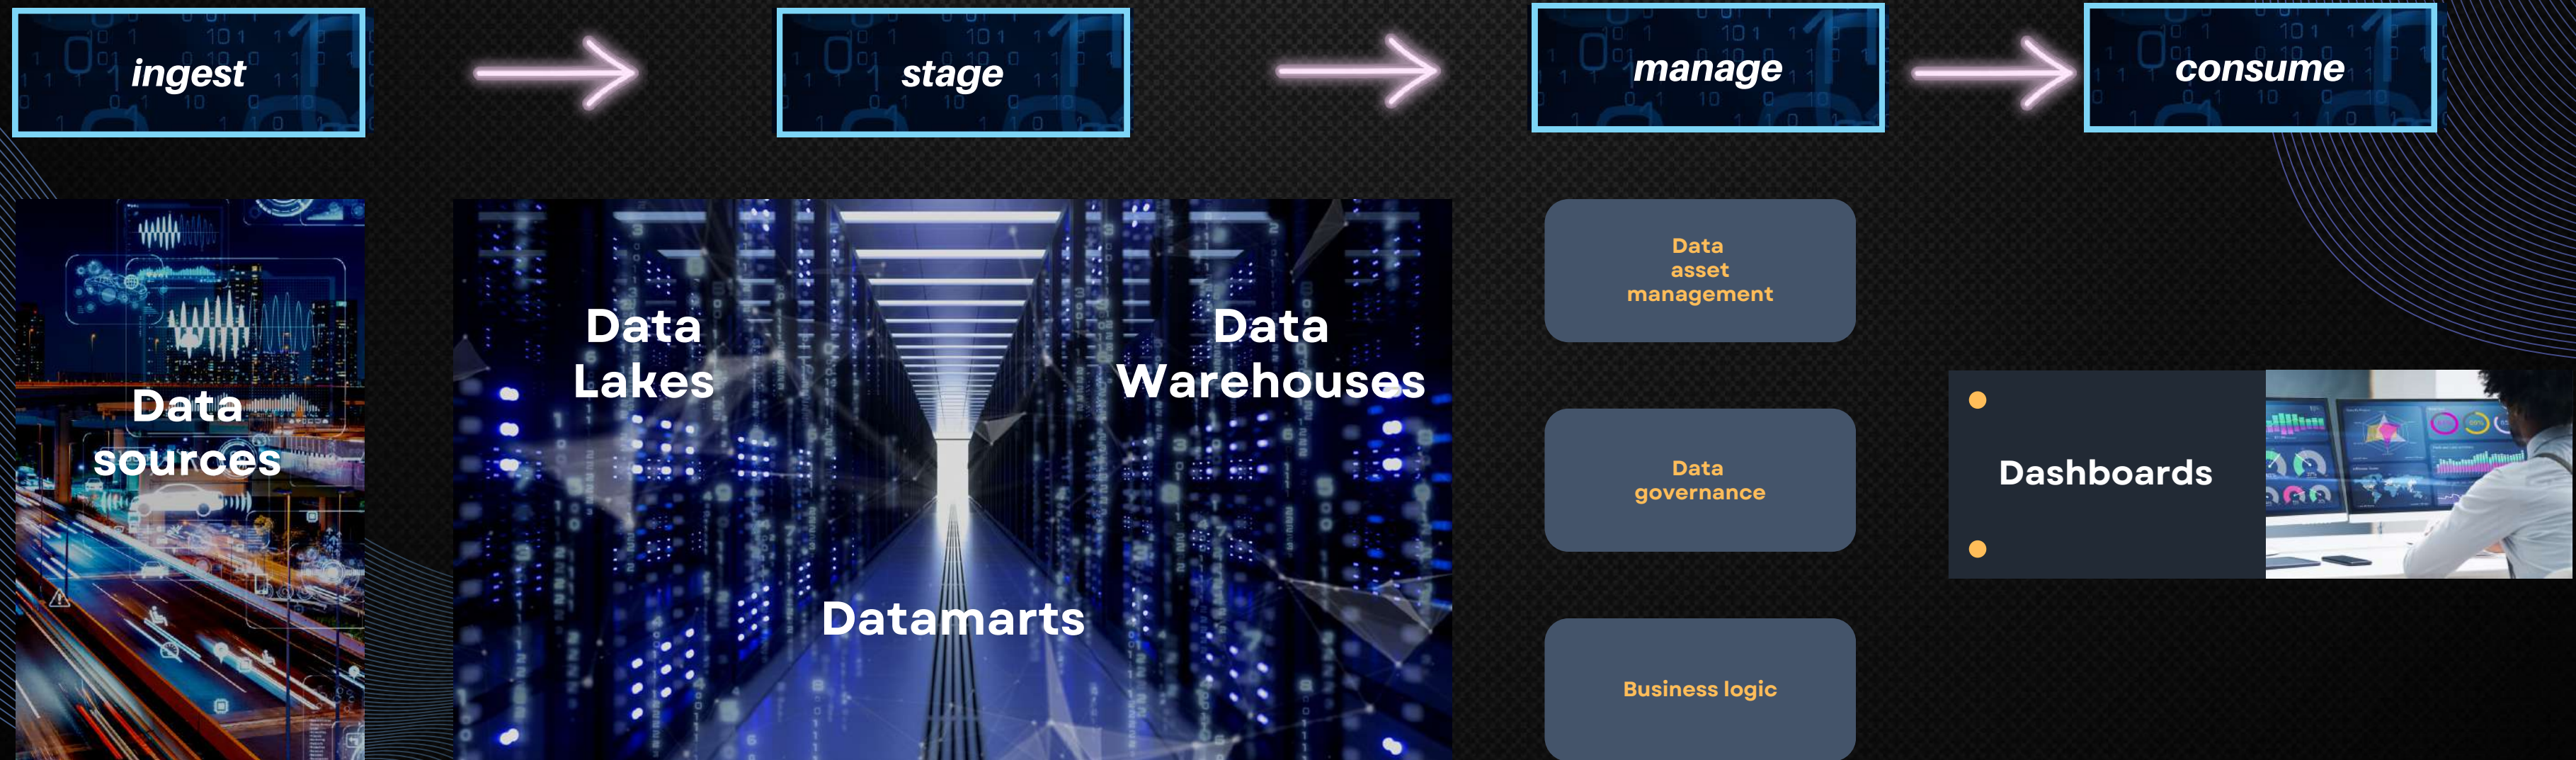
- REAL-TIME, PRODUCTION GRADE, MANAGED APACHE SPARK™ DATA PIPELINES IN MINUTES, USING ONLY SQL
- MASSIVELY ACCELERATED DATA MIGRATION PROJECTS FROM ON-PREM TO CLOUD OR CLOUD(S) TO CLOUD DUE TO LOW CODE / NO CODE DEVELOPMENT AND AUTOMATED DATA MANAGEMENT AND GOVERNANCE OF DATA WHILE IN-TRANSIT
- EASY INTEGRATION OF STREAMING (TRANSACTIONAL EVENT) AND BATCH DATA VIA KAPPA ARCHITECTURE AND AUTOMATED DATAOPS TO POWER YOUR OPERATIONAL ANALYTICS AND MLOPS
- ENABLES ROBUST AND COST EFFECTIVE DELIVERY OF BOTH CENTRALISED AND FEDERATED DATA ARCHITECTURES
- LOW LATENCY, MASSIVELY PARALLELISED DATA PROCESSING, OFFERING INCREDIBLE SPEEDS (>10M TRANSACTIONS PER SEC)
- REAL-TIME INTEGRATION WITH SNOWFLAKE AND MICROSOFT AZURE PRODUCTS

TODAY'S DATA
ANALYTICS ARE
PREDOMINANTLY
BATCH BASED,
DESIGNED FOR
PERIODIC
DASHBOARDS



DATA IS INGESTED
AT FIXED PERIODS
AND STAGED IN A
LAKE/WAREHOUSE
BEFORE
CONSUMPTION VIA
ANALYTICAL TOOLS

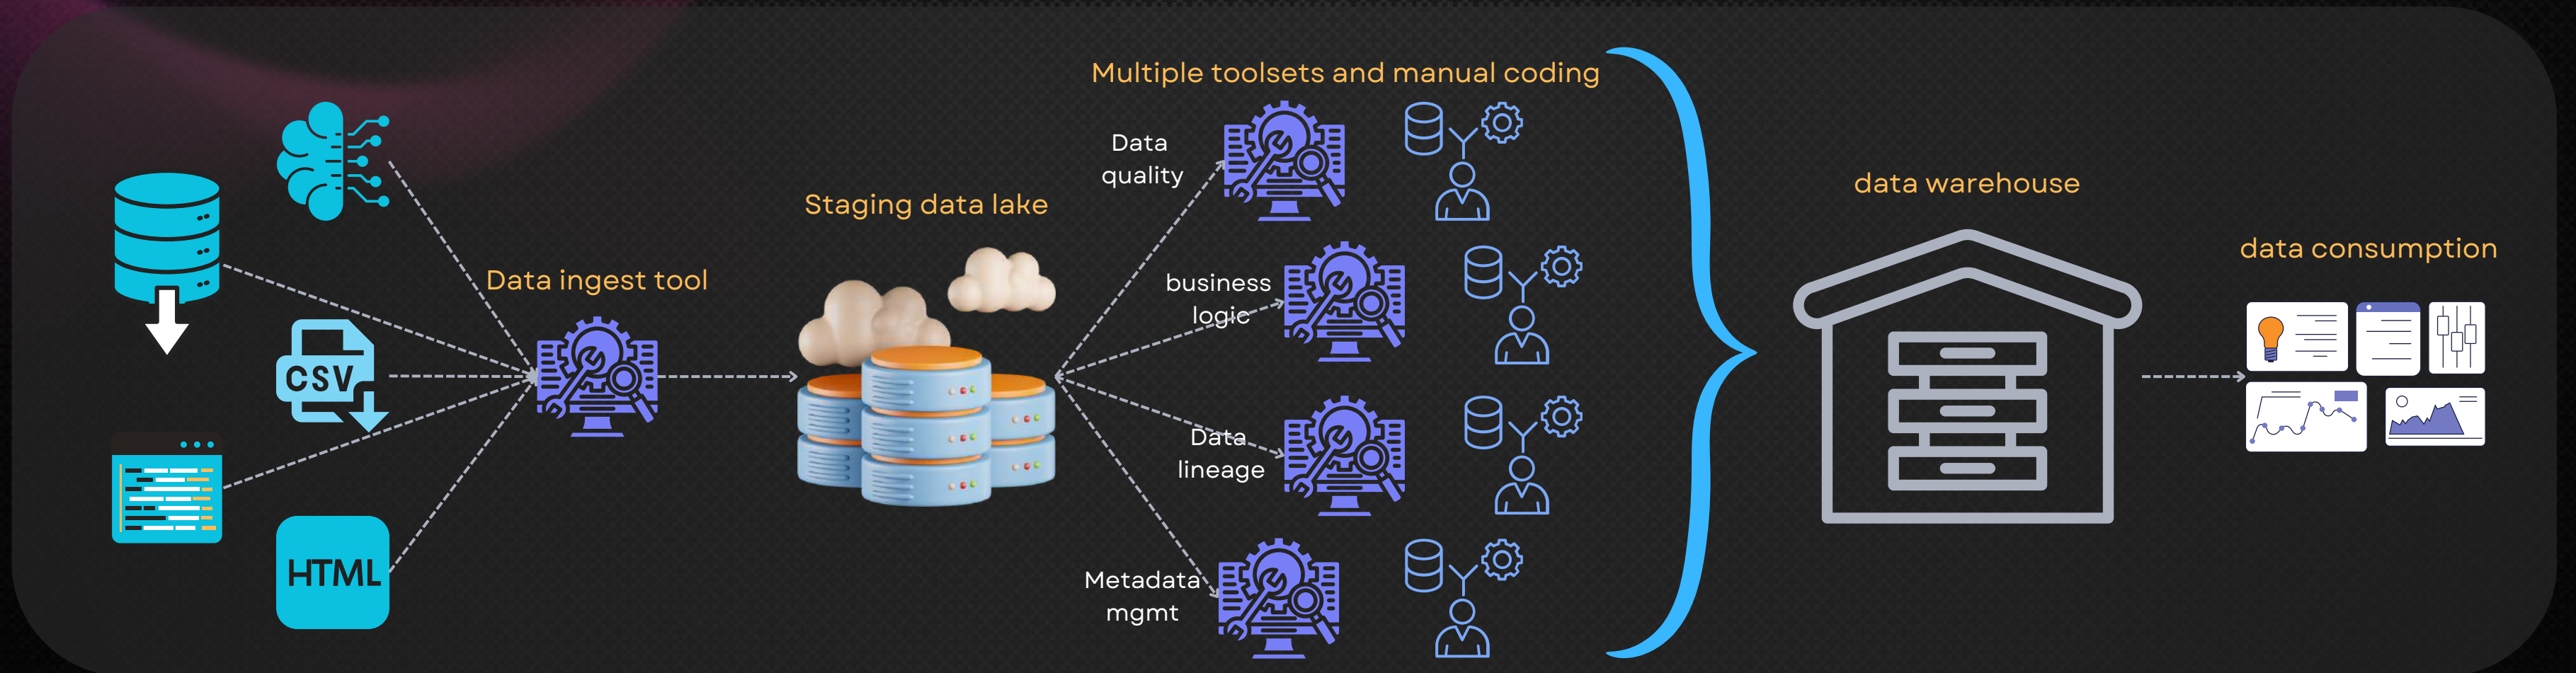
Current state of play is aimed at experimental analytics not production data integration



IT WORKS WELL FOR FINDING EXPLORATORY DATA INSIGHTS BUT BECOMES VERY INEFFICIENT WHEN ATTEMPTING TO SCALE IN PRODUCTION

COMPANIES ATTEMPTING
TO SCALE THEIR
ANALYTICS TO
PRODUCTION LEVEL FACE
A STEEP TECHNICAL
CHALLENGE

Prevalent analytics stack



ARCHITECTURE IS CENTRALISED AROUND AN ON-PREM OR CLOUD WAREHOUSE

INVOLVES MULTIPLE TECHNOLOGIES THAT ARE ADDED ON A MODULAR BASIS

EXTENSIVE ENGINEERING EFFORT TO SCALE, MAINTAIN AND MANAGE

Locked value

BUSINESS DATA IS VERY RICH AND COMES FROM A MULTITUDE OF SOURCES BUT IS OFTEN SILOED, DIFFICULT TO ACCESS AND TENDS TO BE MANUALLY COLLATED AND CURATED BY SMES

No single source of the truth

ORGANISATIONS OFTEN WORK WITH SEVERAL SEPARATE POINT SYSTEMS, AT TIMES RING-FENCED, THAT DO NOT "TALK" TO EACH OTHER AND DO NOT PROVIDE STANDARDISED ACCESS TO THE VALUABLE DATA FOR ANALYSIS - NO SINGLE SOURCE OF THE TRUTH!

Patchwork of solutions

DEPENDENCY ON A MYRIAD OF PROCESSES, TOOLS AND WORKAROUNDS TO ACCESS AND WORK WITH DATA, OFTEN IN THE FORM OF VARIOUS SPREADSHEETS OR PRIVATE SQL TABLES WITH NON-STANDARD FORMATS.

Data quality issues

ANALYSTS SPEND SIGNIFICANT TIME AND EFFORT LOCATING, PREPARING AND VALIDATING RELEVANT DATA BEFORE THEY ARE ABLE TO FOCUS ON INSIGHTS

No true real time data

LIVE STREAMING DATA PROJECTS ARE ESPECIALLY DIFFICULT AND COSTLY (IF NOT IMPOSSIBLE) TO IMPLEMENT AND MANAGE WITH THE CURRENT TOOLS. DATA QUALITY AND GOVERNANCE OF THE STREAMING DATA ARE A NIGHTMARE TO CODE

Slow implementations

MORE COMPLEX DATA PROJECTS, LIKE ATTEMPTING PRODUCTIONISED DATA INTEGRATION OR DEVELOPING A NEW AUTOMATED APP/DASHBOARD/REPORT, REQUIRE EXTENSIVE IT INVOLVEMENT, TIME AND COST. OFTEN, SUCH INITIATIVES NEED TO GO THROUGH LAYERS OF SCOPING, PROCUREMENT, BUDGET APPROVALS, RESOURCE SIGN-OFFS, AND LENGTHY DEVELOPMENT TIMES

YOUR FULL DATA CAPABILITY
CANNOT BE UNLOCKED WITH
THE STATIC ANALYTICS
TECHNOLOGIES AND LEGACY
PROCESSES

YOU ARE **MISSING** VALUABLE
OPPORTUNITIES TO IMPROVE
REVENUE, REDUCE COSTS AND
ENHANCE OPERATIONAL
EFFICIENCY

MAKING DATA INSIGHTS
WORK IN LIVE OPERATIONAL
SETTINGS WITHOUT THE
RIGHT TECHNOLOGIES IS
VERY COMPLEX AND
EXPENSIVE

Revolution in data engineering tech

INTRODUCING THE WORLD'S MOST
POWERFUL PRODUCTION-GRADE
DATA INTEGRATION SOLUTION

IOBLEND IS AN ADVANCED AUTOMATED
REAL TIME ETL/ELT SOLUTION THAT
GREATLY SIMPLIFIES ENGINEERING,
MANAGEMENT AND GOVERNANCE OF
ENTERPRISE DATA

WE ARE THE NEXT GENERATION IN DATA
INTEGRATION TECHNOLOGIES

What makes IOblend special?

IOBLEND IS AN END-TO-END ENTERPRISE DATA INTEGRATION SOLUTION OFFERING **TRUE DATAOPS CAPABILITY** BUILT INTO ITS DNA. IOBLEND IS AN ADVANCED ETL PRODUCT FOR BUILDING AND MANAGING DATA ESTATES WITH PRODUCTION-GRADE DATAFLOWS.

POWERED BY OUR PROPRIETARY ENGINE (**KAPPA ARCHITECTURE**) AND UTILISING THE POWER AND EFFICIENCY OF **APACHE SPARK™**, IOBLEND TAKES DATA ENGINEERING AND DATA MANAGEMENT TO THE NEXT LEVEL.

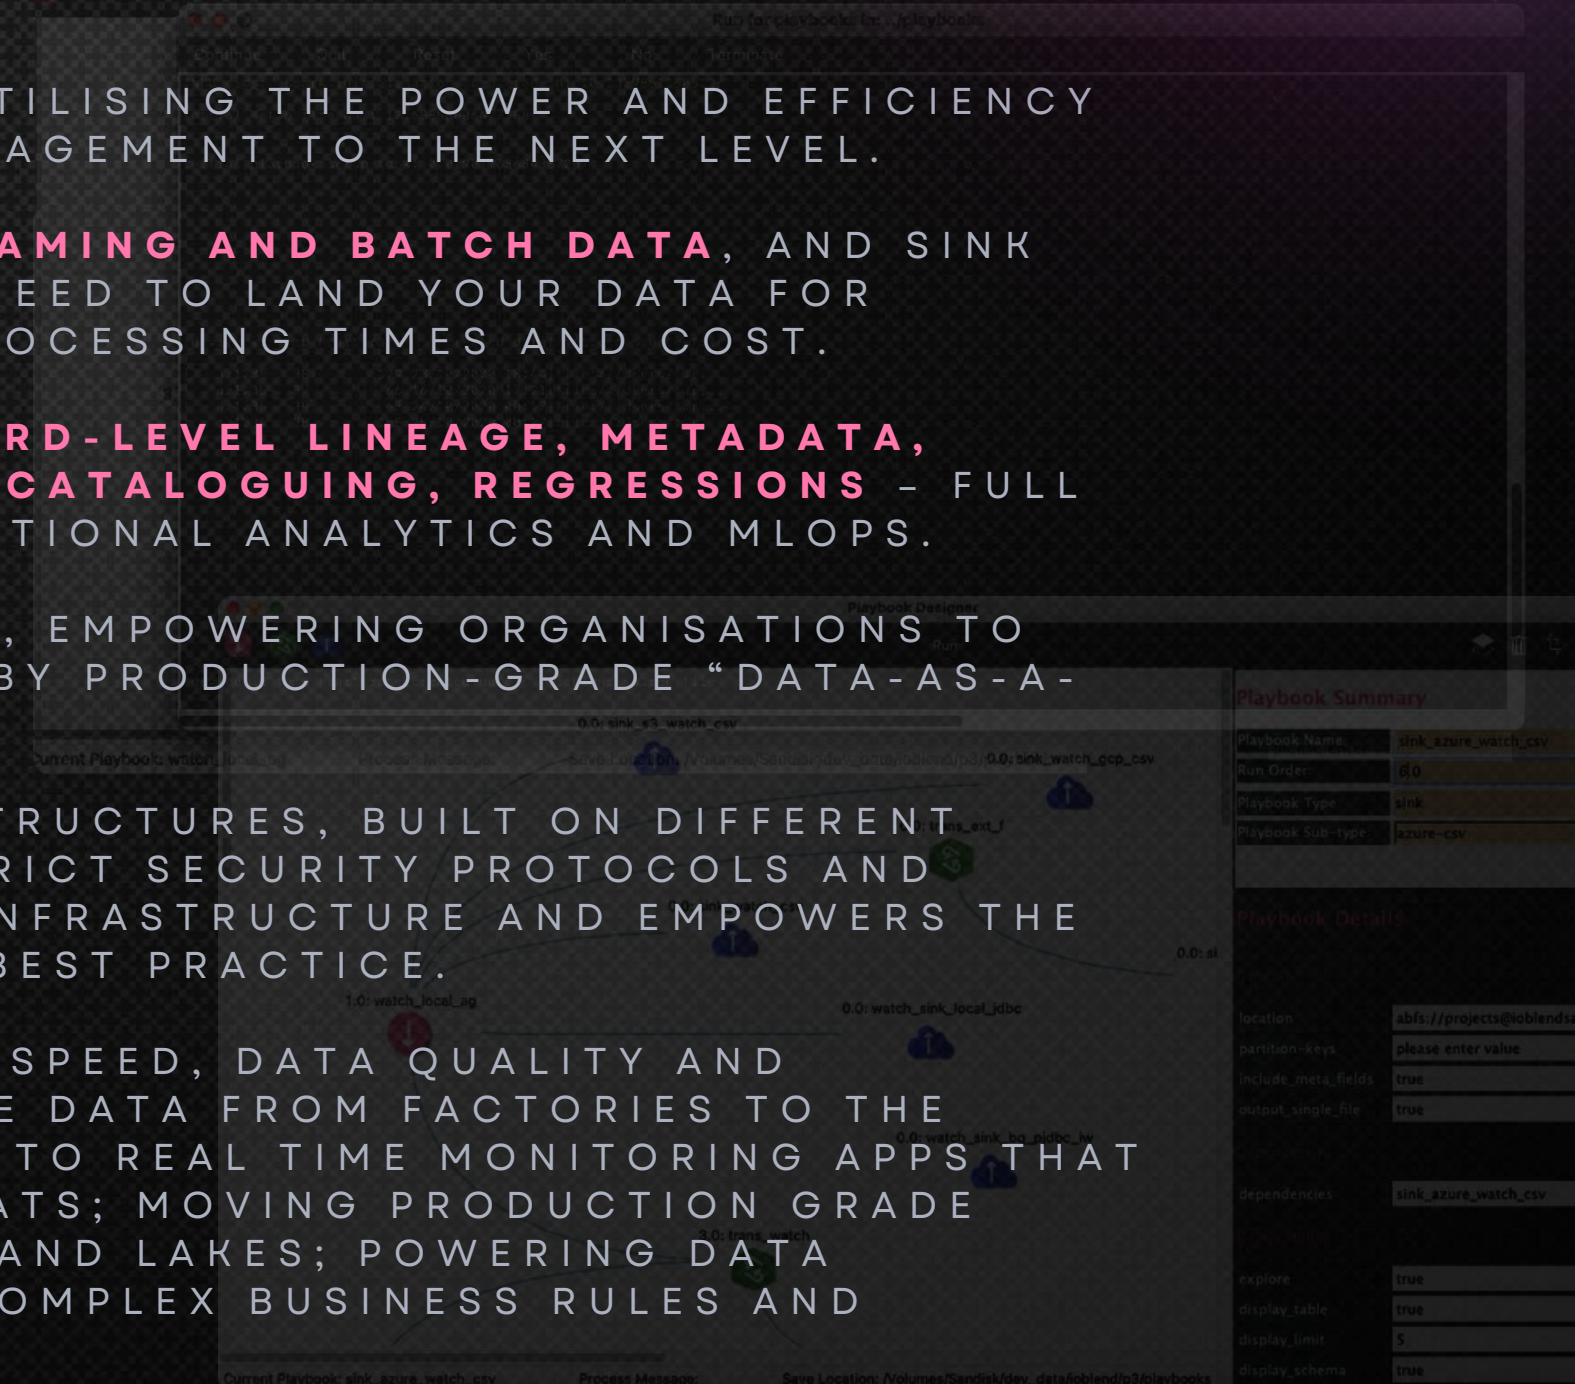
CONNECT TO ANY SOURCE, PERFORM IN-MEMORY TRANSFORMS OF **STREAMING AND BATCH DATA**, AND SINK THE RESULTS TO ANY DESTINATION WITH LITTLE EFFORT. THERE IS NO NEED TO LAND YOUR DATA FOR STAGING – PERFORM YOUR ETL IN FLIGHT, WHICH GREATLY REDUCES PROCESSING TIMES AND COST.

IOBLEND WILL MANAGE YOUR DATA THROUGH ITS FULL JOURNEY – **RECORD-LEVEL LINEAGE, METADATA, SCHEMA, EVENTING, DE-DUPING, SCD, CHAINED AGGREGATIONS, MDM, CATALOGUING, REGRESSIONS** – FULL DATAOPS IN-FLIGHT. THIS ASPECT IS CRITICAL FOR SUCCESSFUL OPERATIONAL ANALYTICS AND MLOPS.

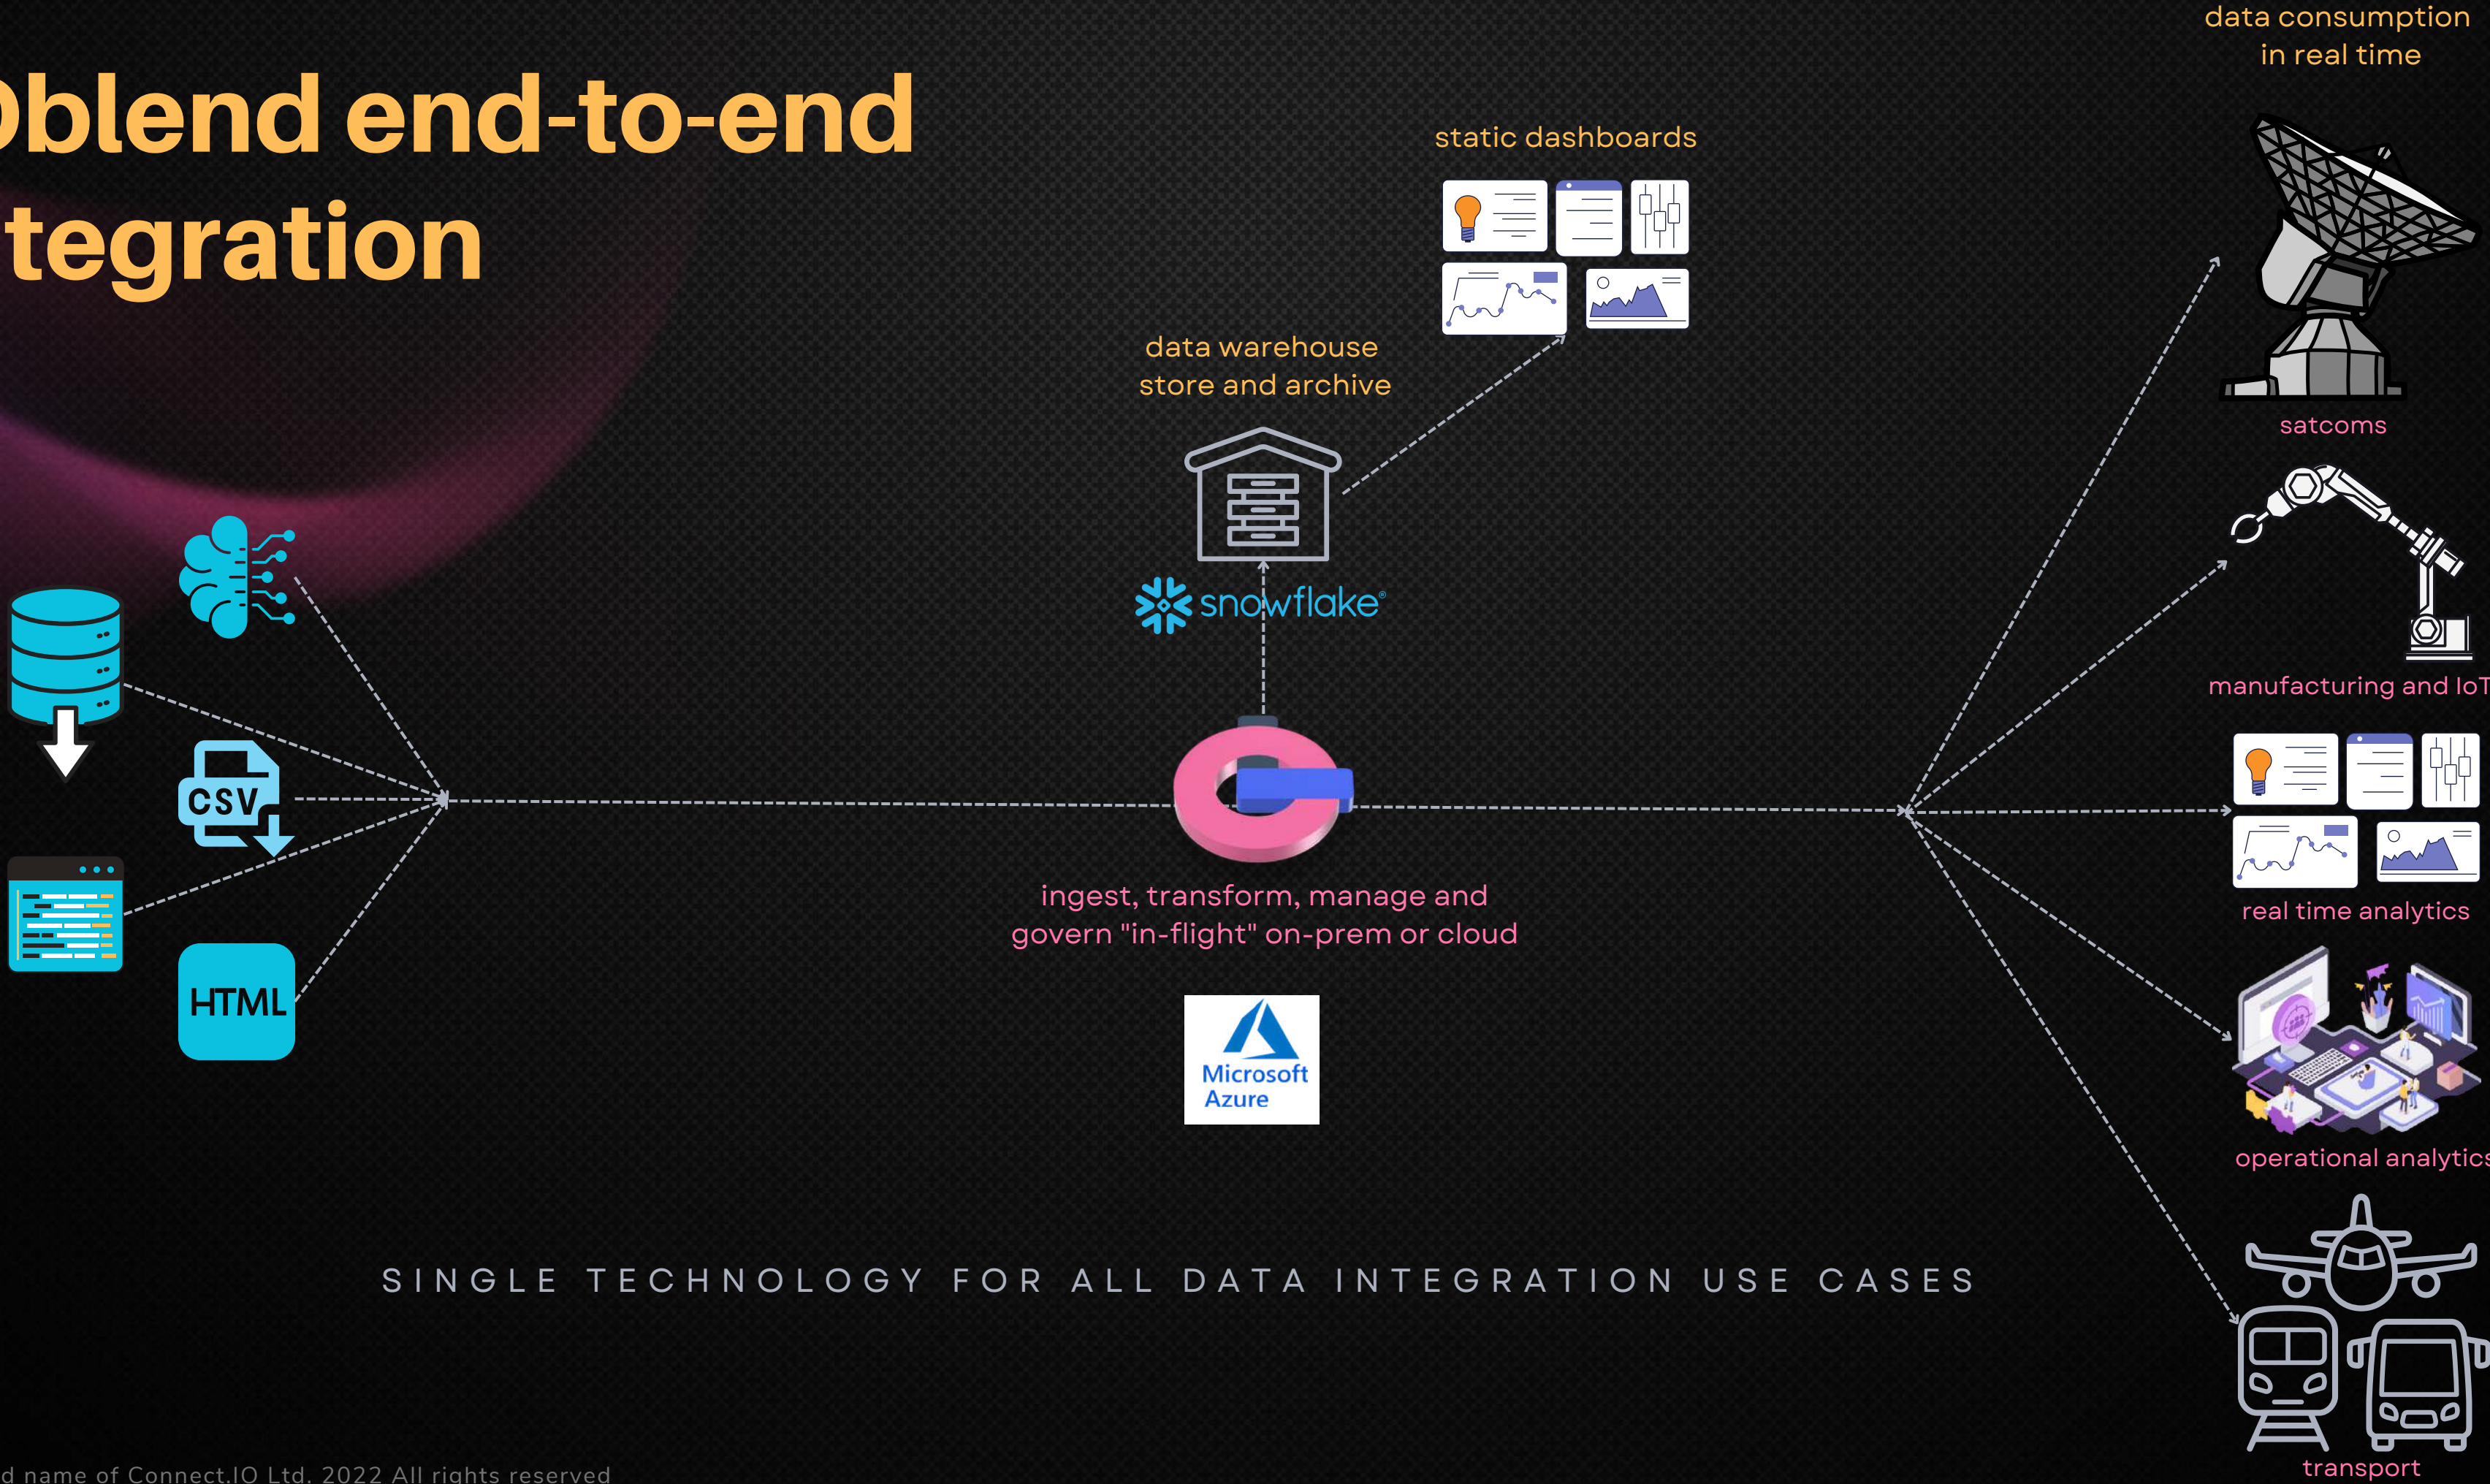
IOBLEND FACILITATES ROBUST **DATA MESH** (FEDERATED) ARCHITECTURE, EMPOWERING ORGANISATIONS TO APPLY DOMAIN-ORIENTED DATA DESIGNS, AUTOMATICALLY SUPPORTED BY PRODUCTION-GRADE “DATA-AS-A-PRODUCT” DATAFLOWS.

SEAMLESSLY CONNECT MULTIPLE DATA DOMAINS ON DIFFERENT INFRASTRUCTURES, BUILT ON DIFFERENT SEMANTICS, UPDATE FREQUENCIES AND FORMATS WHILE OBSERVING STRICT SECURITY PROTOCOLS AND PERMISSIONS. THIS REDUCES THE NEED TO RUN A RIGID CENTRALISED INFRASTRUCTURE AND EMPOWERS THE DOMAINS TO INNOVATE, WHILE PRESERVING COMMON STANDARDS AND BEST PRACTICE.

IOBLEND IS BEST SUITED FOR OPERATIONAL ANALYTICS CASES, WHERE SPEED, DATA QUALITY AND RELIABILITY ARE PARAMOUNT. OUR USE CASES INCLUDE STREAMING LIVE DATA FROM FACTORIES TO THE AUTOMATED FORECASTING MODELS; FLOWING DATA FROM IOT SENSORS TO REAL TIME MONITORING APPS THAT MAKE AUTOMATED DECISIONS BASED ON LIVE INPUTS AND HISTORIC STATS; MOVING PRODUCTION GRADE STREAMING AND BATCH DATA TO AND FROM CLOUD DATA WAREHOUSES AND LAKES; POWERING DATA EXCHANGES; AND FEEDING APPLICATIONS WITH DATA THAT REQUIRES COMPLEX BUSINESS RULES AND GOVERNANCE POLICIES.



IOblend end-to-end integration



SINGLE TECHNOLOGY FOR ALL DATA INTEGRATION USE CASES

Data engineering, management and governance like no other

ADD/REMOVE AS MANY DATAFLOW COMPONENTS AS REQUIRED AND CREATE FULLY PRODUCTIONISED DATAFLOWS IN MINUTES, NOT DAYS

CREATE ADVANCED DATAFLOWS FOR STREAMING OR BATCH DATA FROM ANY SOURCE. SINK TO ANY ENVIRONMENT PREM/CLOUD/HYBRID, SYSTEM AND RESPOSITORY - AND FLOW THE DATA BACK TO SOURCE JUST AS EASILY

WE HAVE TEMPLATED AND ANNOTATED ALL DATAFLOW OPTIONS AND MADE IT HASSLE FREE TO CREATE PRODUCTION GRADE PIPELINES AND DATA ESTATES, NO MATTER HOW COMPLEX

IN DEV MODE, YOU CAN INSPECT EACH STEP OF YOUR DATAFLOW DEVELOPMENT AS YOU PROGRESS, MAKING DATAFLOW ENGINEERING A WALK-IN-THE-PARK

PAUSE, AMEND AND UPDATE DATA SOURCES, TRANSFORMS AND SINKS IN LIVE RUNNING DATAFLOWS WITHOUT THE NEED TO STOP THE JOB

METADATA IS AUTOMATICALLY MANAGED, RECORD-LEVEL DATA LOGGING AND CDC ARE APPLIED, AND THE DATA IS DEDUPED - ALL IN MEMORY

EASILY SEE AND EXPORT THE SCHEMAS (CURRENT AND HISTORY)

TEMPLATE AND SAVE THE DATAFLOWS FOR LATER REUSE AND GET THE ENTIRE ANALYTICAL COMMUNITY WORKING WITH THE SHARED KNOWLEDGE TO SAVE TIME AND EFFORT

The screenshot shows the 'Playbook Designer' interface. The main area displays a dataflow diagram with components like '1.0: watch_local_ag', '0.0: sink_s3_watch_csv', '0.0: sink_watch_csv', '0.0: trans_ext_f', '0.0: watch_sink_local_jdbc', '0.0: watch_sink_bq_jdbc_jw', and '3.0: trans_watch'. A 'Run' button is visible at the top right. On the right side, there is a 'Playbook Summary' panel with fields for 'Playbook Name: sink_azure_watch_csv', 'Run Order: 60', 'Playbook Type: sink', and 'Playbook Sub-type: azure-csv'. Below that is a 'Playbook Details' panel with 'General' and 'Debug Options' sections.

The screenshot shows a terminal window with a table of records and log output. The table has columns: device_id, temp, heart_rate, created_date, and ioblend_source_file. The log output shows the number of records in a table and the completion of a fetch operation.

device_id	temp	heart_rate	created_date	ioblend_source_file
az1235	38	129	23/10/2020 18:44	file:///Users/far...
az1235	37	66	16/10/2020 17:21	file:///Users/far...
az1235	39	63	24/10/2020 00:17	file:///Users/far...
cz1234	38	158	25/10/2020 00:17	file:///Users/far...

```
nos of records in table watch_local_ag_secondary_internal is: 4
passfl:- watch_local_ag_secondary_internal fetch completed in 2.2824497494506836 secs
Relocation/purging of files for watch_local_ag_secondary_internal completed in 6.9161387939453125e-06 secs
pid@1: Schema for watch_local_ag_secondary..
{'device_id,StringType,true'}, {'temp,StringType,true'}, {'heart_rate,StringType,true'}, {'created_date,StringType,true'}
pid@1 output from watch_local_ag_secondary..
[device_id|temp|heart_rate|created_date |ioblend_source_file
```

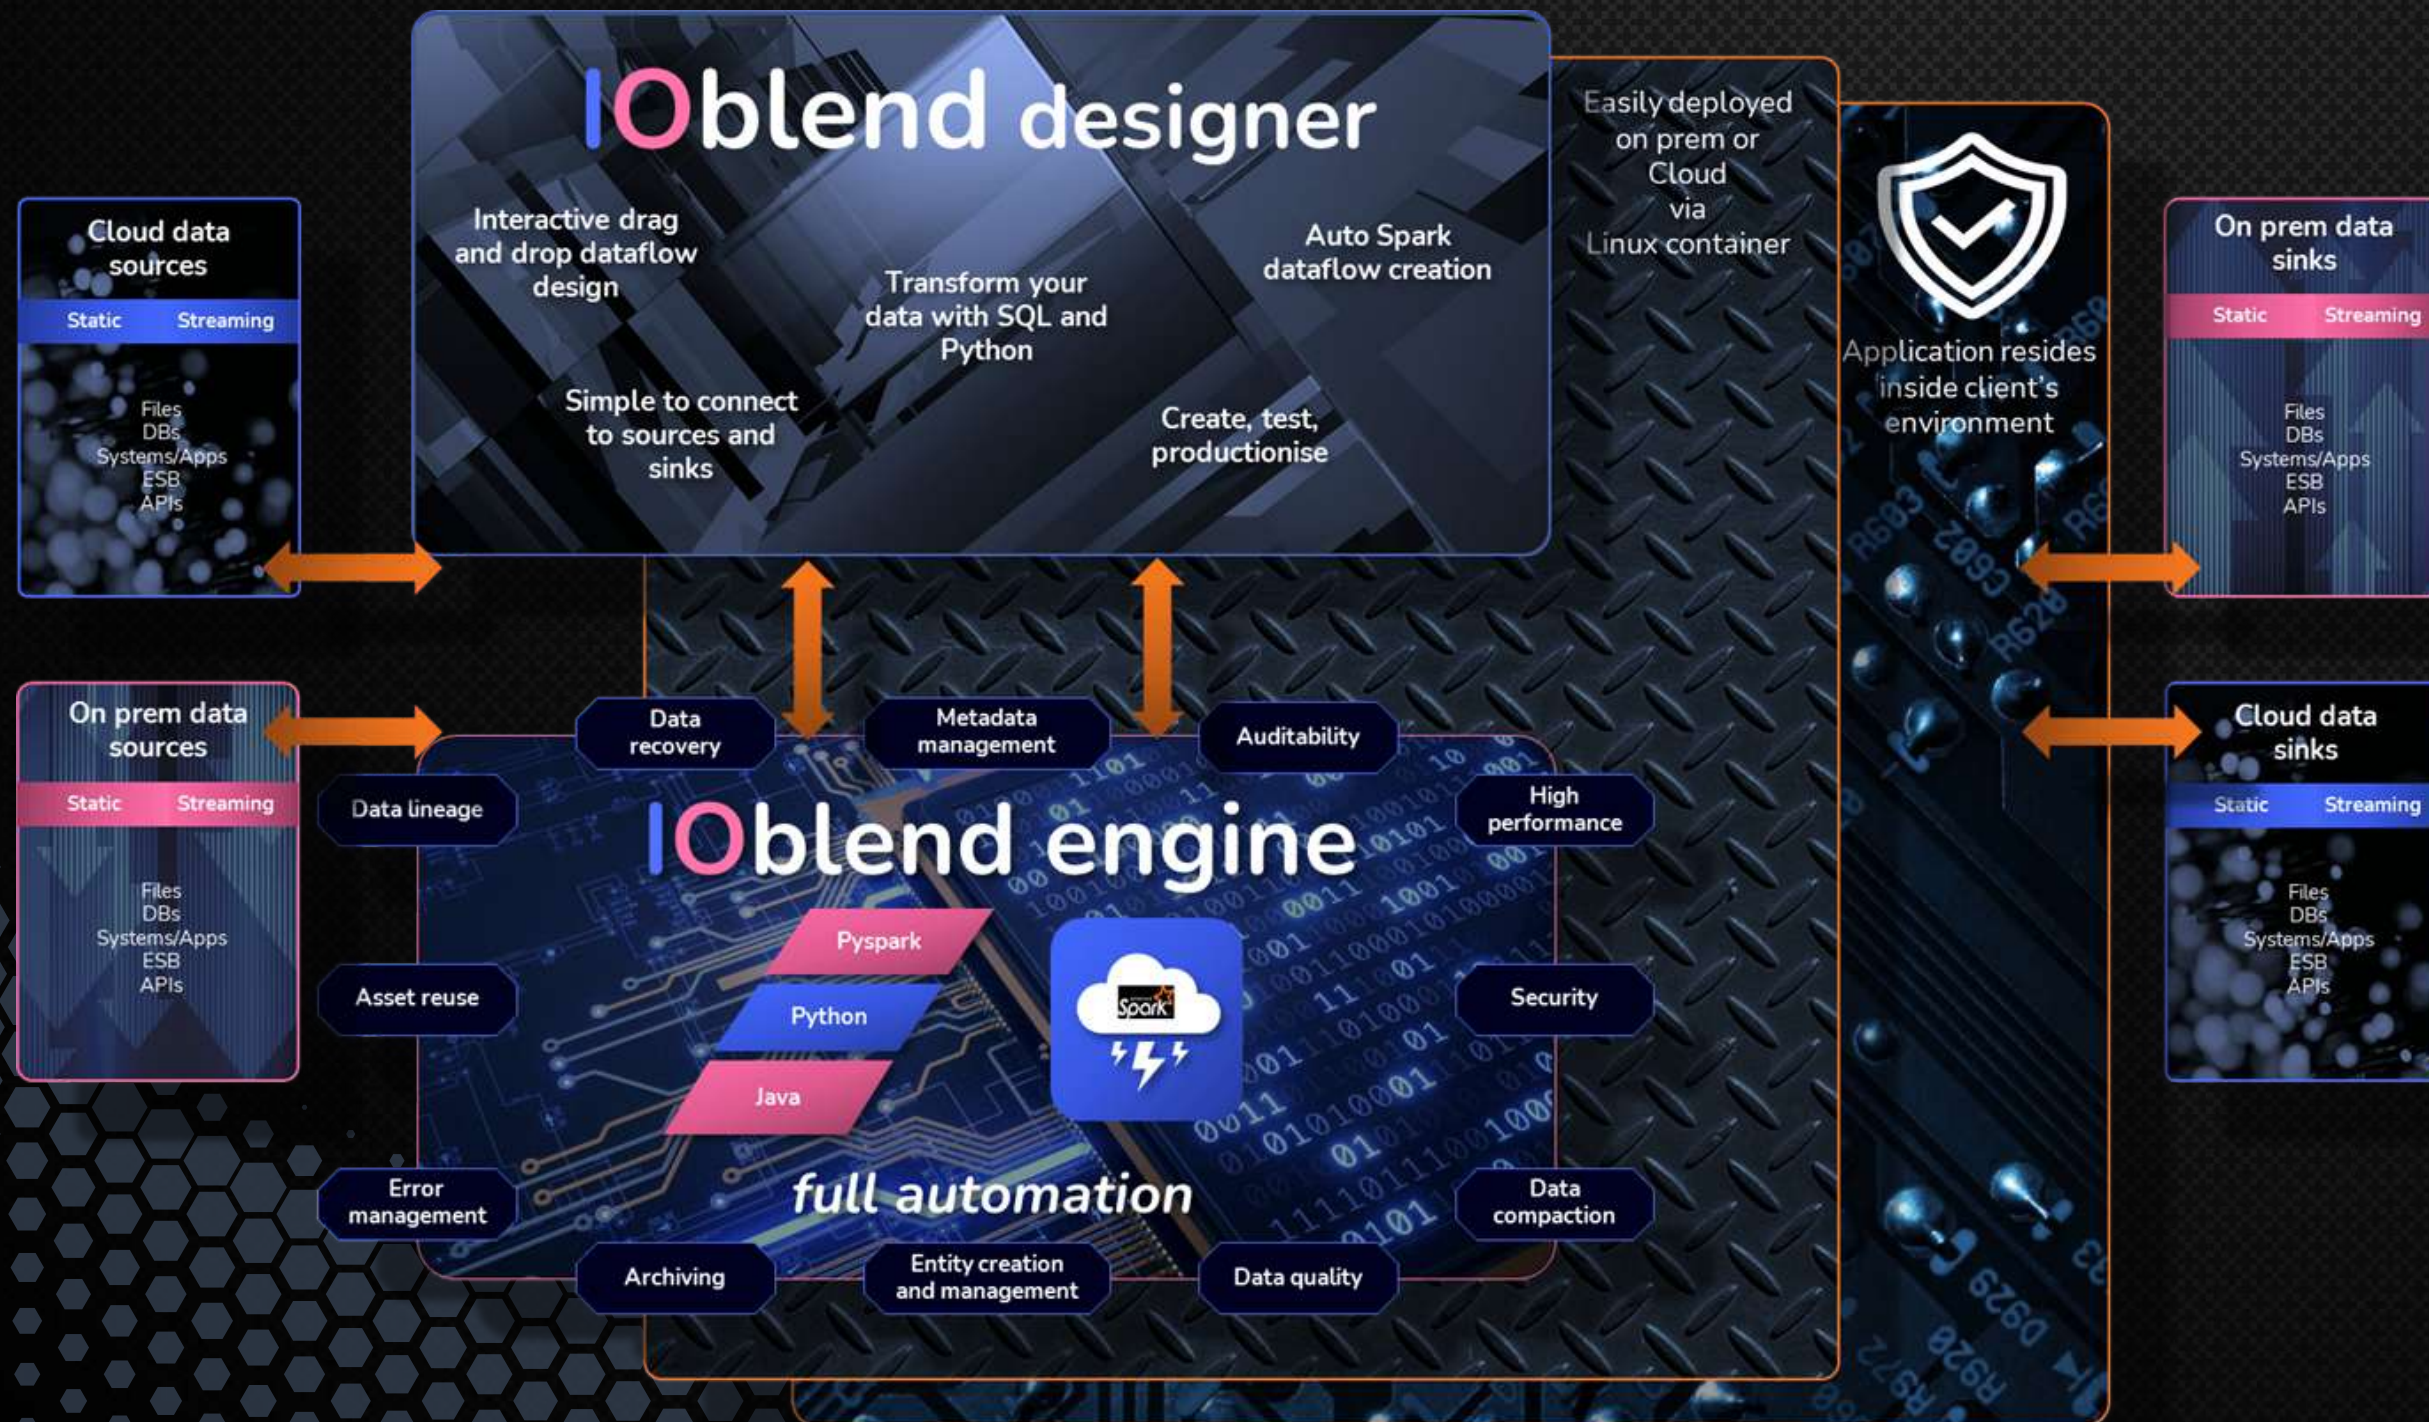
ISV technology partner



TRUSTED BY THE DATA LEADERS

IOblend solution map

THE SOLUTION CONSISTS OF TWO CORE MODULES - **DESIGNER** AND **ENGINE**



IOBLEND DESIGNER IS A VISUAL TOOL FOR DEVELOPING AND MANAGING DATAFLOWS

IOBLEND ENGINE IS A PROPRIETARY FRAMEWORK THAT SITS ATOP OF APACHE SPARK™

DATAFLOWS ARE DEFINED IN PLAYBOOKS, WHICH STORE CONFIGURATION AND THE BUSINESS LOGIC

THE ENGINE DYNAMICALLY CONVERTS THE PLAYBOOKS TO APACHE SPARK™ STREAMING JOBS AND EXECUTES THEM EFFICIENTLY WITHOUT THE NEED TO CODE

IOBLEND SEAMLESSLY FLOWS DATA IN BOTH DIRECTIONS, ACROSS ALL ENVIRONMENTS AND IN REAL TIME, GIVING THE FULL FREEDOM TO MANAGE DATA AS REQUIRED

THE PLATFORM IS EASILY DEPLOYED ON-PREM OR CLOUD

IOBLEND SOFTWARE RESIDES ENTIRELY INSIDE THE CLIENT'S ENVIRONMENT, INHERITING SECURITY PROTOCOLS FOR A COMPLETE PEACE OF MIND



IOBLEND ALLOWS YOU TO
CONNECT TO ANY DATA
SOURCE EASILY

Data sources

CONNECT TO ALL MAJOR FILE SYSTEMS, CLOUDS (AWS, GCP AND AZURE), DATABASES AND LAKES (E.G SNOFLAKE, DATABRICKS), API, FILES, ESB - IF IT HAS DATA, WE WILL CONNECT TO IT!

IOBLEND SUPPORTS TRUE REAL TIME WITH ALL DATABASES VIA OUR UNIQUE STREAMING JDBC CAPABILITY

CONNECT TO MULTIPLE SOURCES SIMULTANEOUSLY, WHETHER BATCHED OR STREAMING, ON THE EDGE, CLOUD, MULTI-CLOUD OR ANY DATA STORES

USE OUR PYTHON TEMPLATES TO CREATE CUSTOM CONNECTORS TO NON-NATIVE SOURCES AND THE SYSTEM WILL AUTOMATICALLY SAVE AND STORE THAT CONNECTOR FOR YOUR FUTURE USE



TRANSFORM YOUR DATA TO THE EXACT BUSINESS REQUIREMENTS IN MEMORY, WITHOUT THE NEED TO STAGE IT - CONSIDERABLY FASTER AND MORE COST EFFECTIVE THAN TRADITIONAL METHODS

Data transforms

CREATE YOUR TRANSFORMS USING SQL OR SCRIPT MORE COMPLEX TRANSFORMS WITH OUR PYTHON TEMPLATES, WHEN THE PROJECT REQUIRES IT

WE HAVE ADDED ENRICHED TRANSFORM FUNCTIONALITY SUCH AS CHAINED AGGREGATES, ENHANCED EVENTING AND REGRESSION OF LATE ARRIVING DATA

IOBLEND AUTOMATICALLY CONVERTS YOUR SCRIPTS INTO IN-MEMORY DISTRIBUTED COMPUTE SPARK TRANSFORMS, GIVING YOU THE CAPABILITY TO TRANSFORM HUGE AMOUNTS OF DATA FAST AND EFFICIENTLY

BUSINESS LOGIC IS YOUR CORE EXPERTISE, SO WE GAVE YOU THE BEST MEANS TO APPLY IT TO YOUR DATA AS YOU SEE FIT. BUT TOOK AWAY ALL THE COMPLEXITY OF PRODUCTIONISING YOUR DATAFLOWS

Data sinks

IOBLEND SUPPORTS NATIVE CONNECTIONS TO AWS, GCP AND AZURE CLOUDS AND ANY ON-PREM/HYBRID REPOSITORIES

SINK YOUR DATA TO ANY DATABASE, DATA LAKE OR DATA STORE, TEMPORARY OR PERSISTENT, WITH JUST A FEW CLICKS

FEED YOU LIVE SYSTEMS, EDGE DEVICES AND DASHBOARDS WITH ROBUST DATA - IOBLEND SUPPORTS THEM ALL

FLOW YOUR TRANSFORMED AND ENRICHED DATA (BATCHED AND STREAMING) BACK TO YOUR SOURCE SYSTEMS JUST AS EASILY - CREATE A COMPLETE ANALYTICAL LOOP AS NEEDED



DROP YOUR DATA
WHEREVER YOU REQUIRE
EASILY AND SECURELY

Production grade data management

automatic

Automatically build and maintain your data lakes and warehouses on any environment

automatic

Automatically creates a searchable metadata catalogue for all schemas and dataflows encouraging reuse

automatic

Automatic data lineage at record level, auditing metadata and determining data changes between sources

automatic

Automatic logging and monitoring of errors and associated directed alerts

automatic

Automatic data compaction for enhanced performance and cost reduction

automatic

Automatic data archiving

automatic

Automatic late arriving data management and unlimited event windows - you define the rules and the system takes care of the rest


totally secure

Run IOblend entirely inside your data security protocols for a complete peace of mind

IOBLEND FULLY AUTOMATES AND STANDARDISES ALL COMPLEX DATA MANAGEMENT TASKS, MASSIVELY SAVING YOU TIME AND EFFORT ON DATAFLOW DEVELOPMENT AND DATA ESTATE UPKEEP- NO MORE TECH DEBT

Deployment architecture

IOBLEND CAN BE
DEPLOYED FLEXIBLY ON
ALL ON-PREM AND CLOUD
ENVIRONMENTS





**What makes
IOblend a
winner?**

> 10 mil tps

Handles huge data volumes at low latency across any environment - **we have tested at 10mil tps at <1 sec on an average laptop**

Over 70%

Saving in Data Engineering and Data Management costs vs legacy ETL/ELT

10X

Reduction in development time of data projects vs traditional methods - **we took projects from 9mth to 3 weeks**

Full DataOps

Centralises and standardises all data management and technical governance in one platform

90%

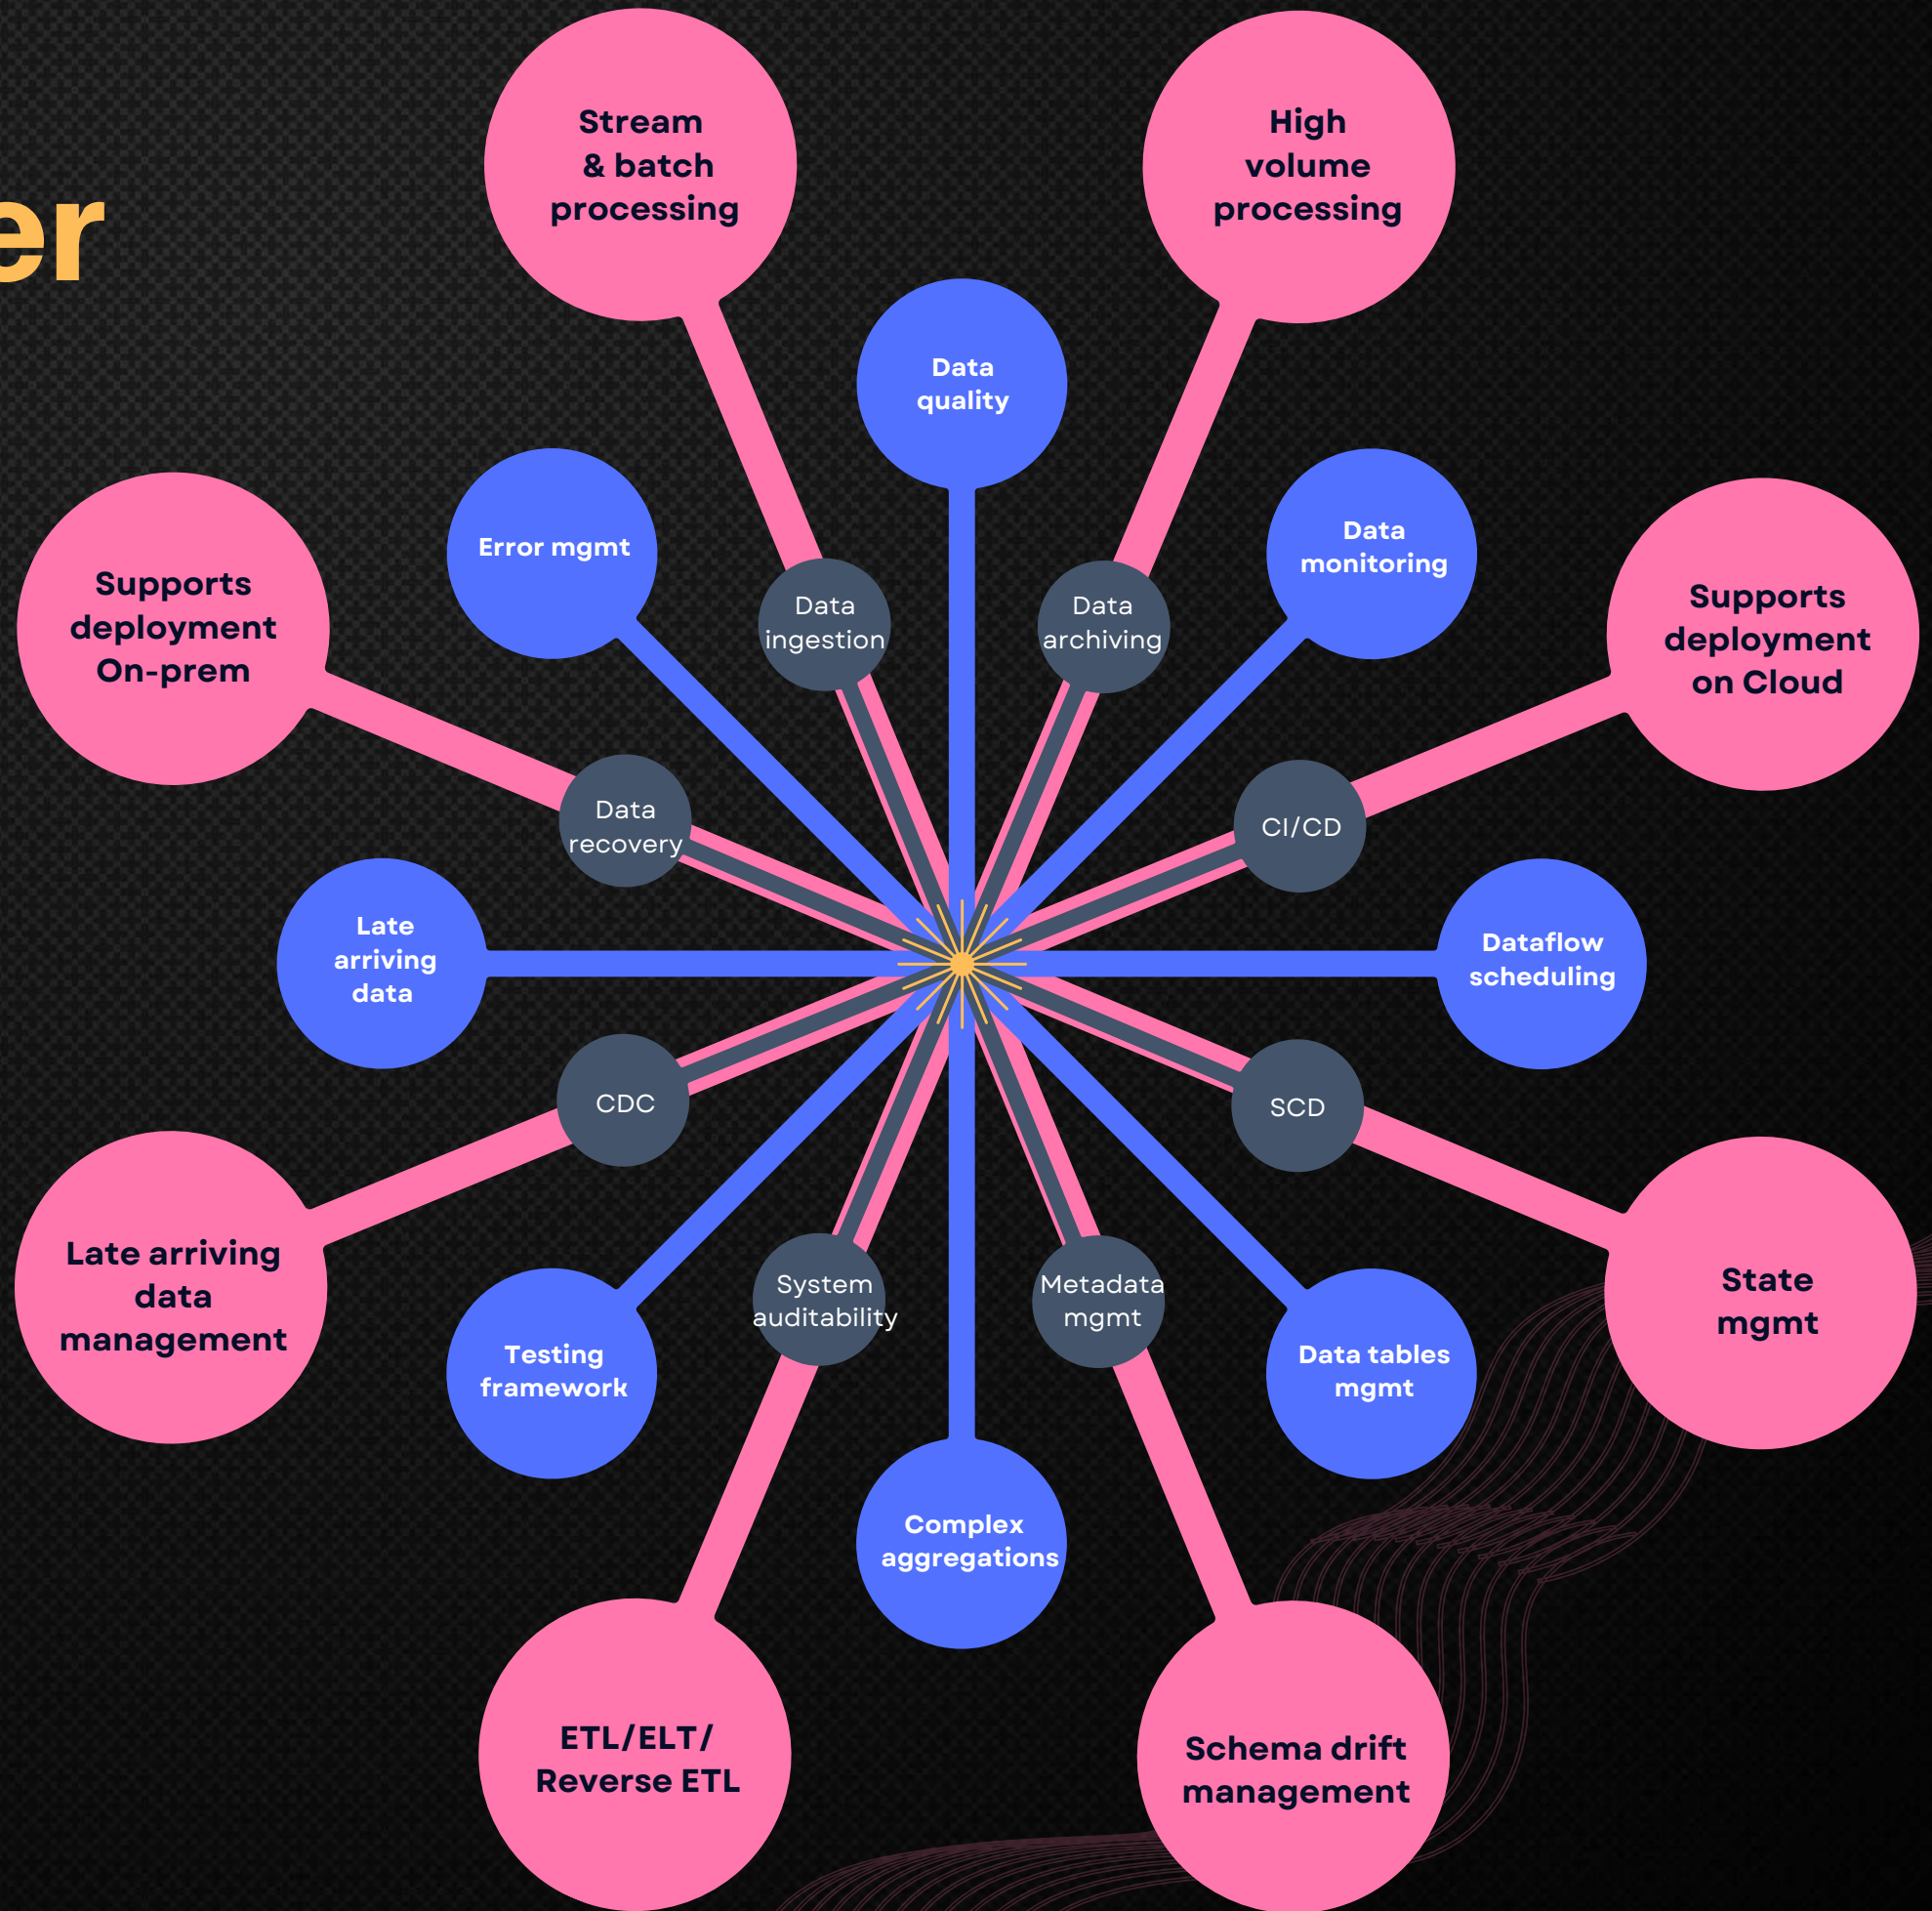
Reduction in data management effort due to the high-level of automation

Data mesh

Easily enables federated data architectures to break your data silos once and for all

Production grade data integration like no other

IOBLEND GIVES YOU THE FULL CONFIDENCE IN YOUR DATA AND HOW IT FLOWS THROUGH YOUR ESTATE BY GIVING YOU THE FULL VISIBILITY OF WHAT IS HAPPENING IN REAL TIME



Covers all data applications

ONLINE STREAMING SERVICES, SATCOMS, AND MOBILE COMMS

PLANT MANAGEMENT, PLANNING AND MONITORING

SMART CITY/HIGHWAYS MANAGEMENT AND IMPLEMENTATION, IOT

REVENUE MANAGEMENT SYSTEMS, NETWORK PLANNING, REAL-TIME DATA ANALYTICS, HEALTHCARE AND CRM

SPACE/SEA/AIR/RAIL MANAGEMENT SYSTEMS AND LIVE FORECASTING

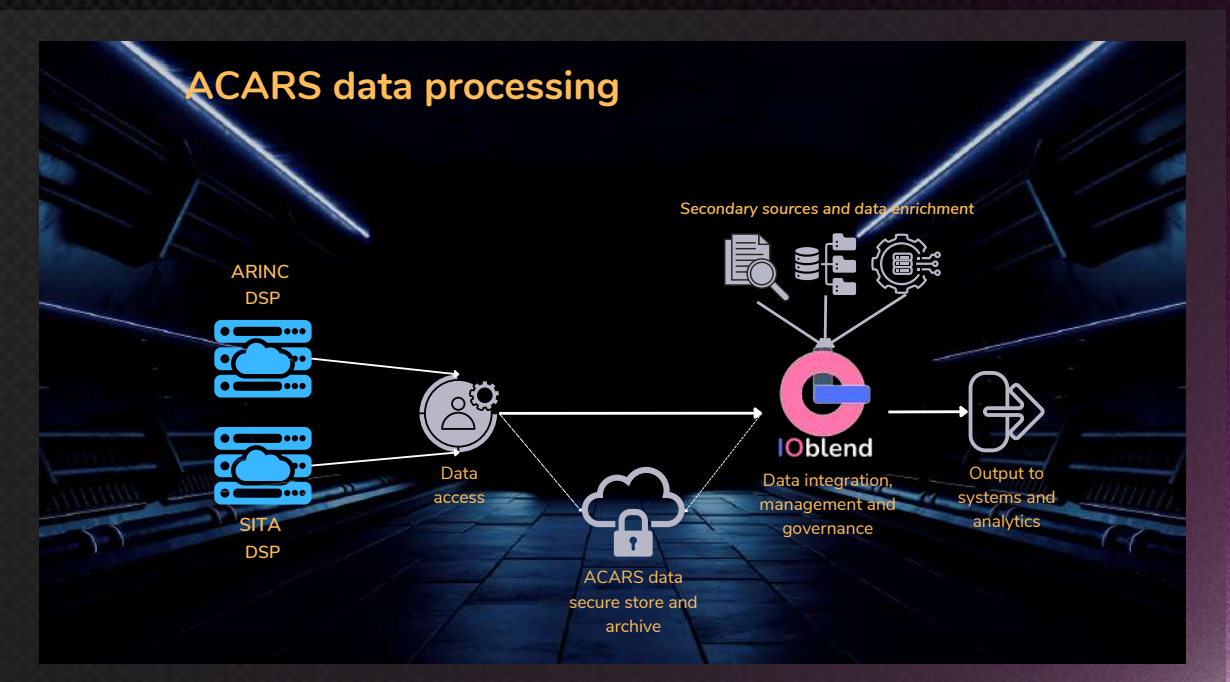
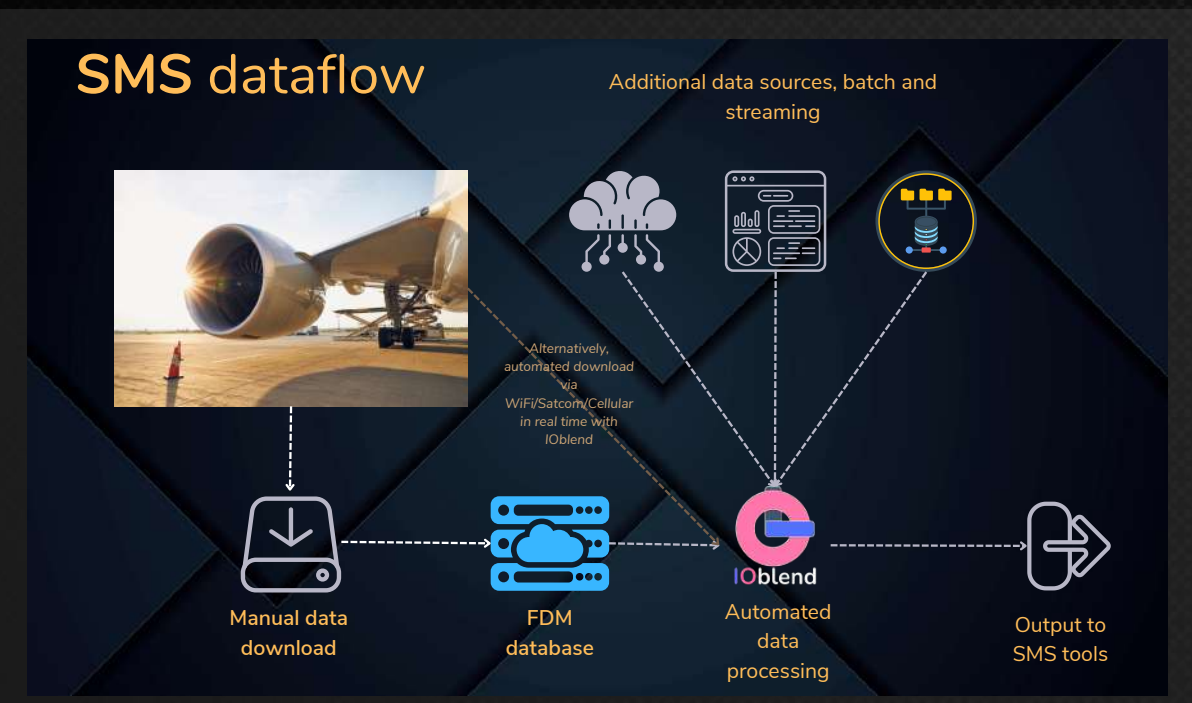
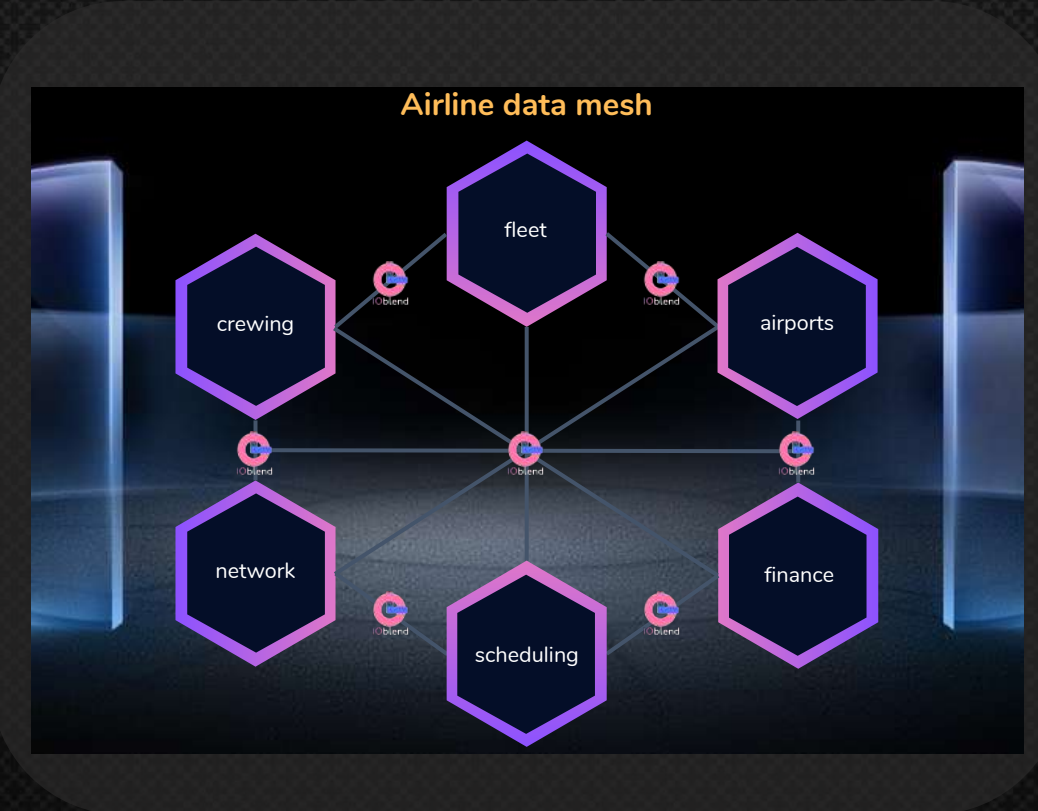
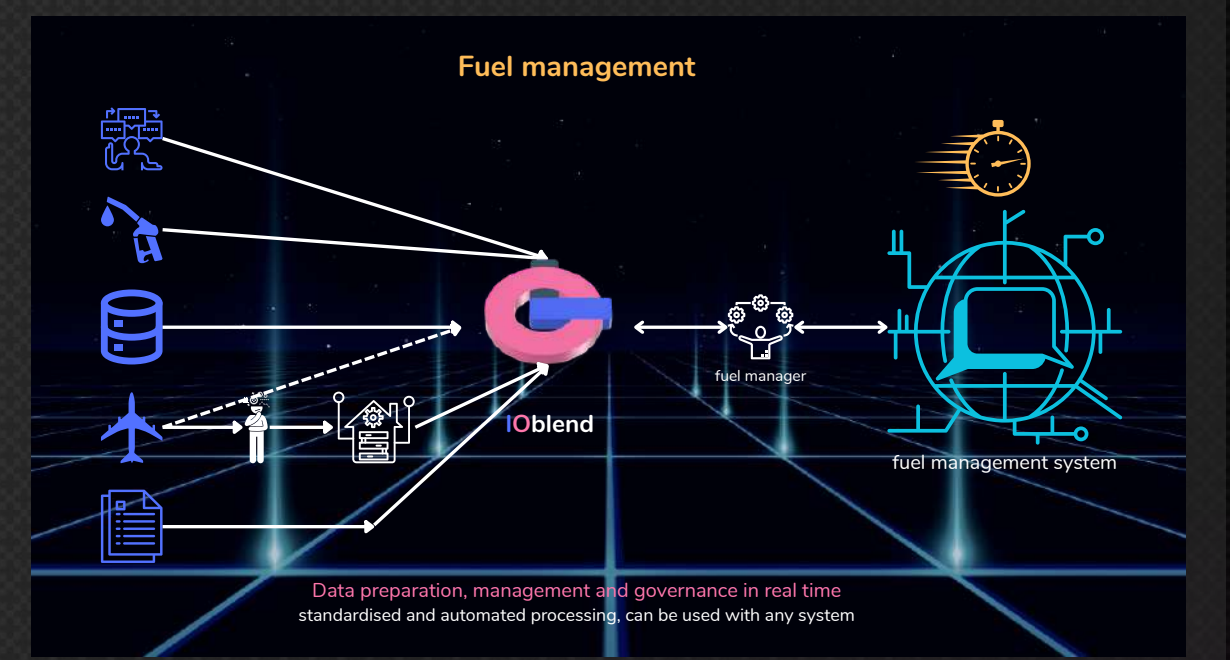

RESOURCE PLANNING AND MANAGEMENT - LIVE OPTIMISATION AND ALLOCATION

CENTRAL OR FEDERATED DATA REPOSITORY CREATION AND MANAGEMENT

DATA MIGRATION, CLOUD IMPLEMENTATIONS, INTEGRATION AND LEGACY SYSTEM UPGRADE PROJECTS

FINANCIAL INSTRUMENTS AND LIVE TRADING DATA MANAGEMENT

Practical application examples



The future of data engineering has arrived

IF YOU ARE IMPLEMENTING ENTERPRISE
DATA PROJECTS, YOU OWE IT TO
YOURSELVES TO DISCOVER WHAT THE ART
OF THE POSSIBLE IS WITH IOBLEND



IOblend

DISCOVER THE POSSIBLE

@IOBLEND.COM

MAKE YOUR DATA ESTATE STATE-OF-THE-ART