



# Single Data View<sup>©</sup> Platform

**Jonathan Summerfield – CEO & Founder**

[jonathan@xiatech.co.uk](mailto:jonathan@xiatech.co.uk)

[www.xiatech.co.uk](http://www.xiatech.co.uk)

[@summerfieldj](https://twitter.com/summerfieldj)



We make technology. Simple.

# Contents



Single Data View® Platform

- Experience and Expertise
- Overview of MicroService / Integration Framework
- Anatomy of a Xiatech MicroService
- Overview of the SDV Data Layer & GDPR admin
- Comparison to a standard Data Warehouse & cubes
- Case Study
- Q & A
- Additional benefits e.g. GDPR & data compliance



# About Xiotech



Single Data View<sup>®</sup> Platform

- Founded 2013 as a consultancy
- Global reach
- Specialists in **B2C & B2B Retail, Supply Chain, Betting & Gaming, Travel, Financial Services, Health Care, Leisure**

## SERVICES:

1. Digital Strategy & Transformation
2. Programme Delivery
3. Software Engineering
4. Big Data/Analytics



# Who we've worked for



# Xiatech Integration Expertise

- ❑ We're experts in APIs, MicroServices, SOA, EAI, and Event Driven Architecture as well as Data, Analytics and Business Intelligence solutions. We've worked on some very large scale integration programmes for household names. We have collectively delivered complex solutions within highly complex programmes and enterprises
- ❑ We've led the integration & information strategy, architecture and platform implementation for companies such as Tesco, Ladbrokes and EasyJet
- ❑ Xiatech Integration Framework and Single Data View Platform is a distillation of what we have learned in 15+ years of integrating systems
- ❑ Our platform has also been adapted to provide companies with a systemised way of tracking, maintaining, reporting and sustaining data compliance

# Our software

1. Real-time Integration platform
2. Real-time 360<sup>0</sup> view of your data
3. Industry-ready Reporting & Dashboards
4. Automated GDPR compliance

**License-free**, secure, cloud based, open source, resilient & scalable



## Single Data View© platform

Single Data View® Platform

# SDV Integration Platform – overview

- The SDV platform is a fully managed service and requires zero maintenance from the client
- The integration platform provides the SDV Operational data store out of the box at no additional cost, providing a 360<sup>0</sup> view of corporate data and API-based access
- In addition, the data is encrypted at rest, providing GDPR assurance on any customer data stored with the SDV. A GDPR automation utility can be added at an extra fee depending on interfacing systems
- A full SLA document will be provided detailing the support and maintenance cover
- The SDV integration platform allows for the addition of the SDV data lake and reporting suite but can be connected to any Data Warehouse platform

# Information challenges today...



Single Data View® Platform

- Timeliness of data
- Data silos
- Different formats
- Duplicated data
- Complex System
- Differing technologies
- Company acquisitions/mergers
- Speed of consumption (velocity) is restricted
- EU Regulatory directives eg GDPR



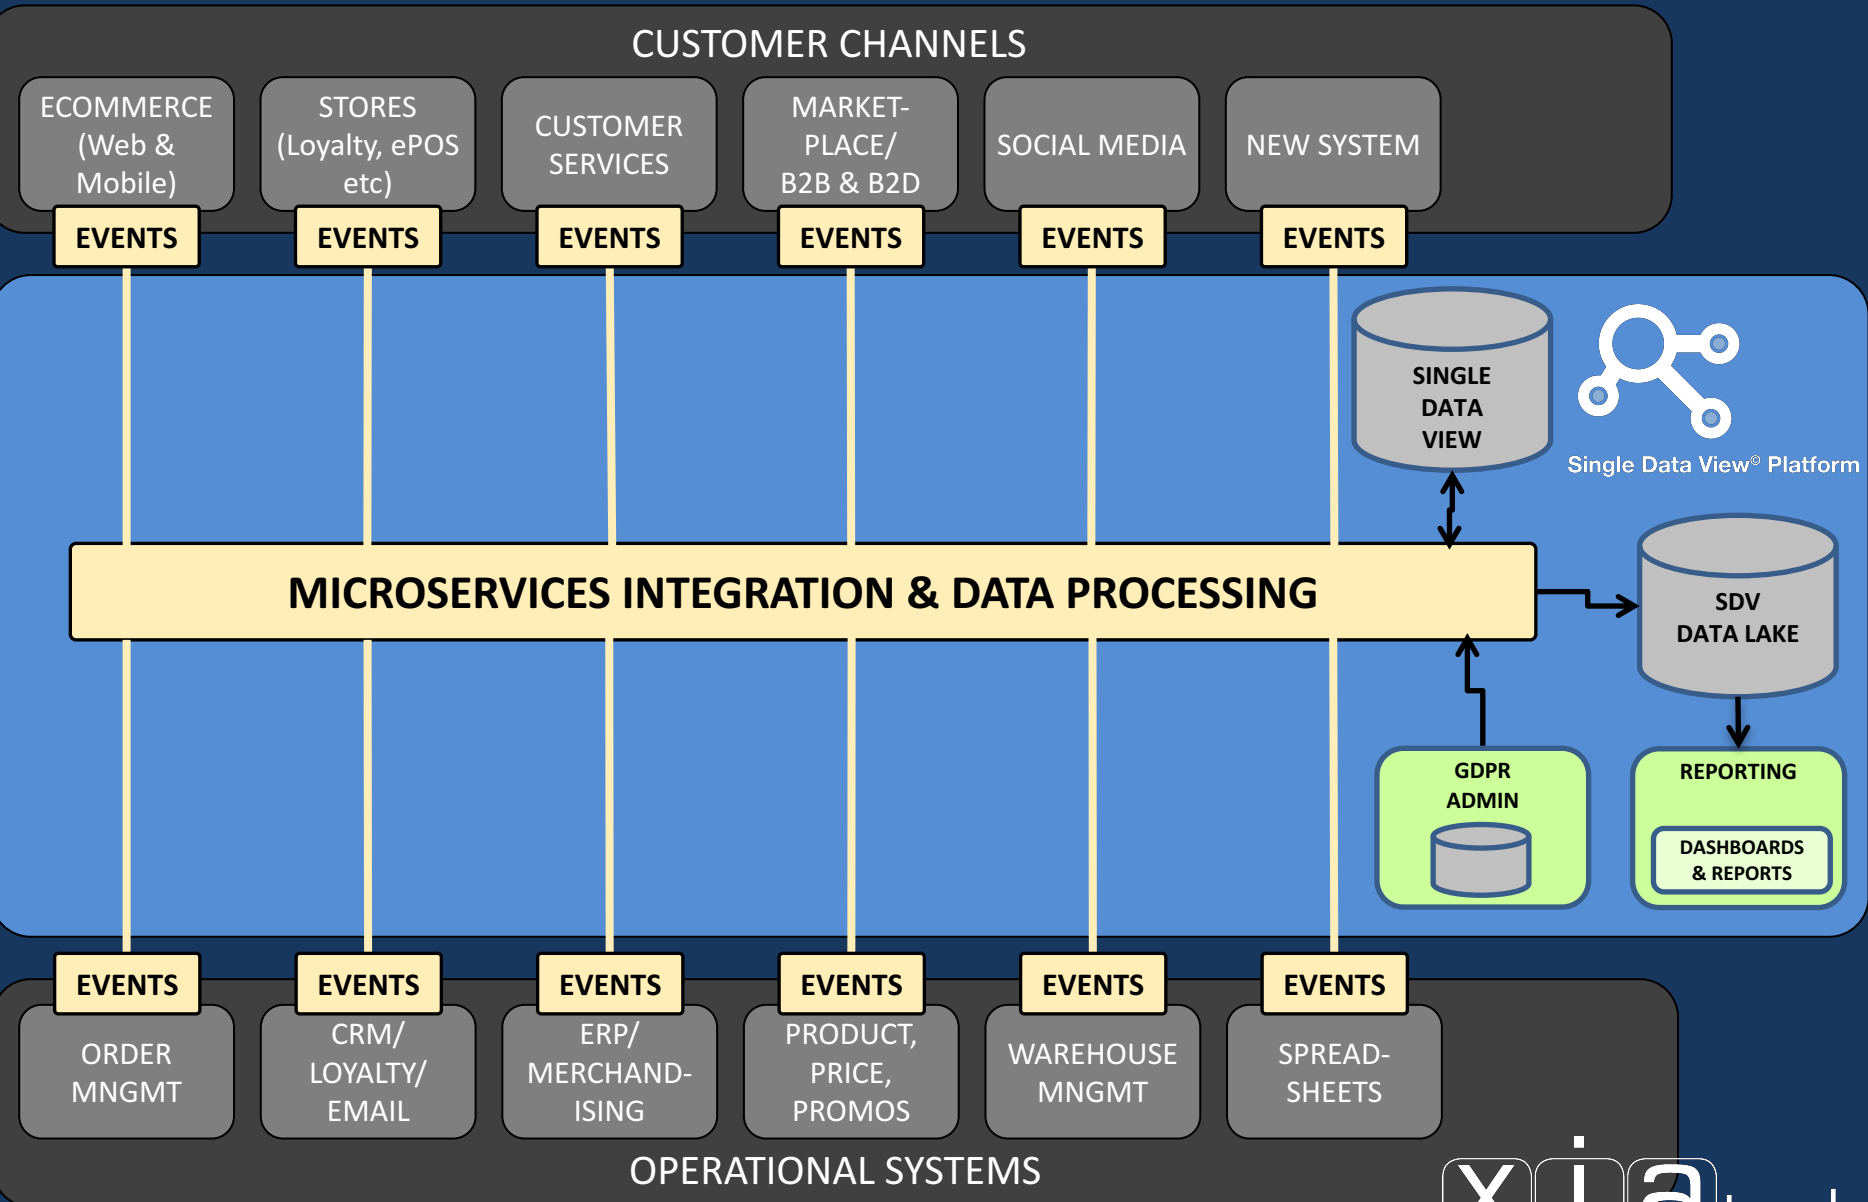



# Why do traditional methods fail...

- Technology & business process constrained
- Not-real time
- Non-GDPR compliant
- Expensive & complex to change
- Not contextually aware (not event based)
- No opportunity to react to data changes
- Difficult to capture unstructured data
- The Data Warehouse and CRM systems maintain the silos



# Our Single Data View © platform...





# Why is it cool for all companies?

Single Data View® Platform

- No license fee (open source) and very low TCO
- Can be implemented in under 3 months
- 360° view of all your data across all your channels – in real time
- Can integrate all your systems - even the old ones!
- Provides industry-ready reports and dashboards out of the box
- Uses latest technology (e.g. MicroServices) so we can attract the best talent
- Continuous improvement – our roadmap is driven by the industry

# WHY COMPANIES CHOOSE OUR SINGLE DATA VIEW<sup>®</sup> PLATFORM



**Secure &  
GDPR  
compliant**



**360° data &  
Industry ready  
real-time  
business decisions**



**Quick to  
implement with  
pre-built  
connectors**



**Cloud based &  
license free**




**Highly scalable &  
resilient**



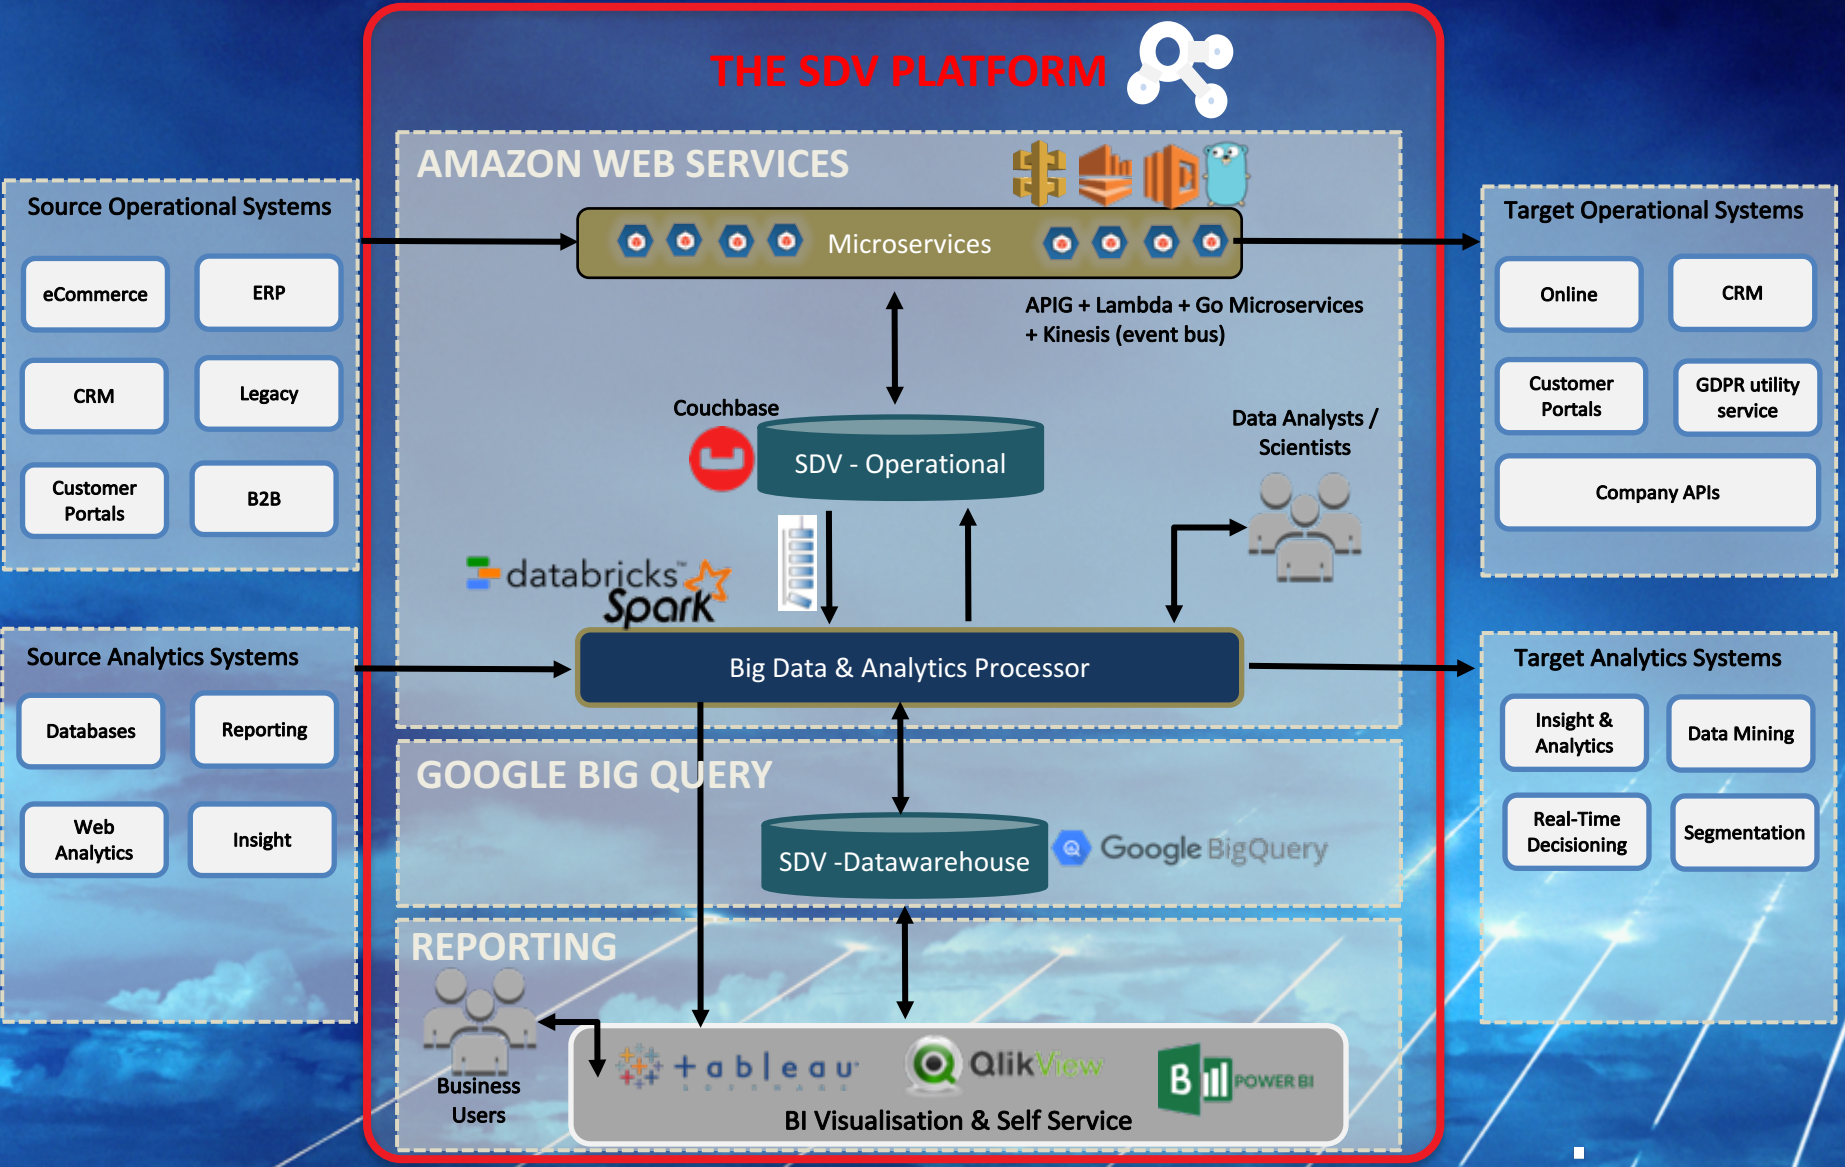
**Single Data View<sup>®</sup> Platform**

**xia**tech:  
consulting

We continue to build integrations into various platforms...



# Our SDV architecture with Reporting Infrastructure



# How does the software work?

*Simply put, we gather your operational data from all your data sources in real time, consolidate it and hold it centrally. From there, the opportunities are endless...*




**GATHER                      INGEST                      ORGANISE                      PREPARE                      EXPLORE**



Single Data View® Platform





Single Data View® Platform

# SDV Integration framework



# Xiotech MicroService / Integration Framework #1

- ❑ **Pre-built Technology Adapters** for connecting to COTS, Databases, SFTP, Web Services/APIs, Message Buses etc..
- ❑ **Built-in support** for Caching, Compression, File Processing, Cryptography, Job Scheduling, Authentication, Logging, Monitoring & Alerting, Service Discovery
- ❑ **Secure** – data encrypted at rest and in-transit. OAuth2 supported. Secrets (passwords etc.) stored in a software Vault
- ❑ **Elastic Scalability** – automatically add/remove new instances of services based on current load. New instances ready in seconds
- ❑ **Cloud Native** – operates in the cloud and can communicate with systems on the ground or in other cloud environments



# Xiotech MicroService / Integration Framework #2

- ❑ **Heterogeneous** – any system or service can interact with our services or subscribe to events, regardless of technology
- ❑ **DevOps** – automated environment creation, continuous integration and deployment
- ❑ **Low Barrier to Entry** – easy for a developer to use. The only pre-requisite is some programming experience
- ❑ **Isolation** – each service is independently managed, deployed, scaled. Allows for micro-deployments
- ❑ **Event Enabled** – services publish events for other systems (including SDV) and services to consume or react to



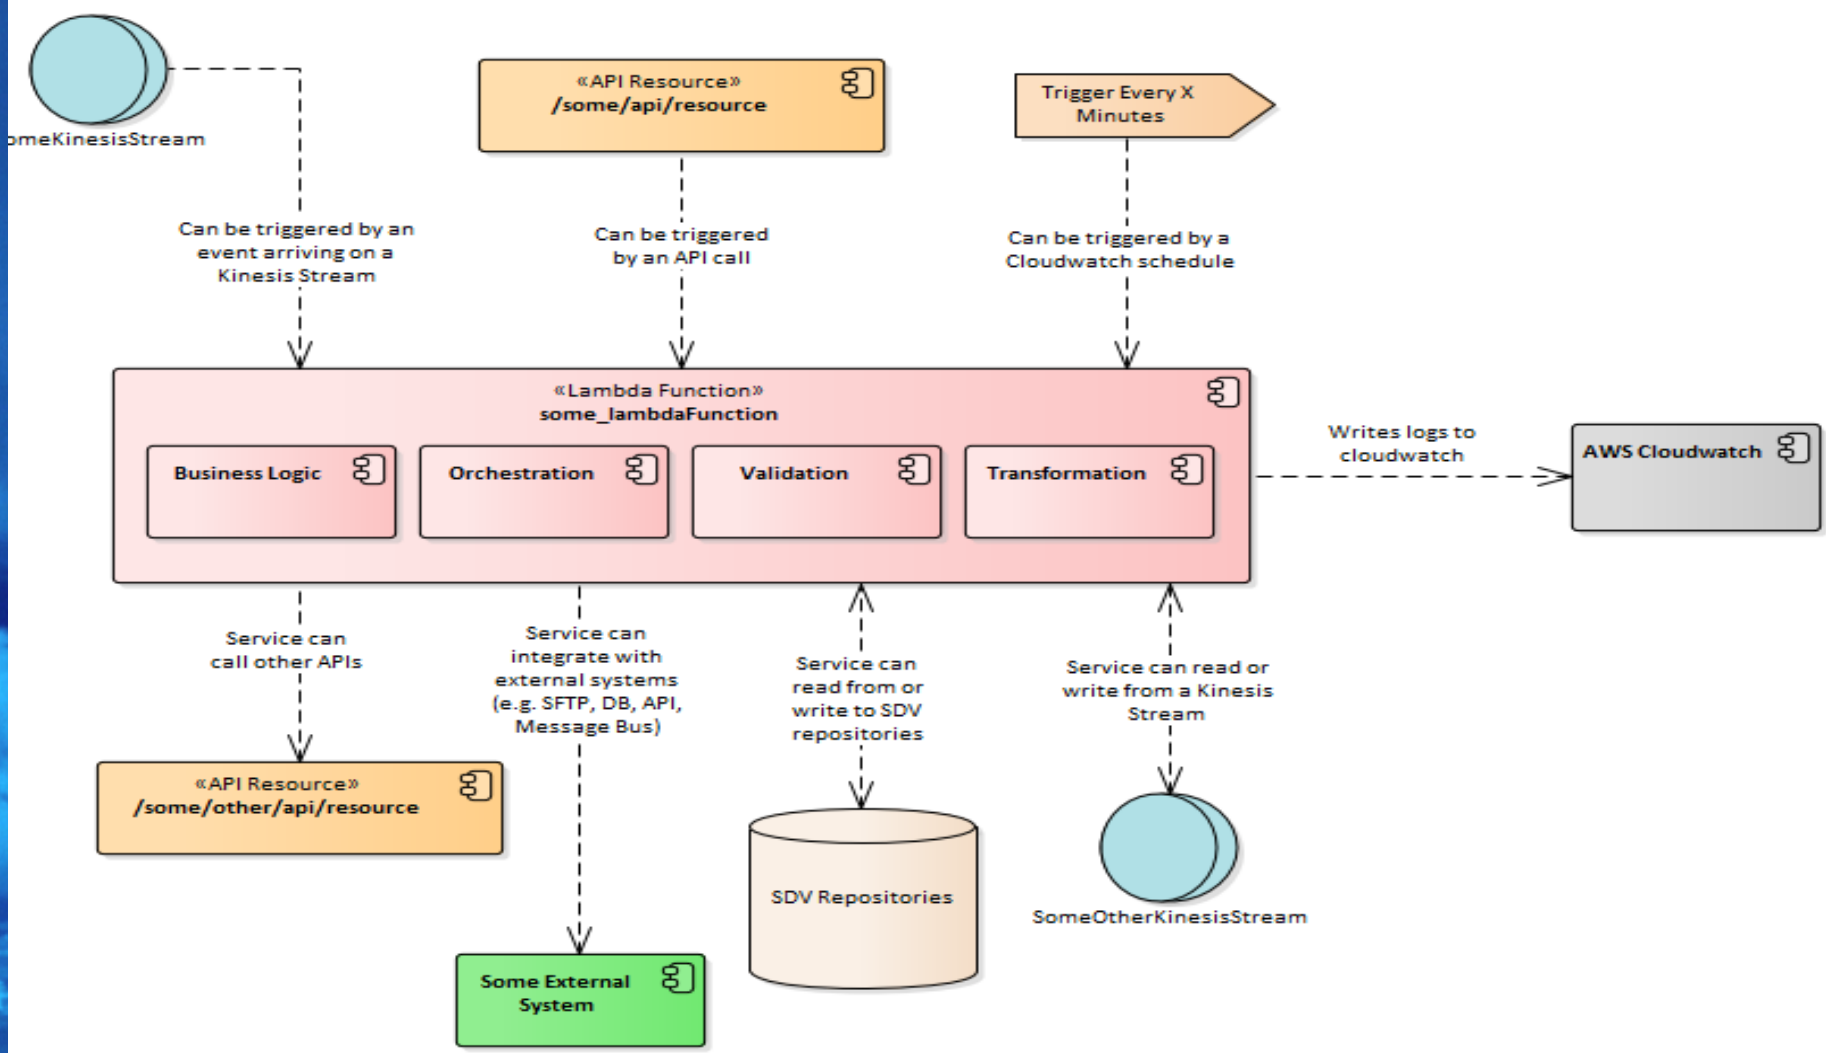
# Xiotech MicroService / Integration Framework #3

- ❑ **Connectors** - reusable integration for business systems such as campaign, e-commerce, and analytics platforms
- ❑ **Easy Maintenance** – Using AWS Serverless stack means we have no servers to maintain and patch, no clusters or load balancers to manage
- ❑ **Fast to Deliver** – A simple service (e.g. receive a message, transform, DB lookup, send to system X) will typically take 1-2 days to develop and test
- ❑ **High Performance** – Services execute fast, and the platform can easily scale to many thousands of executions per second



# Anatomy of a Xiatech MicroService

This is what a Microservice looks like in our platform



# Anatomy of a Xiotech MicroService

- ❑ A MicroService can have 1 or more trigger methods:
  - ❑ An HTTPS API Call
  - ❑ An Event or message on a queue
  - ❑ A Schedule
  - ❑ A DB change event in a database
- ❑ The code executes as an AWS Lambda function. This means that in response to any of the triggers above, AWS will invoke our service. We don't need servers or even Docker containers.
- ❑ All logs, including detailed trace information and execution times are captured by Cloudwatch where we setup monitoring and alert rules
- ❑ Events get published to Kinesis streams, where they can be used to trigger other services or get consumed by the SDV (or other applications)

# Xiotech Single Data View - Integration

How is it different to traditional ESBs such as Mulesoft:

Capability	ESB (e.g. Mulesoft)	SDV Microservices
Data Latency	<ul style="list-style-type: none"><li>❑ Real-time</li></ul>	<ul style="list-style-type: none"><li>❑ Real-time</li></ul>
Risk and Complexity	<ul style="list-style-type: none"><li>❑ ESB is an extra layer of software to build, install, support, monitor, maintain, plus regular software releases and patches that need planning with downtime etc.</li><li>❑ Usually complex to implement and maintain</li><li>❑ Requires operational resource to maintain, monitor and support with regular maintenance necessary</li><li>❑ Usually complex infrastructure to manage, and expensive consultants to execute</li></ul>	<ul style="list-style-type: none"><li>❑ Much simpler – almost <b>zero setup required</b></li><li>❑ Use a cloud-native infrastructure directly – <b>one less layer of overhead</b> to deal with (no ESB required)</li><li>❑ <b>No / very little ops</b> – leave heavy lifting to the cloud providers and their SRE teams</li><li>❑ <b>Proven model at scale</b> - Microservices used by the likes of LinkedIn, Amazon, Facebook, Google, Uber – none of whom use an ESB</li><li>❑ <b>Zero downtime</b> required</li><li>❑ Each Microservice is independent from the other - Allows frequent releases of Microservices with zero impact on other Microservices</li></ul>
Data	<ul style="list-style-type: none"><li>❑ Typically transports and transforms the raw data with no context. The reason for the change is lost</li><li>❑ Data is lost forever post-transport</li><li>❑ Playback of messages is complex, with a risk of losing key data such as customer orders</li><li>❑ Lookups and enrichment rely on external systems providing appropriate APIs – e.g. how would you enrich a despatch message with product and image URLs for a customer email?</li></ul>	<ul style="list-style-type: none"><li>❑ Data is transported and transformed but <b>the event is captured so context is delivered</b></li><li>❑ Data is held persistently for resilience, lookup, audit, playback &amp; support – <b>we don't lose data</b></li><li>❑ <b>Single Data View</b> is created out of the box for zero additional cost</li><li>❑ Rich, fast (in-memory cache) data-source and API for <b>lookups / enrichment of all key entities</b></li><li>❑ Events carry data but also <b>enable simple implementation of business requirements</b>. 'When A happens, I want to do X, Y, and Z'</li></ul>

# Xiotech Single Data View - Integration

How is it different to traditional ESBs such as Mulesoft:

Capability	ESB (e.g. Mulesoft)	SDV Microservices
Scale	<ul style="list-style-type: none"><li>❑ Potentially time consuming and expensive capacity planning and provisioning of infrastructure ahead of time</li><li>❑ Additional overhead of clustering, load-balancing, or service discovery</li><li>❑ 'Black Friday' nightmares if not properly planned for</li></ul>	<ul style="list-style-type: none"><li>❑ Autoscale - <b>elastic capacity on demand</b> (compute and storage) when you need it</li><li>❑ Just don't think about scale - AWS will <b>autoscale</b> it for you, including <b>DR/failover out of the box</b></li><li>❑ <b>No unexpected outages</b> due to lack of capacity</li><li>❑ Microservices <b>use the cloud natively</b> so no extra ESB overhead required</li><li>❑ No load balancing or service registry required</li></ul>
Speed to market	<ul style="list-style-type: none"><li>❑ Need to define and build data-models for all key business entities. This takes time to get right.</li><li>❑ Need to develop lookup services for many of these entities, meaning additional integration into source systems (puts a dependency on systems outside of integration)</li><li>❑ Environment creation is typically more complex and time-consuming to create and maintain if not using a serverless infrastructure</li><li>❑ Additional lookup services required – these will not be out of the box</li><li>❑ Longer test cycles as everything is new</li></ul>	<ul style="list-style-type: none"><li>❑ <b>We guarantee implementation in three months</b> (or some of your money back!)</li><li>❑ <b>Industry-ready, out of the box</b> - over 150+ industry-ready events and services available (incl. Data API)</li><li>❑ <b>Time-to-change is very fast</b> &amp; often, new data from Microservices just flows through to receiving system &amp; can be used straight away</li><li>❑ Changes to Microservices can be <b>deployed on the fly in real-time with zero downtime</b></li><li>❑ <b>Built in integrations</b> with other key cloud services – leverage ecosystems for better solutions and new innovations</li><li>❑ APIs, core business services, event processors are battle tested – <b>fewer bugs and therefore faster to market</b></li><li>❑ <b>Fully automated environment creation</b> – spin up a new dev or test environment in hours not weeks</li><li>❑ Many vendors try to claim their software enhances business agility. Xiotech's SDV genuinely does</li></ul>


# Xiotech Single Data View - Integration

How is it different to traditional ESBs such as Mulesoft:

Capability	ESB (e.g. Mulesoft)	SDV Microservices
Cost	<ul style="list-style-type: none"><li>❑ Usually more - Pay, usually too much, for over provisioned capacity 'just in case', and catering for things like DR (a whole separate environment!)</li><li>❑ Large procurement processes and capex investment to flex infra</li><li>❑ Headcount or contractors cost to manage/operate ESB</li></ul>	<ul style="list-style-type: none"><li>❑ <b>Pay per use</b> – either per request or per time period (millisecond / hour). No overprovisioning, just pay for what you use.</li><li>❑ <b>Fully managed service</b> – no operations resource required</li><li>❑ <b>Very low TCO</b> – we are never beaten on price</li><li>❑ <b>Full DR, failover &amp; resilience</b> all within the price</li></ul>
Understanding	<ul style="list-style-type: none"><li>❑ Logging, monitoring, and alerts will likely be possible but they will need to be specified and configured for every service</li></ul>	<ul style="list-style-type: none"><li>❑ <b>Deep understanding</b> of what is happening in every service. Detailed monitoring and alerting</li><li>❑ <b>Detailed history</b> of entities (order, customer etc.) as a trail of events with full context of what changed</li></ul>
Future/R&D	<ul style="list-style-type: none"><li>❑ Relies on Mulesoft to develop product roadmap</li><li>❑ If you want to upgrade your ESB, you take what is on offer from the vendor, which may include breaking changes and will often be at significant extra costs</li></ul>	<ul style="list-style-type: none"><li>❑ Microservices framework leverages a huge <b>open source community</b> including likes of <b>LinkedIn, Facebook, Google, Uber</b></li><li>❑ Our roadmap is directly linked to our customer requests</li><li>❑ We'll evolve your services for you and <b>automatically keep you on the latest version</b> so you can <b>benefit from framework enhancements</b></li></ul>



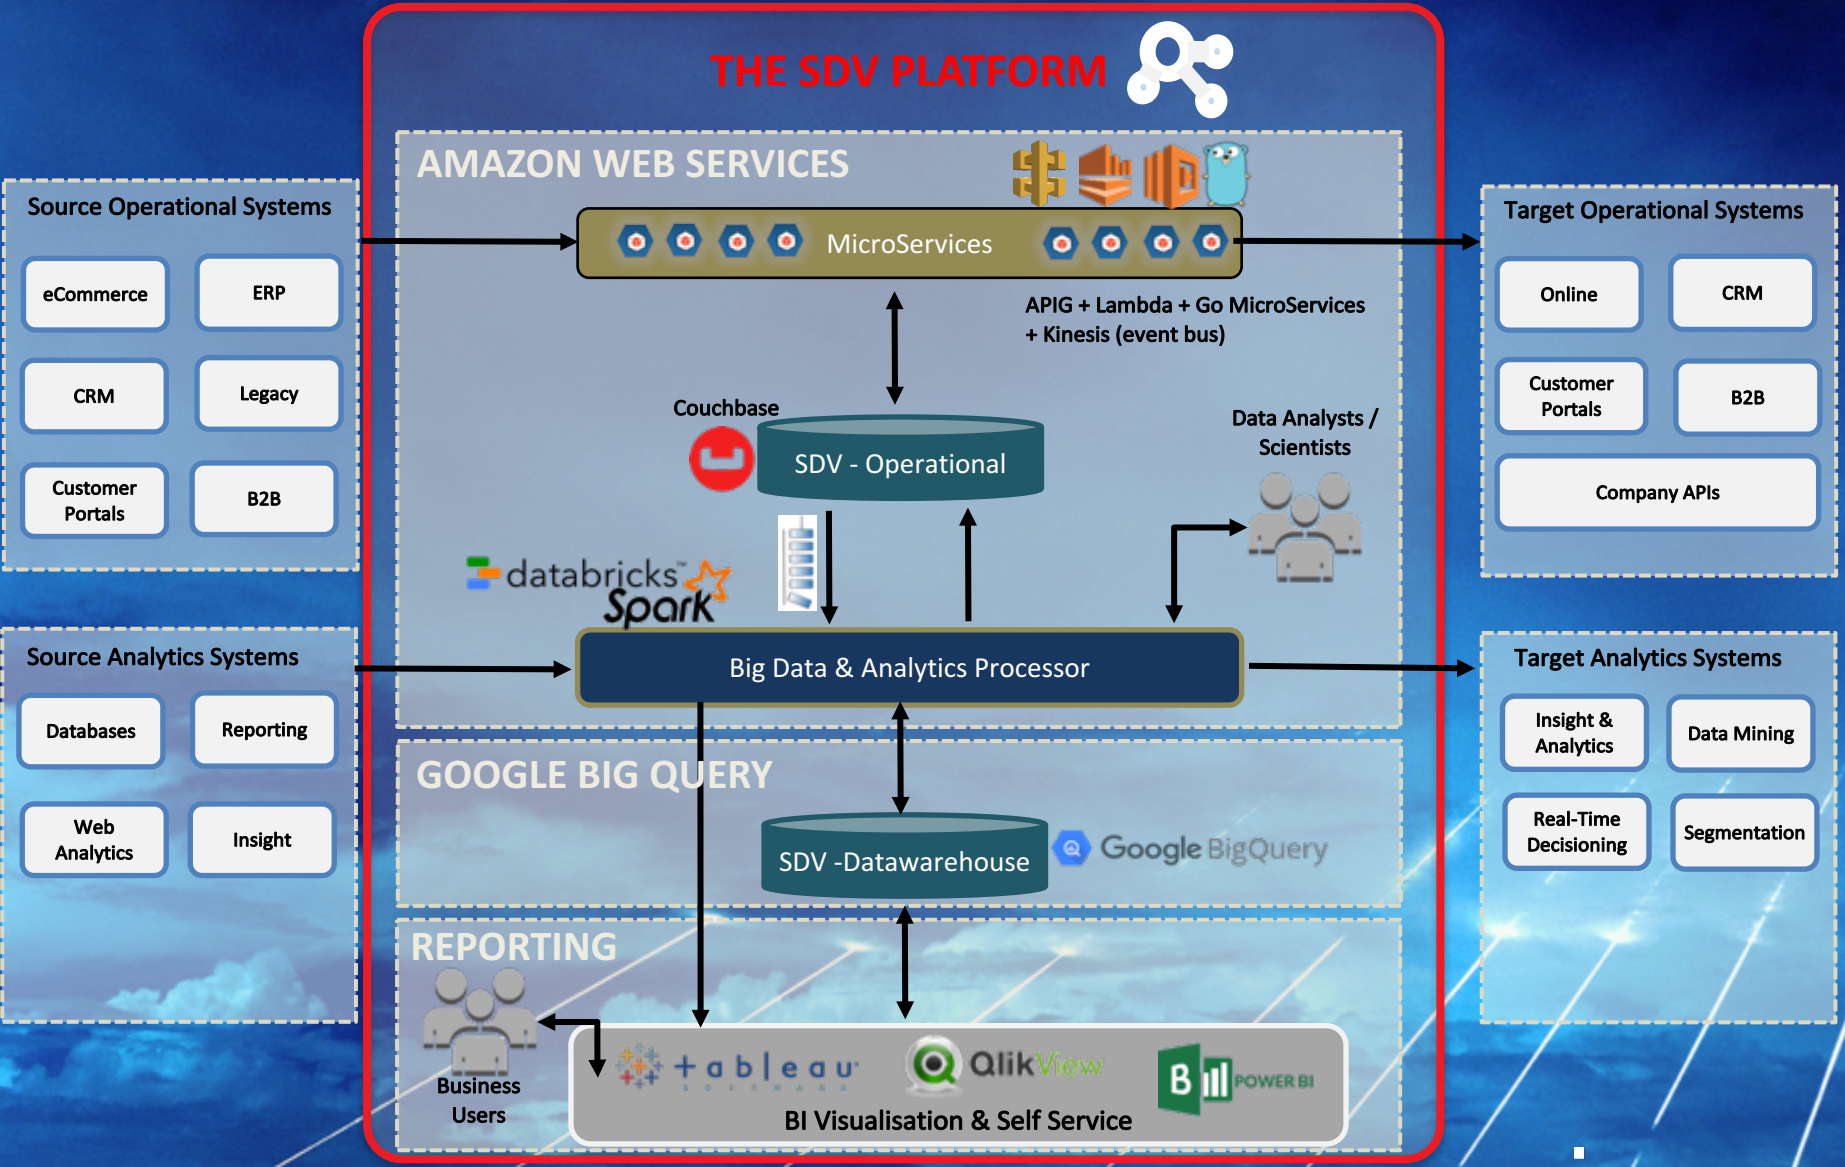




Single Data View® Platform

# SDV Data Layer

# SDV Architecture Overview



# Xiotech Single Data View - General

- ❑ **Industry Ready Data Models** - the SDV contains models for customers, orders, products, price, stock, and purchase orders among others. On the rare occasions where these don't match your models we will modify them for you at our cost as part of our roadmap
- ❑ **Industry Ready Events and Handlers** – many of the business events in a company are common across industries (CustomerRegistered, OrderAccepted, ProductUpdated etc..). We have over 150 pre-configured events and handlers to manage the reference data. All can be tailored to your specific use case
- ❑ **Reusable Reports and Dashboards** – The more we are able to reuse the data model and the events, the more reports we will be able to reuse
- ❑ **GDPR Compliant** – The SDV is GDPR ready, and can be used as a mechanism to track and enforce GDPR compliance in other systems



# Xiotech Single Data View - Operational

- ❑ **Built on the Xiotech MicroServices Framework** - the SDV builds upon the benefits discussed previously, and adds the following:
- ❑ **Operational Single View** – a NoSql document database to provide access to the single view reference data (e.g. Customer, Product, Order, Inventory, Price etc..) and the event history
- ❑ **Event Processor** – configurable pipelines of tasks to handle each event type. Pipeline tasks include storing events, creating / merging / deduping / updating reference data, and republishing based on config and rules
- ❑ **SDV API** - access the SDV reference data and events, including insight (e.g. customer propensities for personalisation)



# Xiotech Single Data View – BI & Reporting

- ❑ **Real-time Distributed Data Processing** – Apache Spark-based routines subscribe to all event streams and populate the data lake. Additionally, modern data science techniques (machine learning, predictive analytics, sentiment analysis) can be used to derive and infer new data
- ❑ **Real-time Data ‘Warehouse’** – a modern data lake that is event aware and maintains a deep history of events and changes to data. Updated in real-time, it enables fast decision making in addition to data mining and self-service reporting using BI platforms such as Tableau
- ❑ **Data Round-tripping** – write inferred or derived information and aggregated statistics back to the operational database for use by other systems



# Xiotech Single Data View – Data Lake

- ❑ **Real-time Data ‘Warehouse’** – The business benefit that a client will receive from a cube-based Warehouse, is just the minimum we can give from day one. We have found with existing clients who we are migrating from cubes is that we also deliver the top 10 cross-data source reports they could never achieve with the data e.g. Cross cube analysis.
- ❑ **Real-time Data** - The SDV Data lake and BI solution does not rely on aggregating the data into chunks (cubes) but is delivered over the Raw Data in real time. Think of it everything is in the cube, always!
- ❑ **Reporting Tool** – We use Tableau as it currently supports this raw data much better than PowerBI and other tools as you have to move the data to PowerBI at the cost of time, CPU and loss of data details. We can develop the reports in PowerBI at a small cost (we can discuss how we could fund some of this), but we always suggest looking at our demo to see what has been developed in Tableau first



# Xiotech Single Data View – GDPR

- ❑ **Data encrypted at rest** – The data within our SDV platform is fully encrypted at rest using a variety of techniques and technologies such as HashiCorp Vault, so that it is not directly human readable
- ❑ **GDPR reporting** – Using the SDV technology stack, we have reports that can show the lineage and source of any customer data request as well as the full detail of the customer records\*. The reports will also show performance data in relation to Data Compliance such as the number of GDPR-related Subject Access & anonymisation requests
- ❑ **GDPR 'Forget-Me'** – The SDV technologies provide a method of real-time integration from your source systems. If requested, we are able to 'reverse' the integration to either delete, anonymise or hash (hide) customer data on request. This action creates an audit trail for reporting but also allows this action to be reversed.

\*This relies on integrating your source systems into the SDV platform in order to create a single central view of customer information



# Xiotech Single Data View – Data Lake vs Cubes

What about traditional ‘cubes’?

- ❑ Cubes are a performance ‘means-to-an-end’ in the ‘old world’, predefining multi-dimensional arrays to ease the burden and processing demands of complex analytic queries.
- ❑ While serving an important purpose in previous architectures, it suffers some key pain-points that inhibit time-to-insight:
  - Cubes are built / materialised upfront based on current assumptions and needs at a point in time. What happens when this changes?
  - Choose which data you want up-front, discarding the rest of the data. What happens when the business want a new dimension added?
  - Costs and length of time to build / change a cube before it can be used (projects to deliver, overnight jobs to rebuild every night) can be costly when requirements change
  - Data latency – with lag in the data in the cube, I cant answer ‘what is happening now’?
  - An extra layer that needs to be integrated and tested as per any software development
- The SDV data lake negates the need for any cubes and provides the answers immediately even when requirements change





# Xiotech Single Data View – Data Lake vs Cubes

Our approach – ‘Serverless Event Driven Analytics’ - no cubes and NO discarding of data!

- ❑ Events, in the operational SDV, are core – its just the same data being used differently
- ❑ Rather than building predefined cubes as the data layer for analytics (Analytics-on-write), we prefer deferred ‘Analytics-on-read’ – don’t lock yourself out of data at all.
  - Just query the raw events and refdata in our BigQuery data lake, in its natural form, on demand when needed, however big or small.
  - Our facts are our raw events – orderPlaced, orderItemDespatched, customerRegistered, itemAddedToBasket etc..
  - Reference data is changed at a point in time by an event
  - We create views (not materialised) for reusable SQL Queries to represent data and derive insights for Tableau.
  - Tableau and BQ working in unison.
- ❑ We think of all our datasets as just one big virtual cube – i.e. all our data, at our disposal, query as much or as little you need, with scale not a concern.



# Xiotech Single Data View - BI

How is it different to traditional Warehouses? More detail:

Concern	Old World (On Prem, Relational)	New World (Cloud, BigQuery etc.)
Data Latency	<ul style="list-style-type: none"><li>❑ Daily - Data most commonly is ETL-ed over during nightly batch jobs (sometimes longer if errors occur and the window is missed).</li><li>❑ Decisions are based on yesterdays data, key opportunities missed.</li></ul>	<ul style="list-style-type: none"><li>❑ Seconds / Minutes – Data moves as a continuous stream of events until it lands in the target systems, including the cloud DW where it can be consumed instantly for reports.</li><li>❑ Decisions can be near real-time and more opportunities can be taken advantage of.</li></ul>
Complexity	<ul style="list-style-type: none"><li>❑ Usually complex – from the (multiple source) data ingest to DW via ETL, to transformations within the DB itself and the extra infra needed.</li><li>❑ Natively relational only, not well suited for modern event driven world, with nested / hierarchical data. Very complex to ‘flatten’ Json (schemas) into tables.</li><li>❑ Usually complex infrastructure to manage, and expensive consultants to execute</li></ul>	<ul style="list-style-type: none"><li>❑ Much simpler – ingest the raw event streams (already cleaned event data) generated by the SDV MicroServices straight into the DW, removing complex ETL burden from DW.</li><li>❑ Same data sets used across SDV Ops and DW</li><li>❑ Use a cloud-native DW that is designed for modern datasets – BigQuery just accepts hierarchical Json data and you can SQL over it still!</li><li>❑ No / very little ops – leave heavy lifting to the cloud providers and their SRE teams.</li></ul>
Scale	<ul style="list-style-type: none"><li>❑ Usually hard, time consuming and expensive From capacity planning to provisioning infra ahead of time, the lead times can be long.</li><li>❑ Limited by your own skills and expertise to scale effectively, or that can be affordably sourced.</li><li>❑ ‘Black Friday’ nightmares</li></ul>	<ul style="list-style-type: none"><li>❑ Autoscale - elastic capacity on demand (compute and storage) when you need it</li><li>❑ Just don’t think about scale - let the likes of Google / AWS autoscale it for you</li><li>❑ No unexpected outages due to lack of capacity.</li></ul>

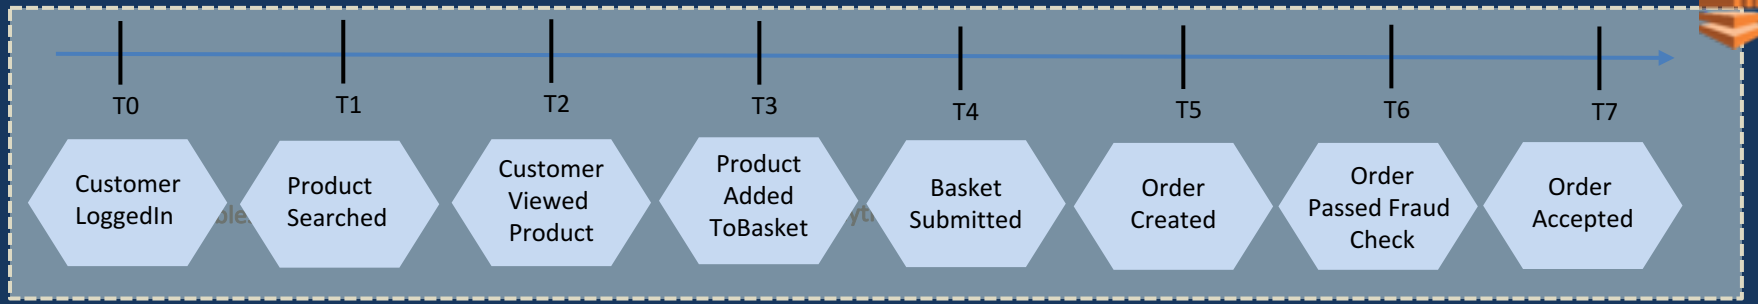
# Xiotech Single Data View - BI

How is it different to traditional Warehouses? More detail:

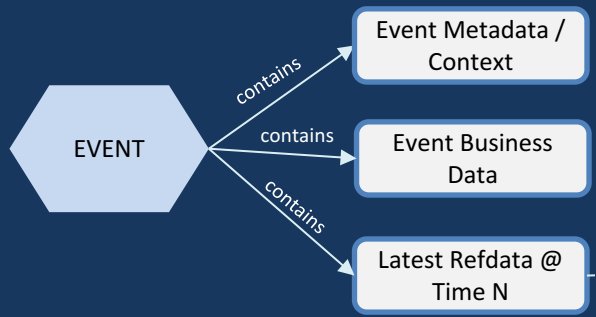
Concern	Old World (SQL Server / OBIEE)	New World (Big Query, Spark)
Flexibility	<ul style="list-style-type: none"><li>❑ Time-to-change is lengthy, often involving projects to deliver changes across the stack</li><li>❑ Schema changes tend to be rigid and painful and have ripple affects across the various DW layers.</li><li>❑ DW are usually 'siloesd' away from the operational world, and integrations can feel bloated and complex. Often changes have to span multiple teams.</li></ul>	<ul style="list-style-type: none"><li>❑ Time-to-change fast AND often, for example, new data in existing events (e.g. extra fields, objects) from MicroServices just flow through to DW automatically and be used straight away.</li><li>❑ DevOps new streams in days if they don't exist.</li><li>❑ Schema changes much easier – apart from having the same (hierarchical/nested) dataset, BigQuery evolves the table schema for you when it sees new fields in the stream!</li><li>❑ Built in integrations with other key cloud services – leverage ecosystems for better solutions and new innovations.</li></ul>
Cost	<ul style="list-style-type: none"><li>❑ Lots more - Pay, usually too much, for over provisioned capacity 'just in case', and catering for things like DR (a whole separate environment!)</li><li>❑ Large procurement processes and capex investment to flex infra</li><li>❑ Headcount or contractors cost to manage/operate infrastructure</li></ul>	<ul style="list-style-type: none"><li>❑ Pay per use – either per request or per time period (second / hour). No overprovisioning, just pay for what you use.</li><li>❑ Little to no infrastructure skills needed. Googles SRE team make sure things work and are secure and reliable (e.g. built in DR)</li></ul>
Licensing	<ul style="list-style-type: none"><li>❑ Yuk - Often licensing lock-in (we paid for 3 years upfront) and limitations (only x cores can be used in total)</li></ul>	<ul style="list-style-type: none"><li>❑ No commitment – PAYG service, stop it if/when you want to and pay no more.</li></ul>

# SDV - Continuous Event Stream


## Time series of events



## Event Structure



## Ref Data Evolution – Immutable, timeseries snapshots in the DW



- Events contain some key metadata (e.g. category, name, time, event time etc..), the actual event data (e.g. an order item to be cancelled) and (critically) the **latest** snapshot of the order POST processing. This is all done by the MicroServices!
- This event data gives the warehouse everything it needs to record the event, and also that latest snapshot post processing (i.e.. the 'refdata')

- The refdata record is immutable and timeseries – we do not change it (unless for exceptional circumstances), it's a fact of what the entity looked like at that point in time.
- A subsequent change (from another event) then appends another record (with a newer timestamp) into BigQuery – this is then the new 'latest'. Thus changes are tracked natively over time

# Xiotech Single Data View – BI & Data Lake

Our approach – Google BigQuery/Spark are key

- ❑ Google Big Query is key, game changing technology - a ‘serverless’ cloud warehouse, with near infinite scale in both storage and compute.
  - ❑ There is NO infra to manage, it is ‘no-ops’ and we simply don’t need to think about it (of course, you still need to tune things like your queries!)
- ❑ It is SQL compliant but with some VERY KEY extensions around nested / hierarchical data.
  - ❑ It can natively ingest our JSON events as-is – no ugly complex ETL to unpack into traditional flat tables.
  - ❑ It evolves the schema dynamically when detecting new fields
- ❑ Spark streams the raw events/refdata into BQ, which itself can do other more advanced analytics on the stream enroute (e.g. machine learning predictive etc.) and interact back to the operational layer (e.g. surface a recommendation)
- ❑ The kicker is its seriously cheap!
  - ❑ \$5USD per TB scanned in queries, pennies per GB stored.....plus there is the free tier
- ❑ With the performance of BQ, we can (and do) perform complex queries on the data in its raw form. It’s really THAT SIMPLE



# Xiotech Single Data View - BI



Single Data View® Platform

Our approach – Tableau for reporting / self-service

- ❑ Tableau as a market leader in self service BI, reporting and ad hoc data exploration is our preferred tool
  - ❑ It has a mature native BigQuery connector, and we have a suite of reports that are built upon our standard data model for reuse.
  - ❑ Such reuse will enable rapid acceleration & adoption of reports like ecommerce trading packs and stock visibility etc..
- ❑ If we were to use Power BI as an alternative to Tableau, the following would be the key considerations:
  - ❑ There is a *beta* BigQuery connector released in August, however it is a beta. Xiotech would need to do a proof of concept on this connector. We strongly believe in BigQuery and not suggest changing this as the DW, so proving this out for Power BI would be key
  - ❑ Current suite of reports would need to be built in Power BI
  - ❑ Power BI is a cheaper option
  - ❑ A second alternative we would recommend is Google Data Studio. This is totally free, and integrates well with BigQuery
  - ❑ If PowerBI is preferred, then we can discuss how we might fund this development as part of our roadmap internally

# Technologies in the Xiotech SDV Platform



- The framework and MicroServices are written in the [Go](#) programming language
- For executing this code, we use [AWS Lambda](#) (Function As A Service)
- For Web API triggers (including public facing), we use [Amazon API Gateway](#)
- For Event / Message triggers we use [Amazon Kinesis Streams](#)
- For Scheduling service invocations, we use [Amazon Cloudwatch Events](#)
- For Log Aggregation, Search, Monitoring, and Alerts, we use [Amazon Cloudwatch](#)
- For Environment creation and configuration, we use [HashiCorp Terraform](#)
- For securing passwords, API keys, and encryption keys, we use [HashiCorp Vault](#)
- For build and deployment, we use [Concourse](#)
- For the operational SDV, we use [Couchbase Server](#)
- For the Distributed Data Processor we use [Databricks](#) (managed [Spark](#))
- For the data warehouse, we recommend [Google BigQuery](#)





# Case Study



Single Data View® Platform

**xia**tech:  
consulting



# Fashion Retailer

## Problem:

- ❑ A global footwear retailer, wished to replace their e-commerce and campaign management platforms and add a new Business Intelligence capability. These systems would need to be integrated within 3 months to each other & into two legacy systems.

## Solution:

- ❑ Xiotech implemented their SDV suite of reusable MicroServices to integrate the new e-commerce platform (Hybris) with the 10-year legacy ERP (Navision), the old eCommerce platform, new Campaign Management system (Redeye), customer services (ZenDesk) & various spreadsheets within three months.

The SDV was also used to surface data from all these systems for BI reporting using the Xiotech SDV platform.

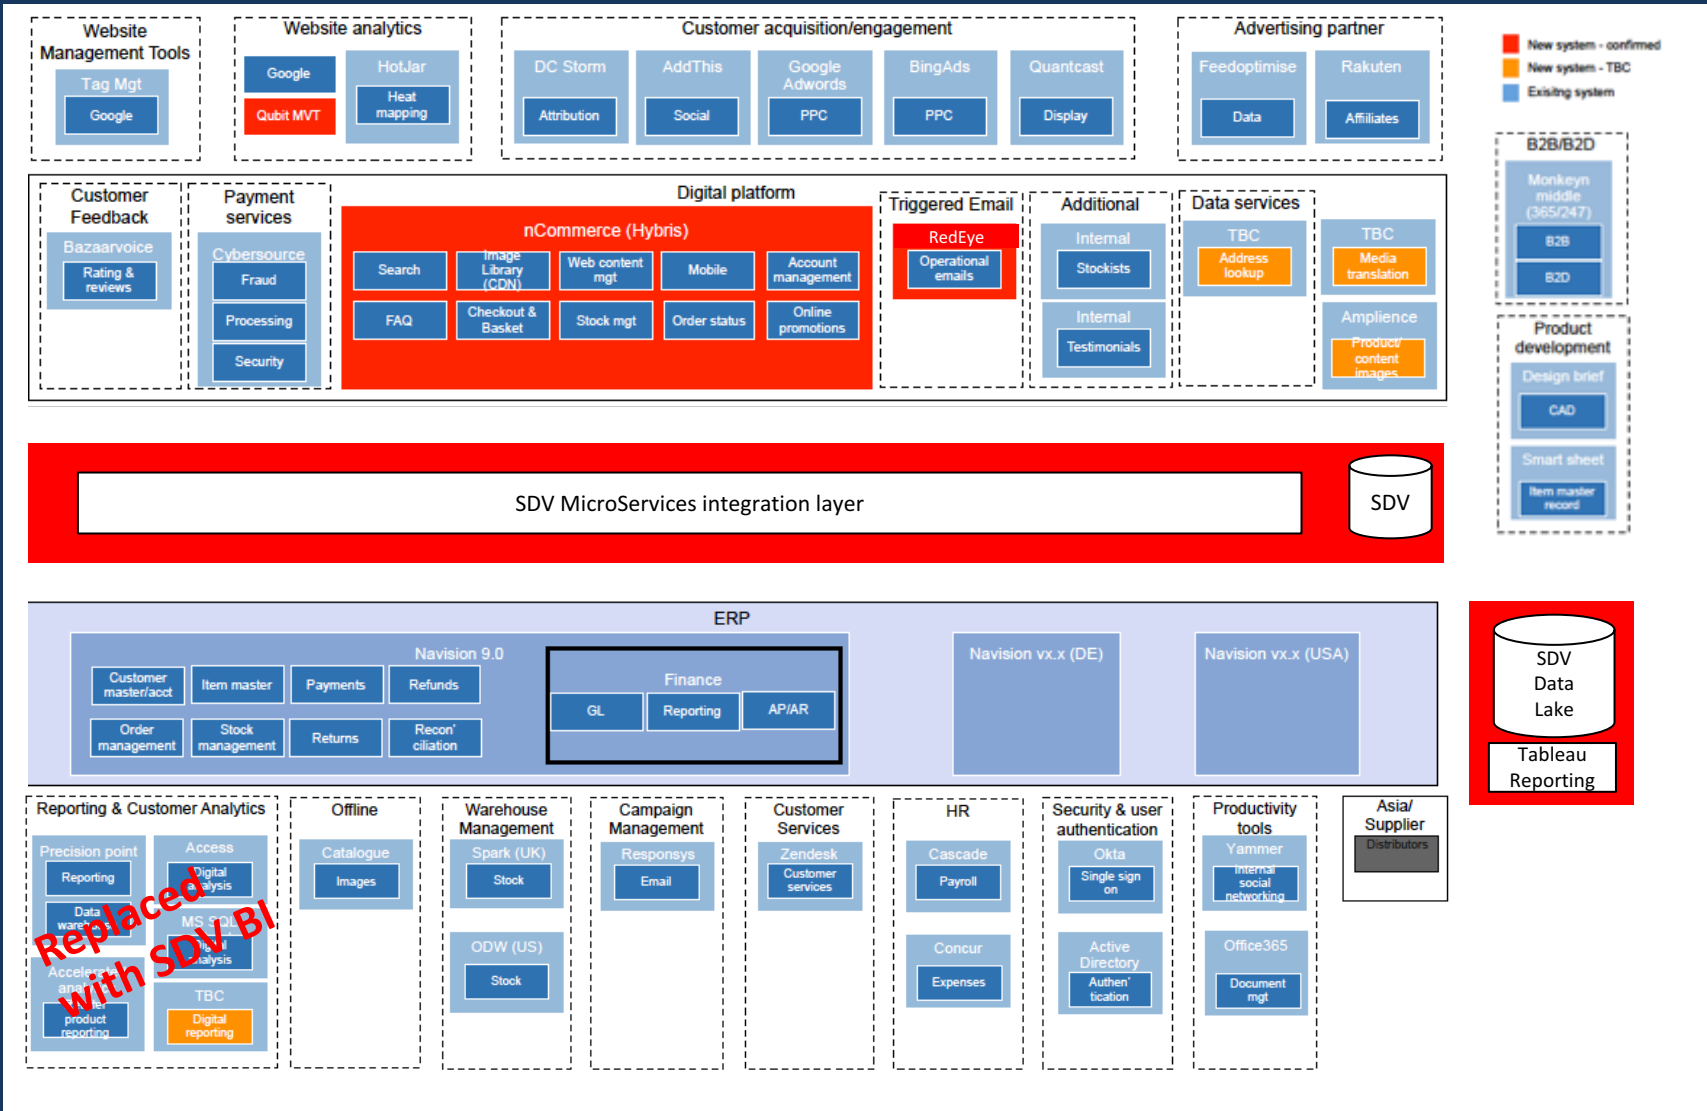
This includes:

- ❑ Full details on Product, Price, and Inventory as well as detailed Merchandise Trading packs
- ❑ Customer Registration and key events (profile updates, logins etc..) as well as full acquisition source data
- ❑ Orders – capture, despatch, cancel, return, refund
- ❑ Financial data such as sales. Discounts, margin, landed cost price etc..
- ❑ Comms – service emails and comms history


## Benefits:

- ❑ Client now has an agile integration platform where new systems/features can be delivered with unprecedented speed
- ❑ Insight into business operations is massively improved through access to previously untapped raw data
- ❑ Detailed Customer insight and analytics has been introduced into the CRM & marketing team for the first time
- ❑ Stock management which had been costing the business due to inaccurate data has now vastly improved
- ❑ Hours of inefficiency through compiling spreadsheets has now been saved, freeing up time for more analysis
- ❑ A highly flexible real-time self-service BI platform on which new reports are being added on a daily/weekly basis
- ❑ Zero minutes of downtime since go-live despite numerous releases of new functionality
- ❑ A highly scalable infrastructure with capacity on demand and automatic failover and monitoring
- ❑ Global adoption throughout the company

# Fashion Retailer - architecture




**Above: The SDV integration platform was key to our clients future state architecture and was delivered within three months**



Single Data View® Platform

# Q&A




Single Data View® Platform



**Q: Will Xiotech be involved in the architecture design workshops?**

**A:** Absolutely – we would encourage this so that we can all agree the scoping and phasing. We would also encourage the use of a regular Technical Design Authority which we will be happy to attend, in order to ensure consistent architecture governance and decision making

**A:** We have a lot of architecture experience; not just integration but enterprise, business, information, application, and infrastructure so would be able to add more value



Single Data View® Platform




**Q: How do MicroServices work when integrating into flat files?**

**A:** The MicroService framework can monitor a file system or SFTP location and process flat files when they appear.

**A:** We also create flat files and upload them to wherever they are needed. Fixed width and delimited (e.g. comma, pipe, tab) files are supported 'out of the box'.

**A:** The creation of 'copy-cat' flat files is a common method for the SDV integration framework when integrating into legacy systems, so that we can implement our SDV platform without the need to ever modify the legacy system estate.




Single Data View® Platform



**Q: How do we ensure the additional benefits don't detract from delivering a robust, strategic integration platform?**

**A:** The implementation of the SDV is incremental and typically takes an Agile methodology. Therefore, we will always start with the Integration of the systems and data flows. BI Reporting can be introduced at any time and can follow the same iterative implementation when required

**A:** The success of the SDV BI Reporting is partially linked to the quality of the data and the integration layer it gets its feeds from. We only start consuming the events into the SDV once we know the services are working



Single Data View® Platform




**Q: Why is the SDV different to more traditional data warehouses or our competition?**

**A:** Most DW platforms are not **real-time** or **event based** so they are not **contextually aware**. The SDV captures data as events & also holds structured & unstructured data. It can therefore piece together **'data-stories'** such as **Customer Journeys** or **inventory movement patterns** which are usually difficult to compile.

**A:** The **costs and time-scales** of implementation are far lower with the SDV, and the **TCO** is a fraction of the usual DW costs due to cloud hosting. Additional changes are simple and cost efficient

**A:** Our solution offers **automated GDPR compliance**, by offering the ability to anonymise personal data on command via an admin interface




Single Data View® Platform



**Q: How has the software been funded?**

A: Completely self financed mostly from cash-flow





Single Data View® Platform



**Q: How does this contribute to a Company's bottom line?**


A: Increases speed to market of new revenue generating systems e.g. eCommerce platform

A: Decreases the cost of new system implementation

A: Total cost of ownership is significantly lower than standard integration & DW platforms

A: Cost of SDV implementation is on average 10 times lower than equivalent projects

A: The analytics available offer deep insight through cross data-set analytics. This provides companies with the ability to optimise their channel trade through otherwise unavailable detail




Single Data View® Platform



**Q: What is the typical implementation timeline for the SDV?**

A: This does depends on the number of systems involved and whether we can use our standard adapters. However, it can be as quick as three months. We tend to implement the SDV incrementally and adding new systems can take a matter of days.



Single Data View® Platform



**Q: Why is the SDV so innovative?**

A: We have developed the SDV using leading edge technologies developed by the likes of Google, LinkedIn and Facebook. We have also relied entirely on the open source nature of these products, so we benefit from inheriting and contributing to all the new features and support from the community

A: The architecture is unlike any tools we have seen and the innovation in the product allows you to do things that would normally take at least three times longer to achieve

A: We have numerous reusable out of the box adapters to many commercial software products e.g. eCommerce platforms & ERP systems so implementations can be very quick

# Real-time event based integration...

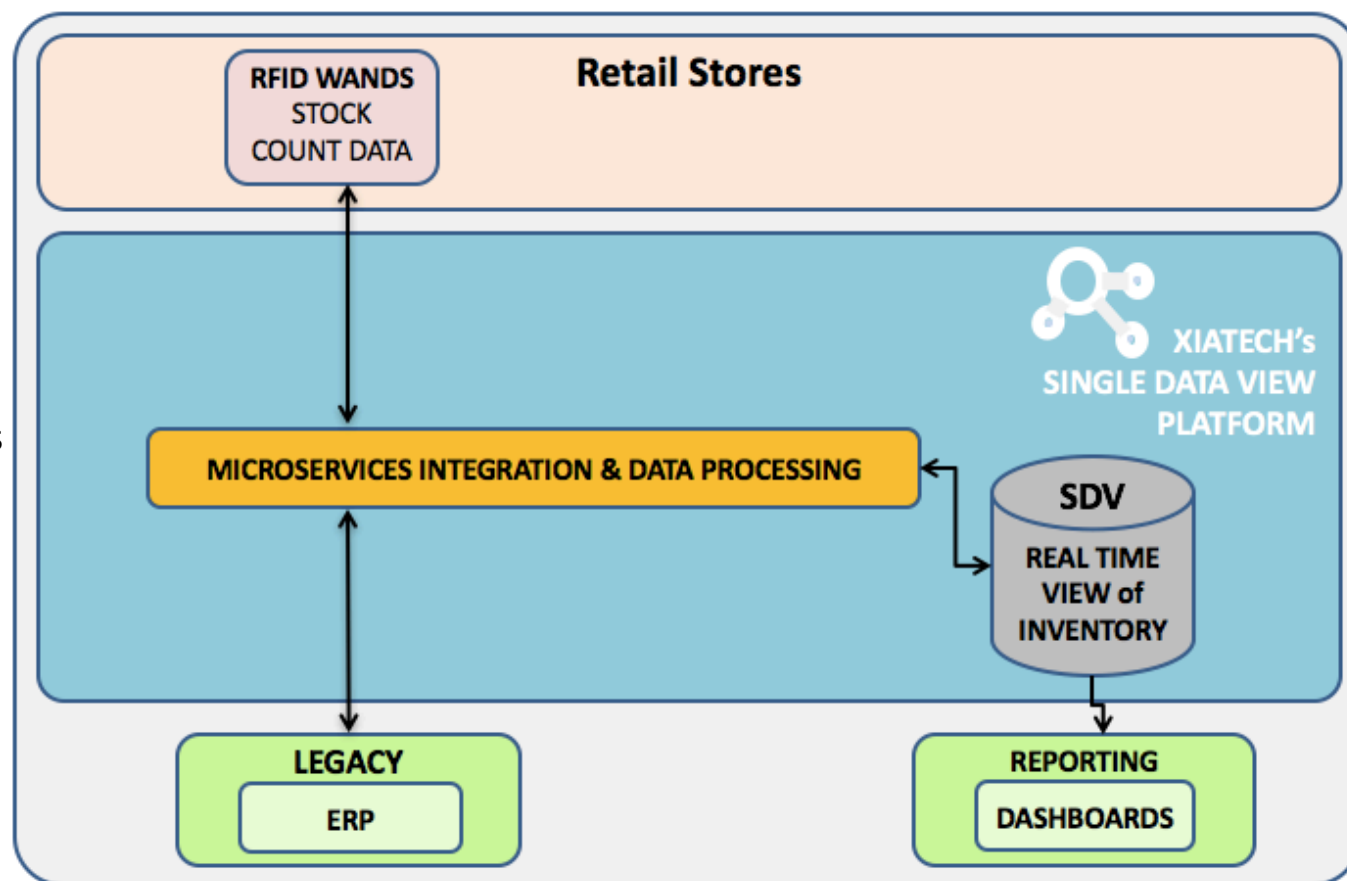
- Inventory management/visibility was the far-and-away leader (37 percent\*) of top priorities for Retail supply chain/logistics spending in 2017.
- To meet customer experience expectations, retailers must be able to locate inventory anywhere in the supply chain — be it on a store shelf, in a warehouse, or on a delivery truck.
- \* <https://www.retailitinsights.com/doc/retail-tech-resellers-where-will-your-customers-invest-in-0001>



# Case Study:

## Fashion Retailer

### RFID data exchange



- **The Problem**

RFID data produced 1000's transactions per second from stock counts

- **The Solution:**

**Xiatech's** Single view of Inventory (in the cloud) & messaging platform:

- **Rapid Design to Deployment in 3 months**

Agile development allowed rapid deployment

- **The Technology**

- License free open source technology such as GoLang, Kinesis, Databricks Spark, CouchBase
- Cloud based (Amazon Web Services and Google Cloud)
- Near **real-time** MicroServices event-based data streaming & processing technology

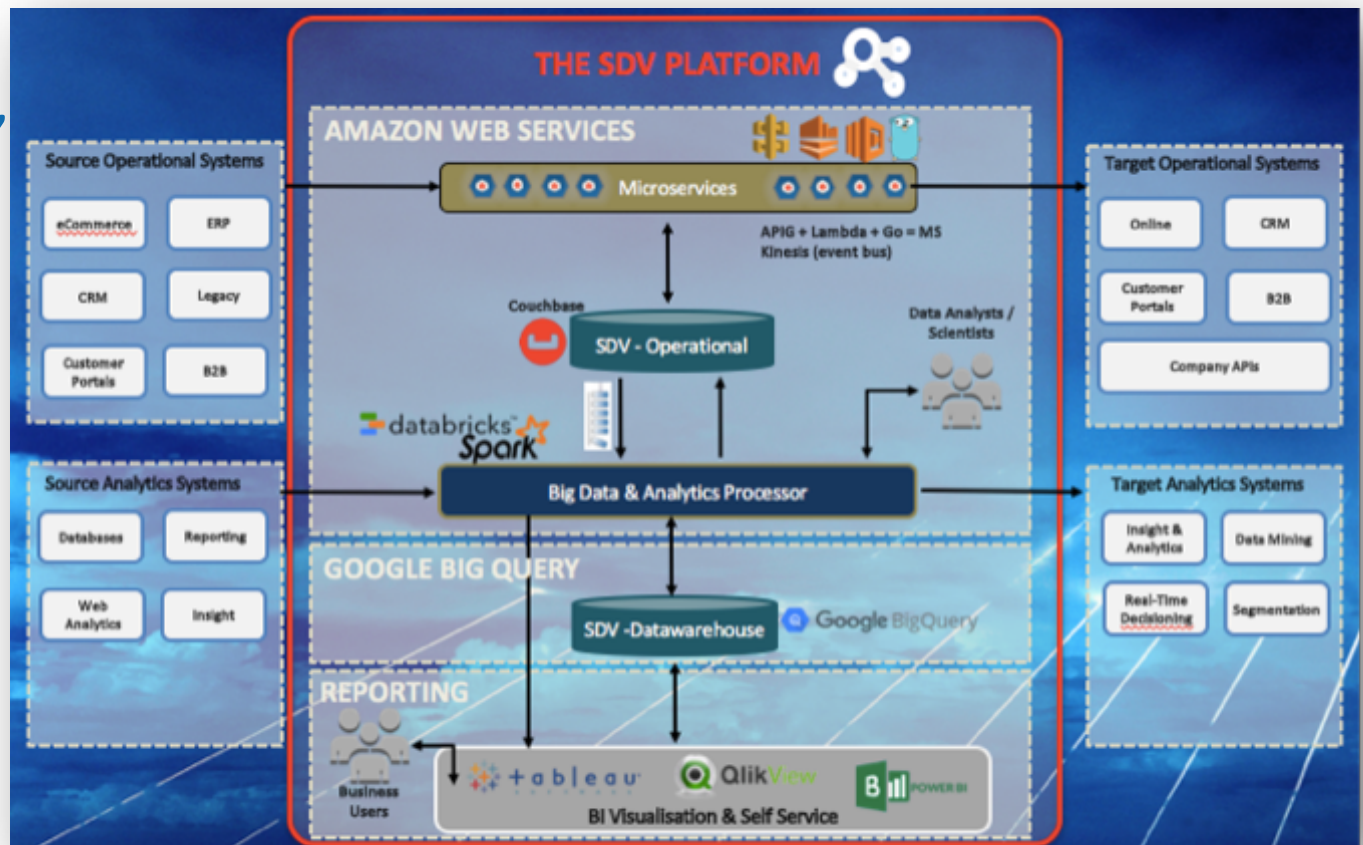
# Single view of Data...

- Centralised, consolidated single version of the truth of any data
- Easy to change or add new data
- Can expose the SDV to any system
- Real-time, event-based view of data



# Case Study: Childrens' Lottery, Scotland

## Single View of Customer



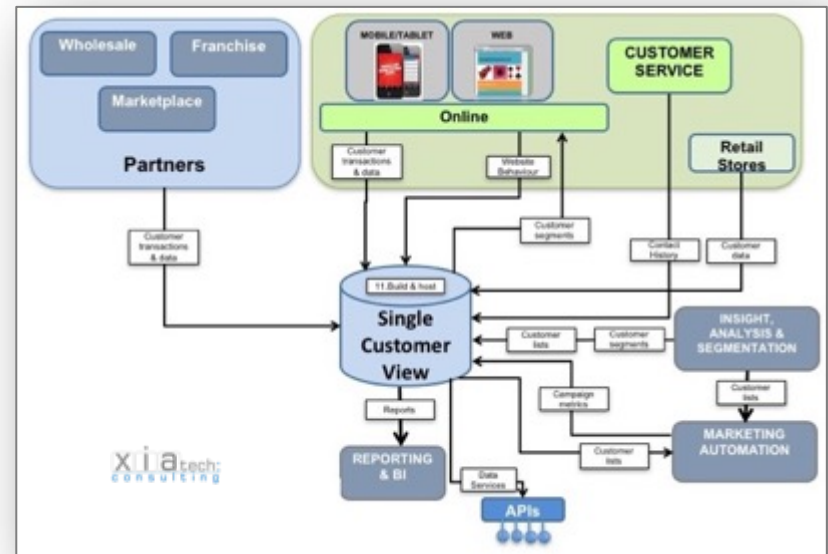
- **The Problem**  
Multiple versions of customer data across many channels  
Introduction of new Lottery system would worsen the situation
- **The Solution:**  
Single view of Customer Data (in the cloud) for insight, analytics & segmentation, for CRM & Marketing:
- **Rapid Design to Deployment:**  
New systems added within days

# Case Study:

## Fashion Retailer, London

### The Problem

- Multiple versions of corporate data across many channels
- Inconsistent data across legacy systems & reporting
- Inability to obtain insight & analytics across the customer base
- Lack of single version of truth of other data eg Inventory, Product, Price
- Information was out of date due to slow batch interfaces (at least 24 hours)
- No self service reporting capability



### The Solution: Single view of Corporate Data (in the cloud):

- Xiatech were engaged to conduct a Discovery process with a view to a future design
- The proposed architecture from Xiatech presented a real-time ingestion of data from all operational system (over 7)
- Data would then be held persistently in a central Big-Data repository, providing a unified corporate view of all data such as Customer, Product, Inventory
- The tooling would provide the ability to perform in-day, real-time reporting using standard Reporting tools
- The nature of the design would allow the solution to seamlessly scale on demand, especially in peak trading periods such as Black Friday

### Rapid Design to Deployment in 3 months

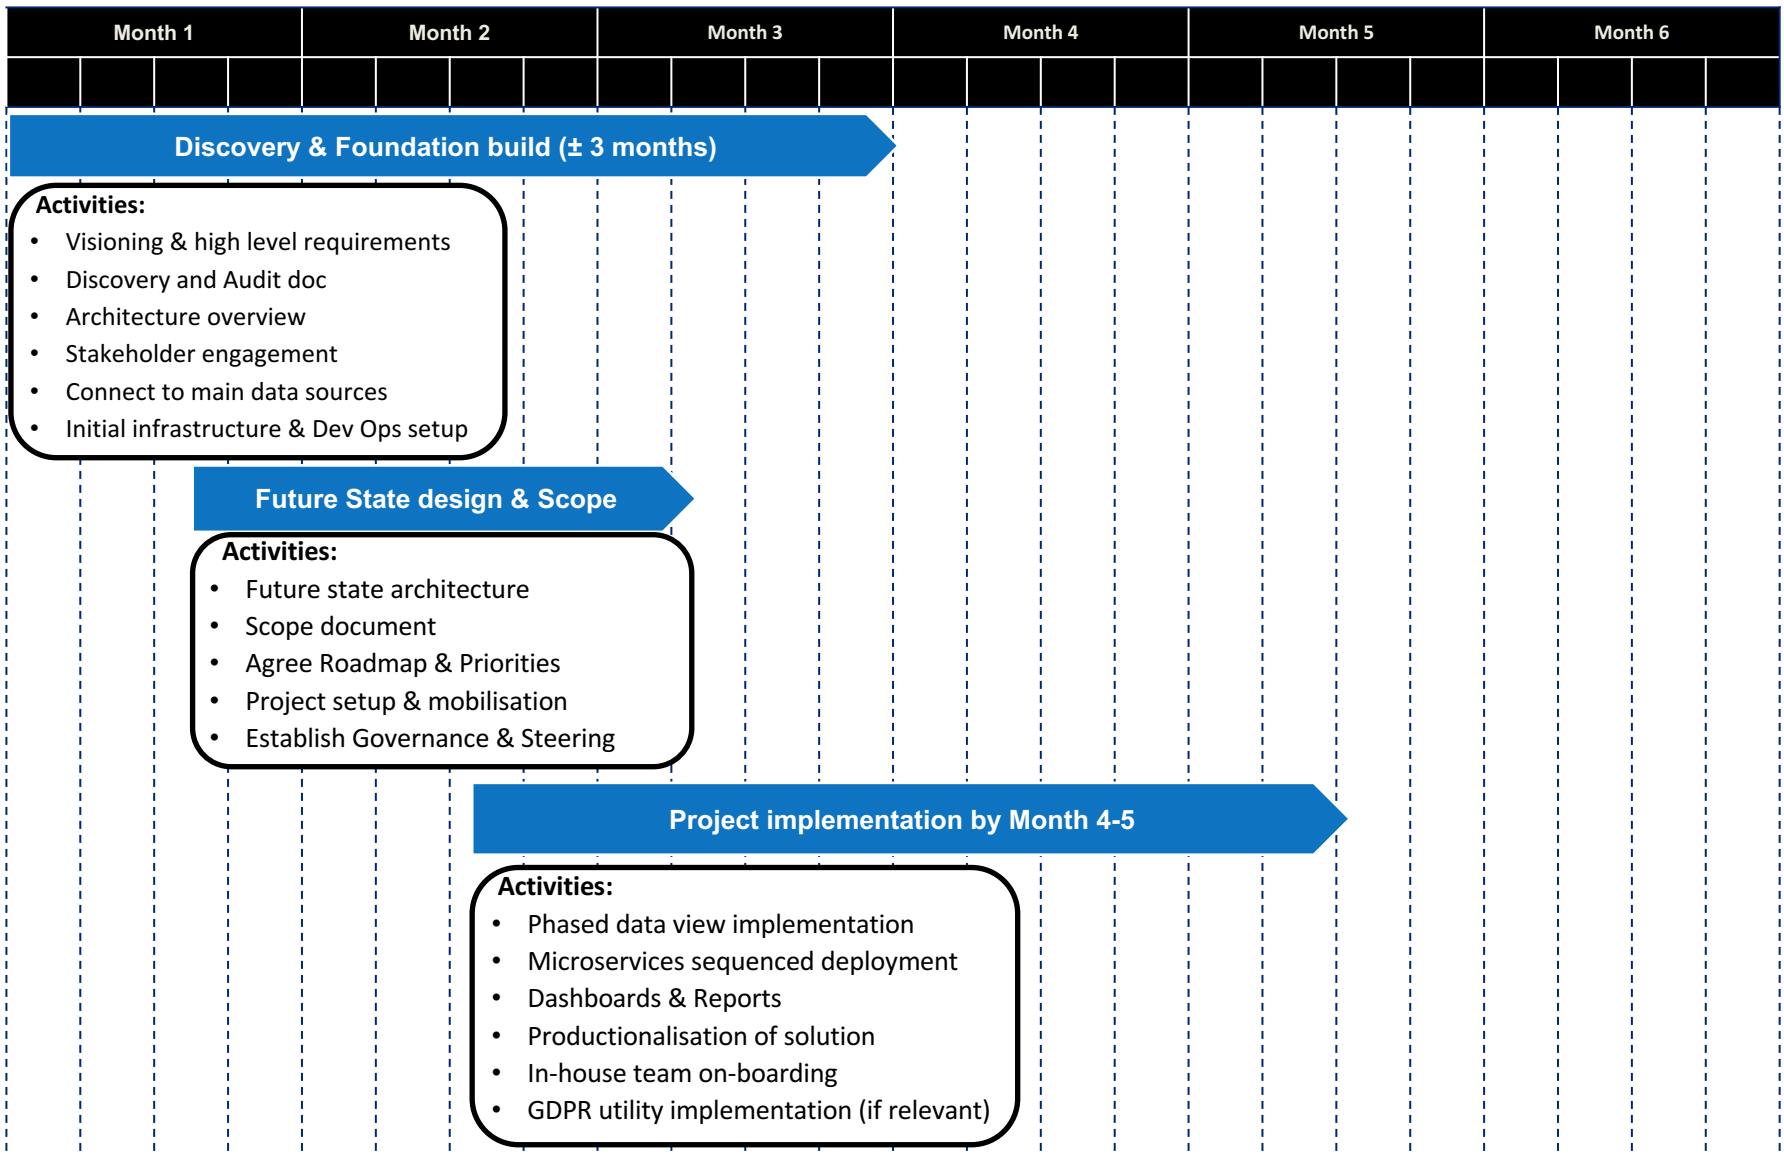
- Following a 4 week Discovery, a proof of concept was developed by Xiatech and integrated in a few weeks to 3 main systems
- The solution was then rolled out iteratively across the business (to the remaining systems) over a 6 month period, using an Agile project team

### The Technology

- License free (or nearly-free) Open Source technology such as GoLang, Kinesis, Databricks Spark, CouchBase, Google Big Query
- Cloud based & Infinitely scalable on demand (through Amazon Web Services and Google Cloud)
- Near **real-time** event-based data streaming & processing technology
- Open standard APIs (using Microservices)



# Our typical implementation approach



# Data Compliance

# GDPR

(GENERAL DATA PROTECTION REGULATION)

- Consent
- Breach
- Right to Access
- Right to be forgotten
- Privacy by Design



# Our Frameworks

Xiotech have developed a number of frameworks to help companies with numerous challenges:

- Programme Assurance/Delivery Assurance
- Design Assurance
- Change Management
- IT & Business Architecture
- IT capability
- Due Diligence
- These are now supported by a new **Data Compliance** framework in preparation for GDPR.



# Data Compliance Assurance

Xiotech have developed a framework which can be used for reviewing and developing data related programmes, governing them and ensuring they meet the new data regulations.



# General Data Protection Regulation

The UK currently relies on the Data Protection Act 1998.

It will be superseded by the new General Data Protection Regulation (GDPR) from May 25<sup>th</sup> 2018.

Simply, GDPR will give people more control over how their personal data is used.



# How it will affect you?

GDPR will affect both the '**Data Controller**' and the '**Data Processor**'.

Data must be processed lawfully, transparently, for a specific purpose and only retained for as long as is necessary.

This affects both new and legacy data.



# Why is it so important?

Currently the ICO can apply fines of up to £500,000 for contraventions of the Data Protection Act 1998.

However, under GDPR, fines can be up to **€20 million or 4% of global turnover**, whichever is the greater.


Talk Talk's record £400,000 fine in 2016 could have been £59 million under GDPR!





# Data Compliance Framework

Our Data Compliance Framework focuses on the operational and commercial activities of your business but also considers both the technical and legal implications.



# Data Compliance Framework

We use our framework to review your current programmes against each key pillar :

## Corporate Goals

Business Vision

Business Goals

Stakeholders

Commercial Implications

Technology Strategy

## Data Collection

Data Processing

Types of Data Collection

Consent Practices

Privacy Notices

Audit & Quality Control

## Data Usage

Contact Strategy

Data Supply Chain Review

Profiling & Automated Decision Making

Digital audit

Access to Personal Data

## Data Management

Suppression & Erasure

Data Retention

Privacy Impact Assessments

Archive, Portability & Transfer

Data Security & Breach Notification

## Technical & Legal

Information Governance

Data Protection Training

Industry Regulation

ePrivacy Regulation

Legal Consideration and Review

# Our Approach

We offer a three stage approach to helping you approach the journey to GDPR compliance:

**1. Discovery**

We will audit the 'what, where & how' of your data from an Operational, Commercial, legal & technical perspective

**2. Prioritise & Plan**

We will help you prioritise, plan & cost the actions required

**3. Remediate**

If required, we will manage the implementation of the prioritised recommendations



# Our software

Xiotech have developed an open-source, near-license free platform, providing companies with a **real-time single view of their data** such as customer information.

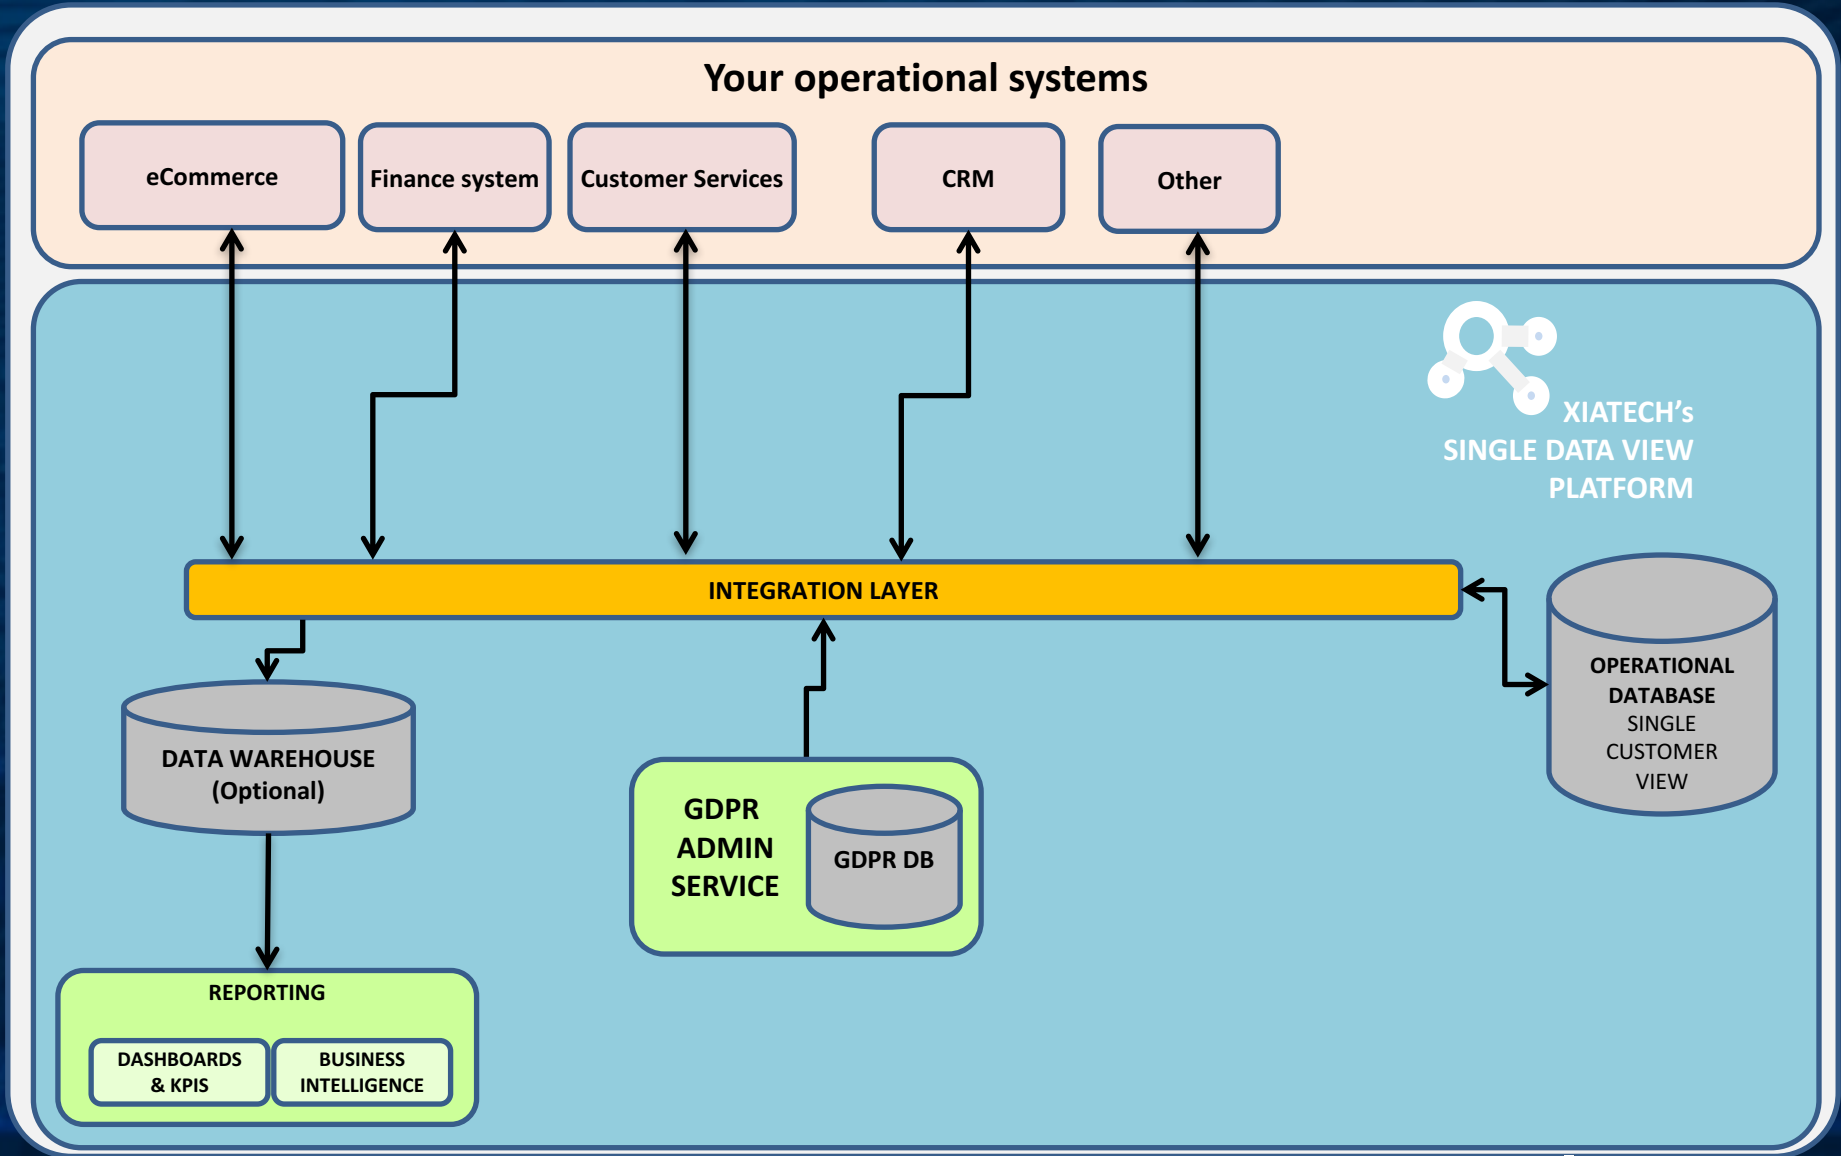
We have adapted this platform to enable an automated method of providing **GDPR compliance** to your systems which can be implemented in a few weeks



Single Data View® Platform

# Single Data View© platform

# Our software sits at the heart of your systems to provide GDPR automation



# WHY COMPANIES CHOOSE OUR SINGLE DATA VIEW<sup>®</sup> PLATFORM



**Secure &  
GDPR  
compliant**



**360° data &  
Industry ready  
real-time  
business decisions**



**Quick to  
implement with  
pre-built  
connectors**



**Cloud based &  
license free**



**Highly scalable &  
resilient**



**Single Data View<sup>®</sup> Platform**

**xia**tech:  
consulting



**Xiatech have been awarded a number of industry accolades in recognition of our work in this field**

EΥΧΑΡΙΣΤΩ TÄNAN HVALA GRACIAS DZIĘKUJĘ  
GRAZIE ありがとう MERCI TACK

**THANK YOU** DIAKUIU  
PALDIES

ACIU TANK DANKE DANK U WEL ДЗЯКУЮ

СПАСИБО 谢谢 OBRIGADO KIITOS

TESEKKUR EDERIM

diolch