

# Modern Data Architecture


Alexey Grishchenko

# About me


## *Enterprise Architect @ Pivotal*

- 7 years in data processing
- 5 years with MPP
- 4 years with Hadoop
- Spark contributor
- <http://0x0fff.com>


# How it started...



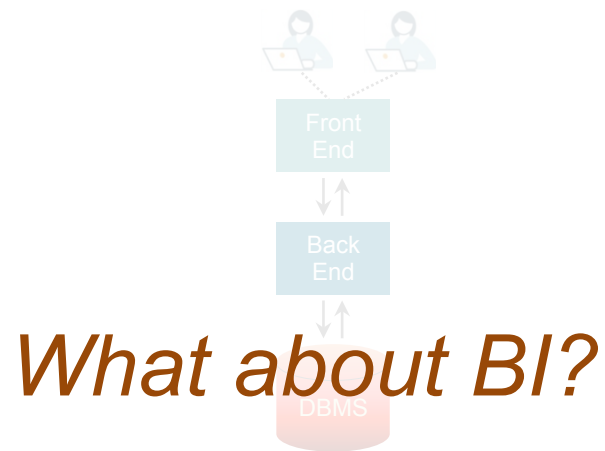
# How it started...



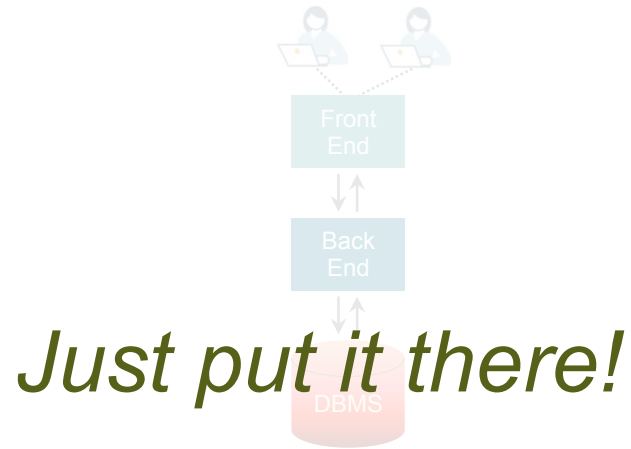
# How it started...




# How it started...



# How it started...

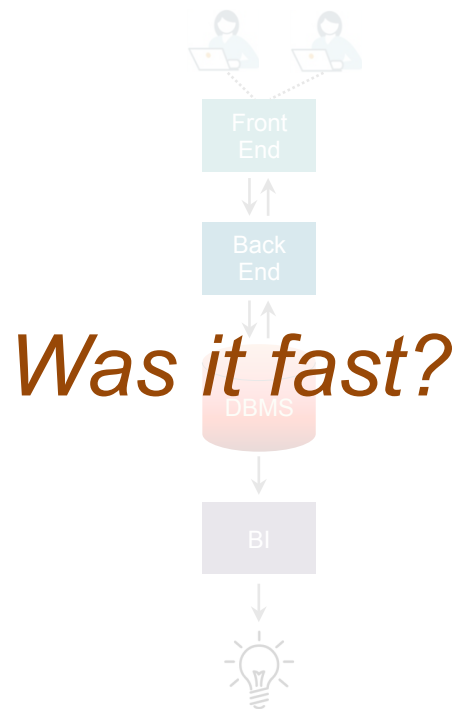


# How it started...







# How it started...



# How it started...

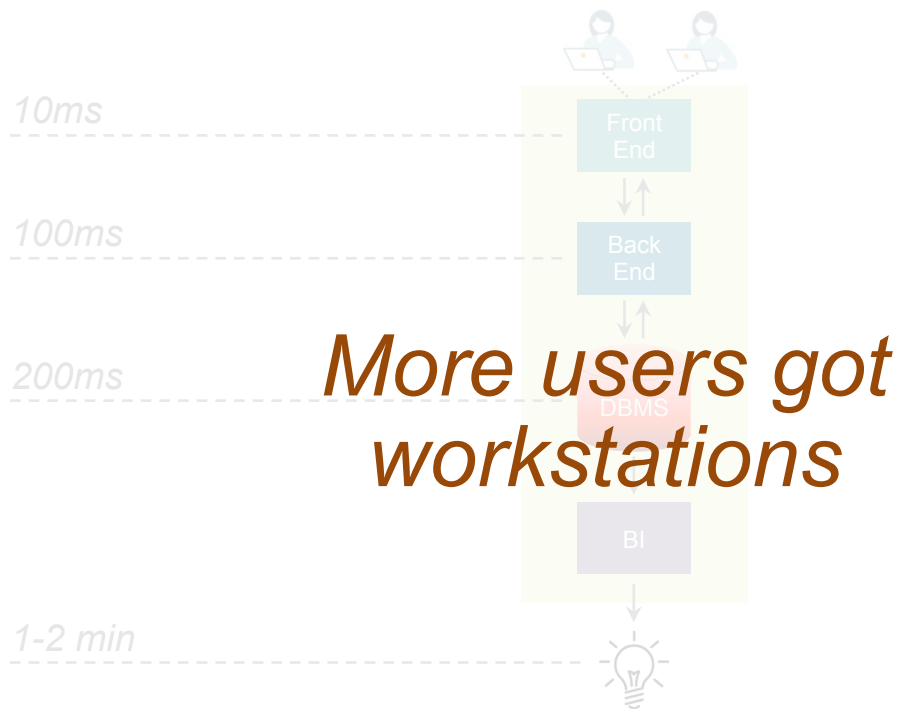


# How it started...




*yes, single server...*

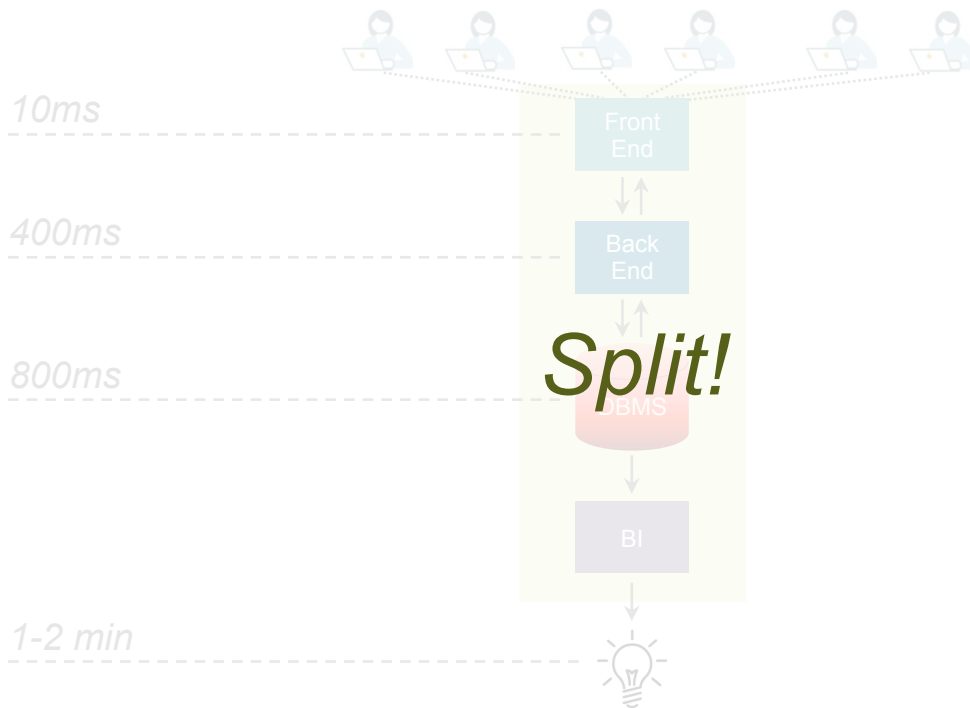
# First Issues



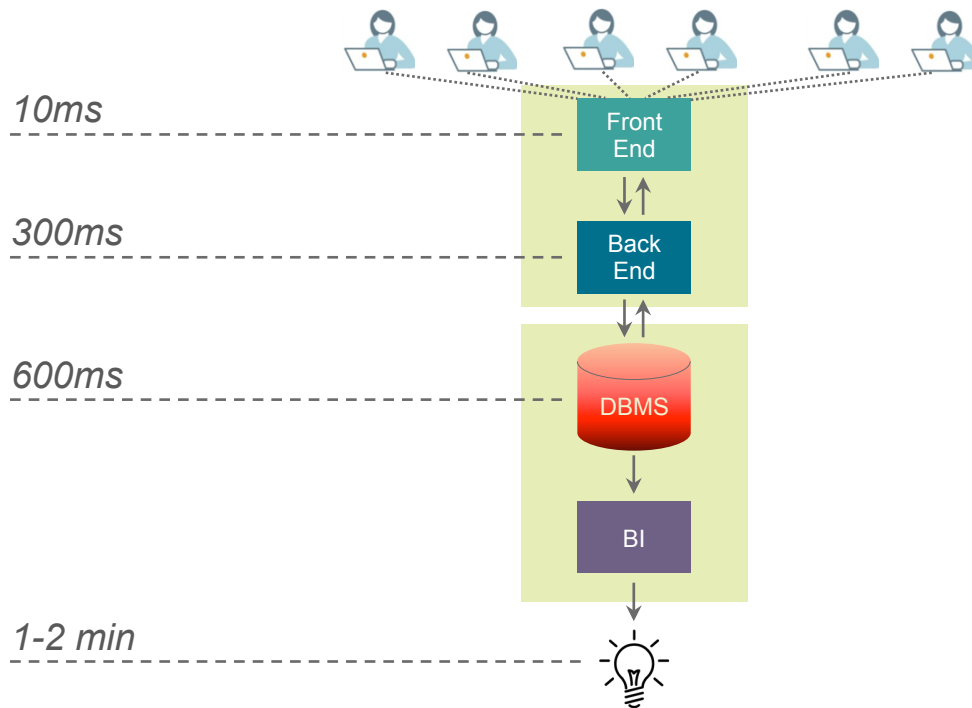
# First Issues



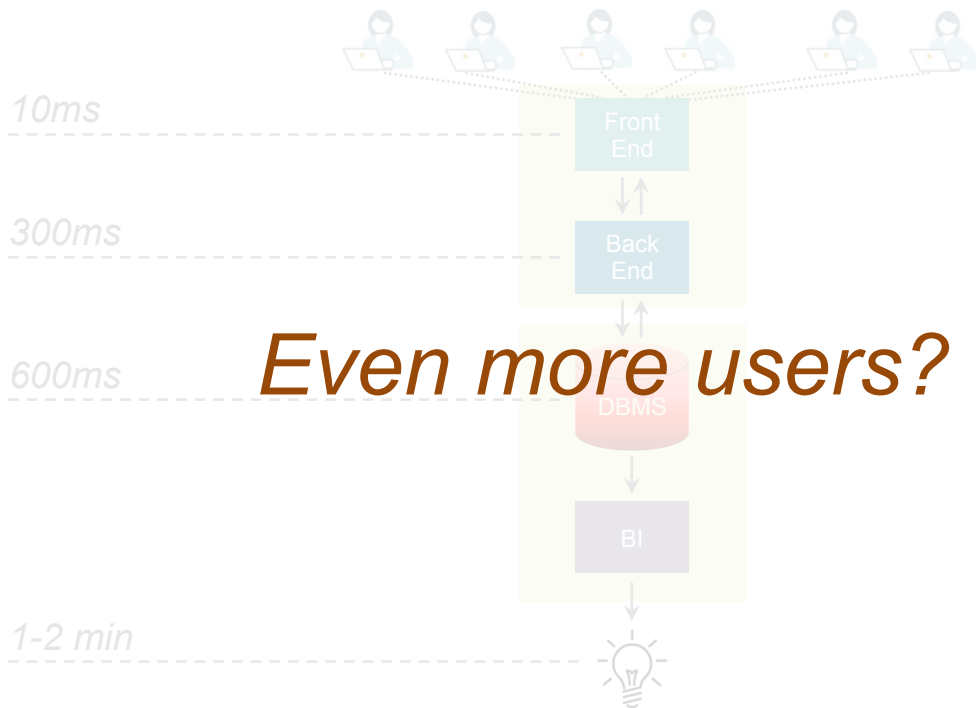
# First Issues



# First Issues




# First Issues






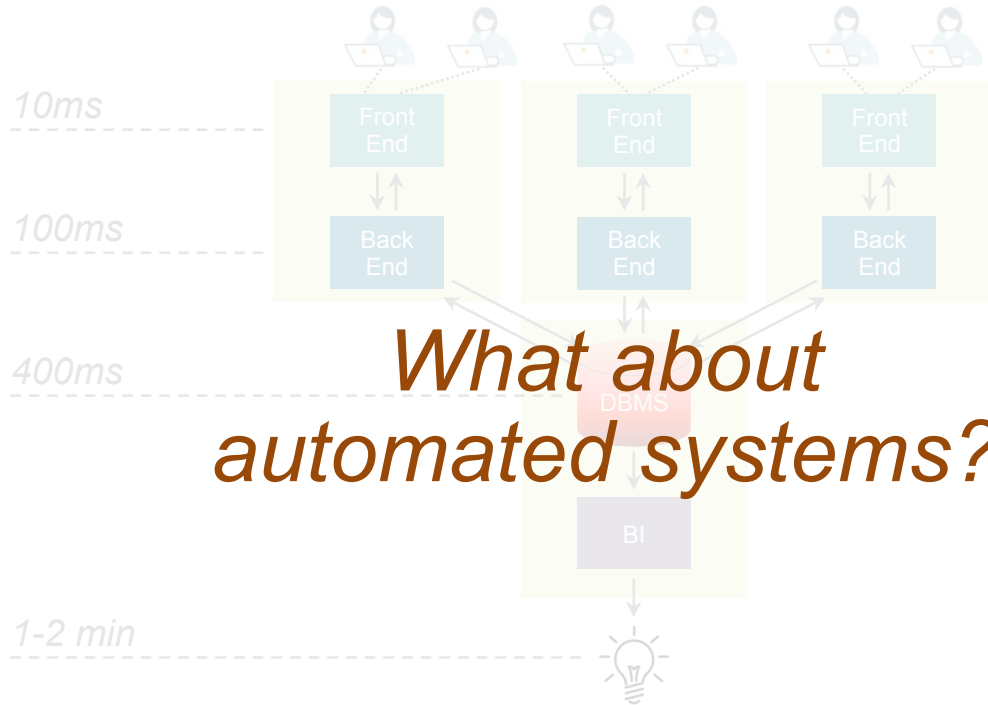
# First Issues



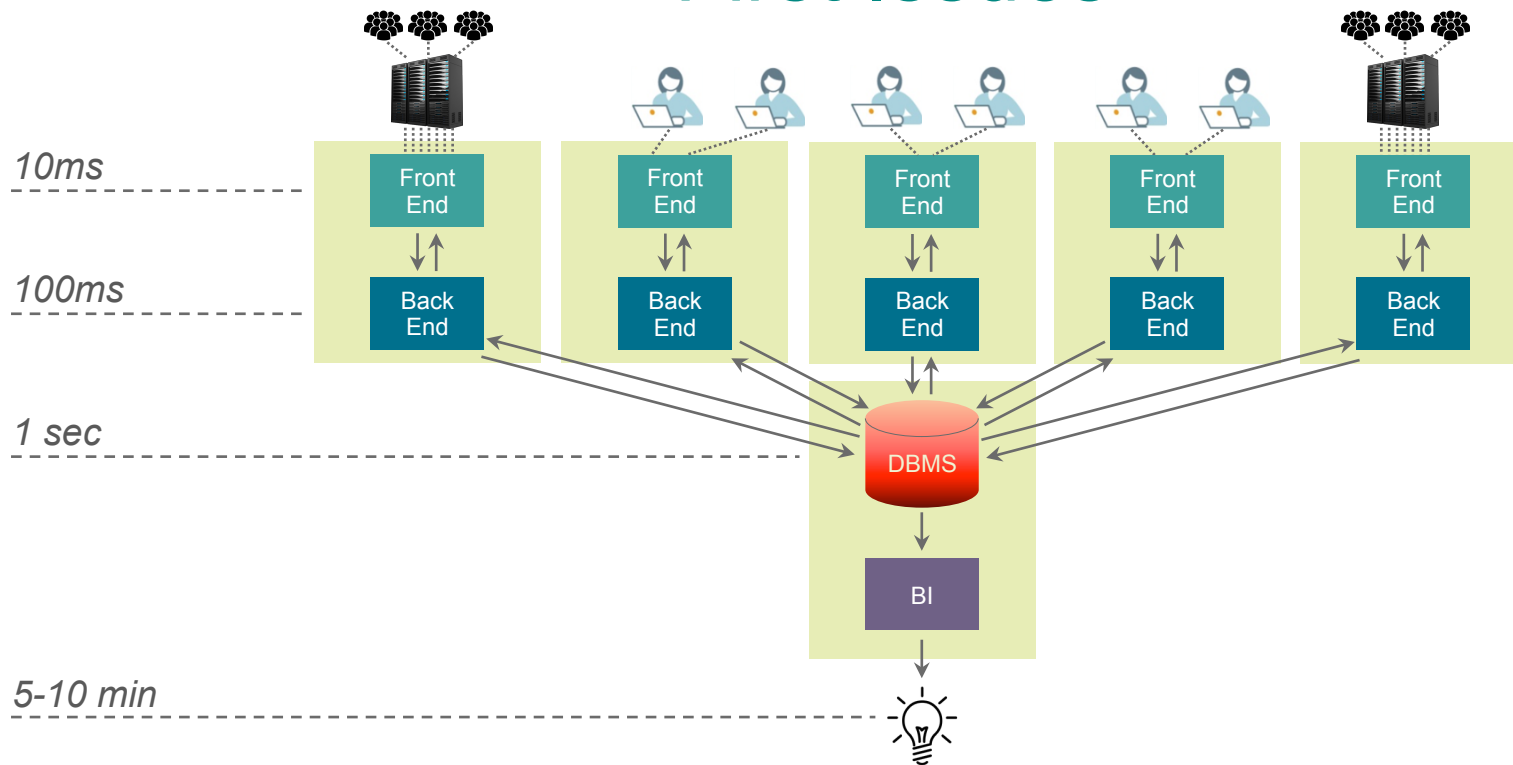
# First Issues



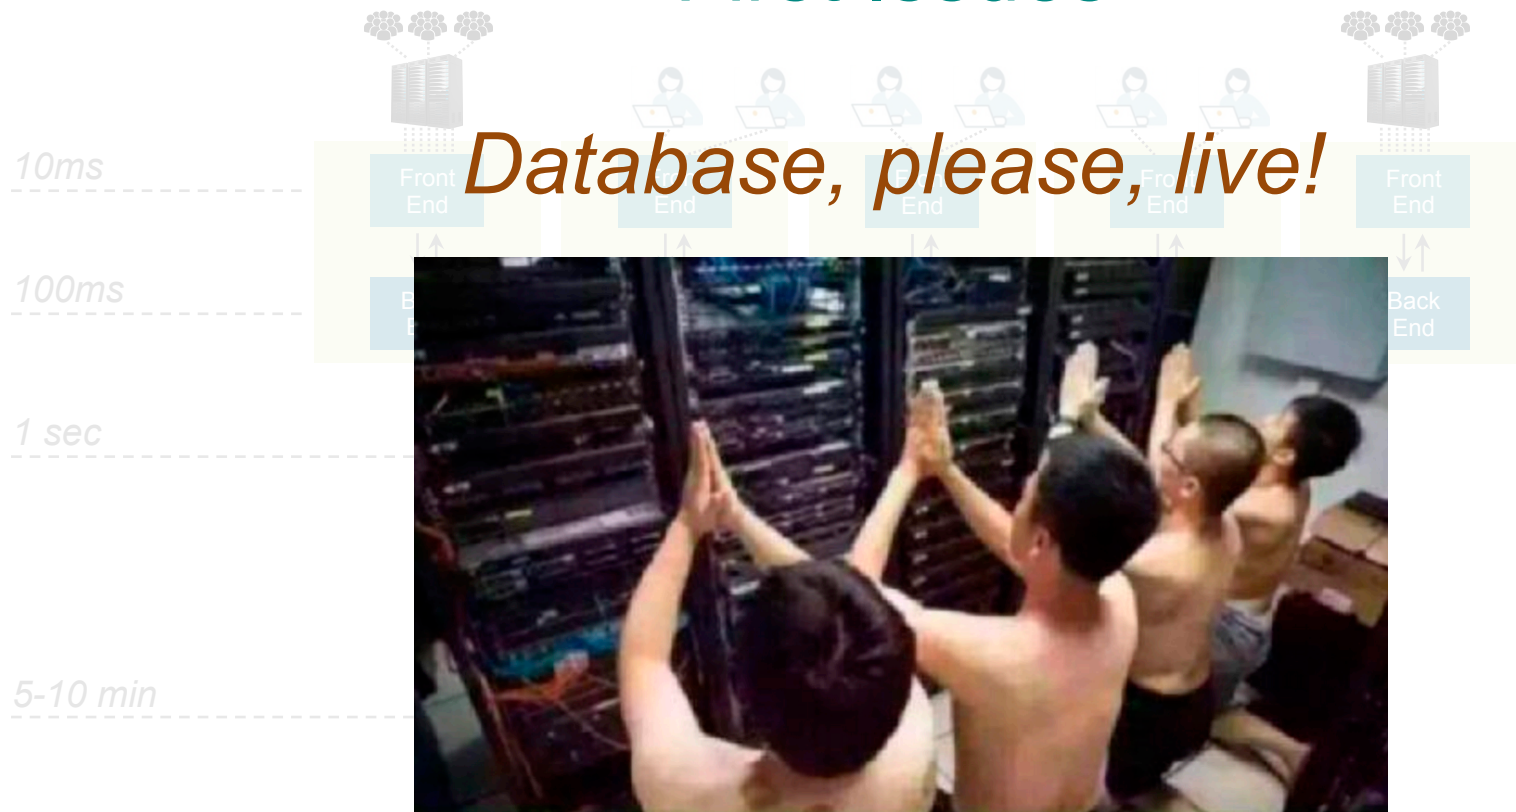
# First Issues



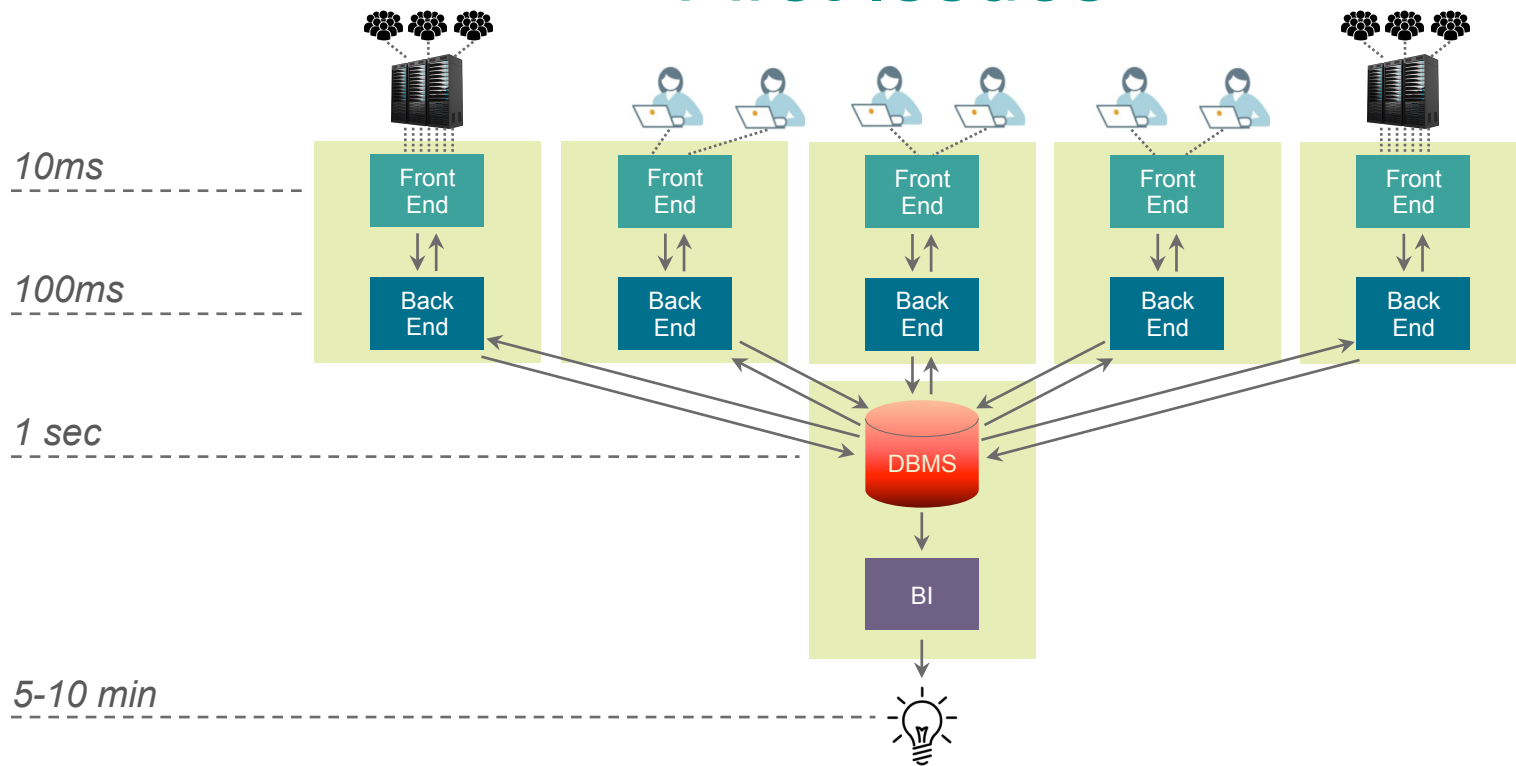
# First Issues



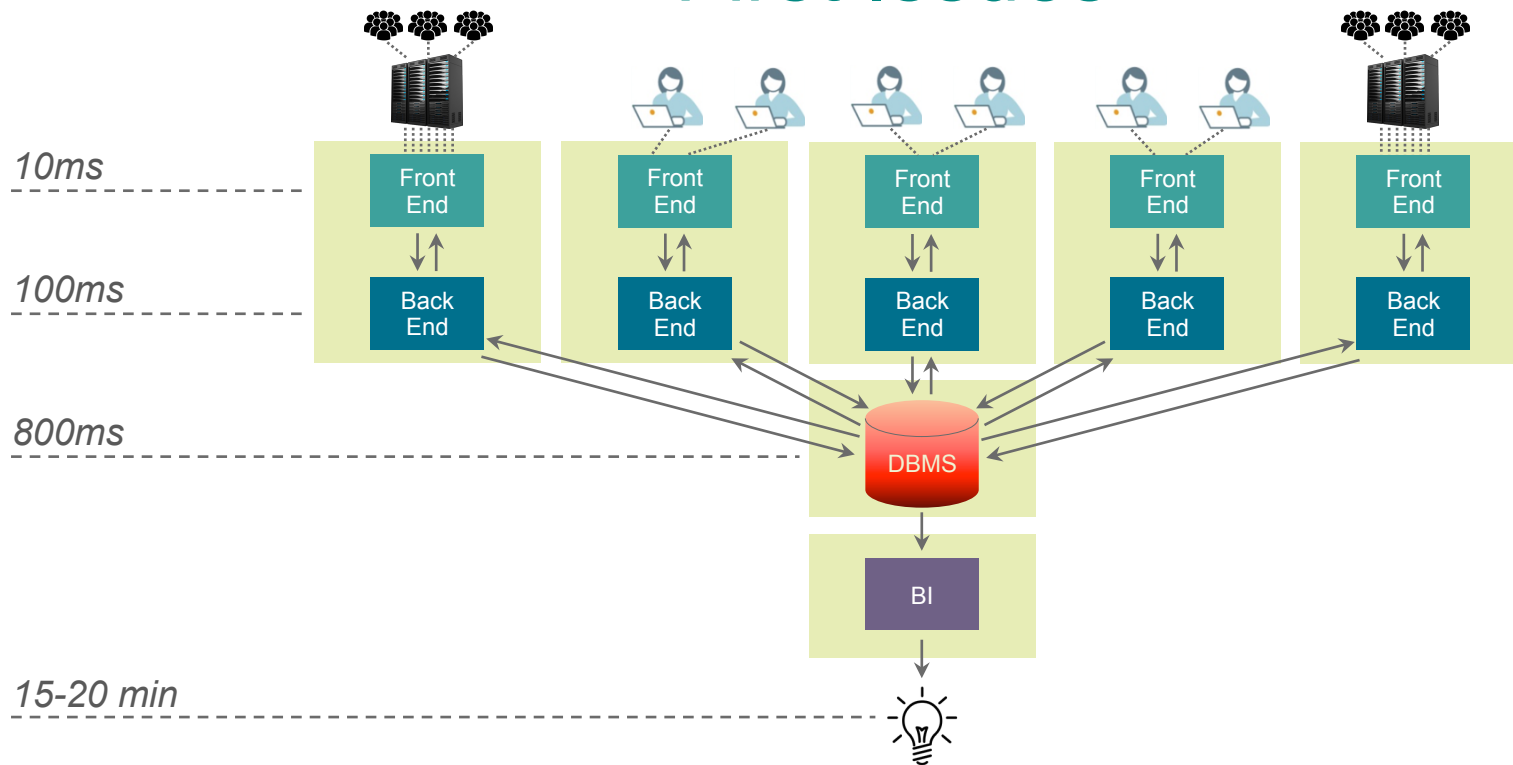
# First Issues



# First Issues



# First Issues

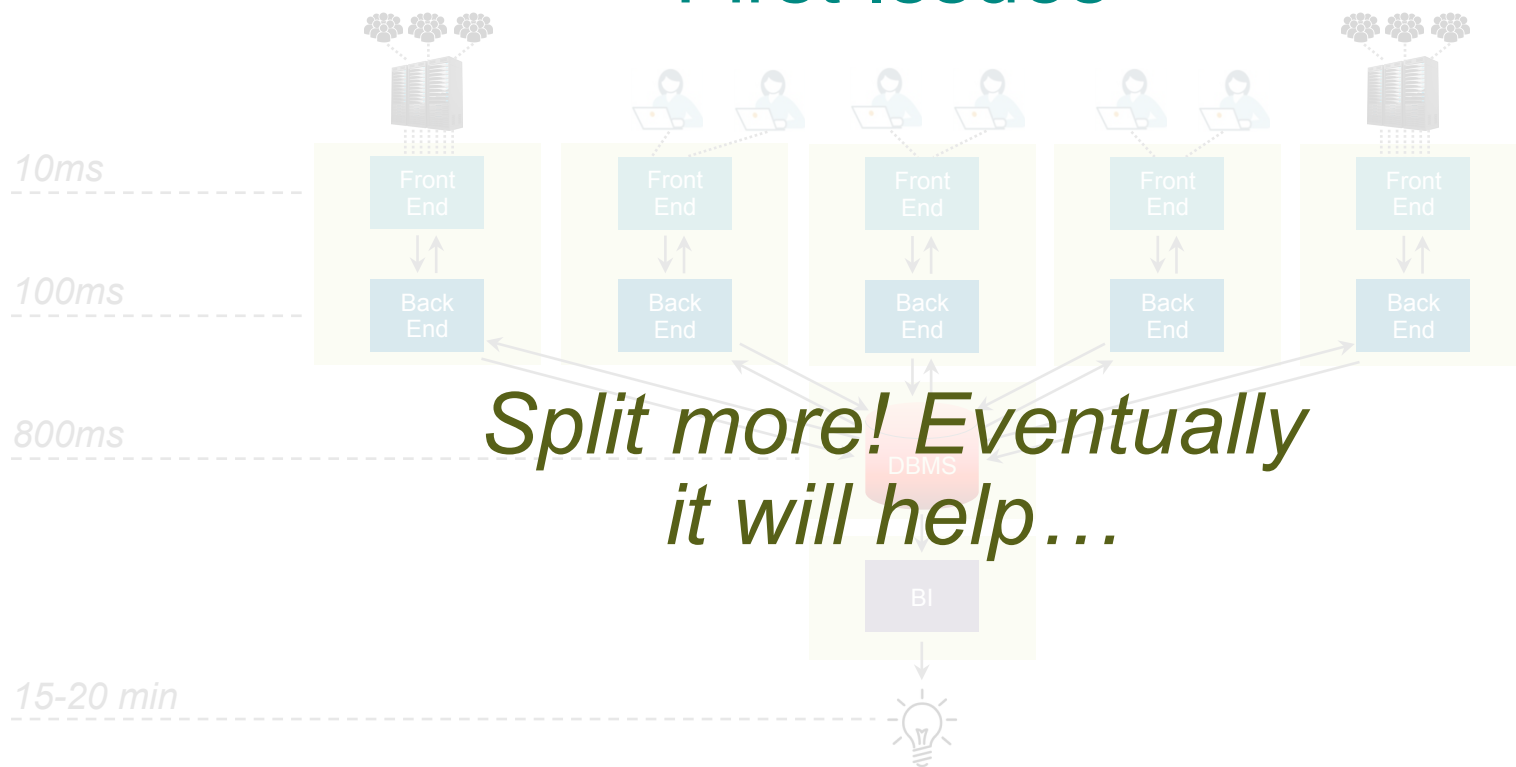


# First Issues

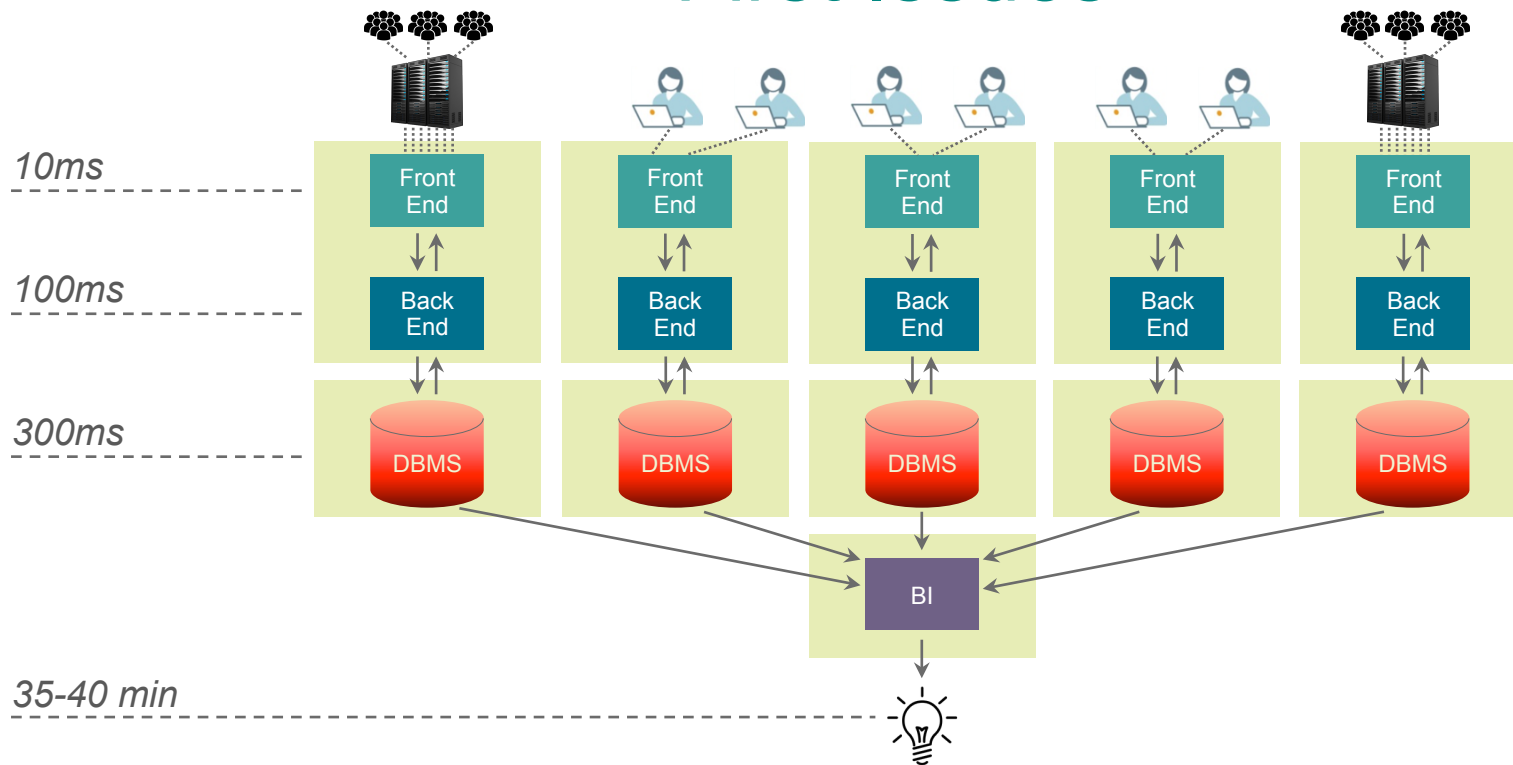




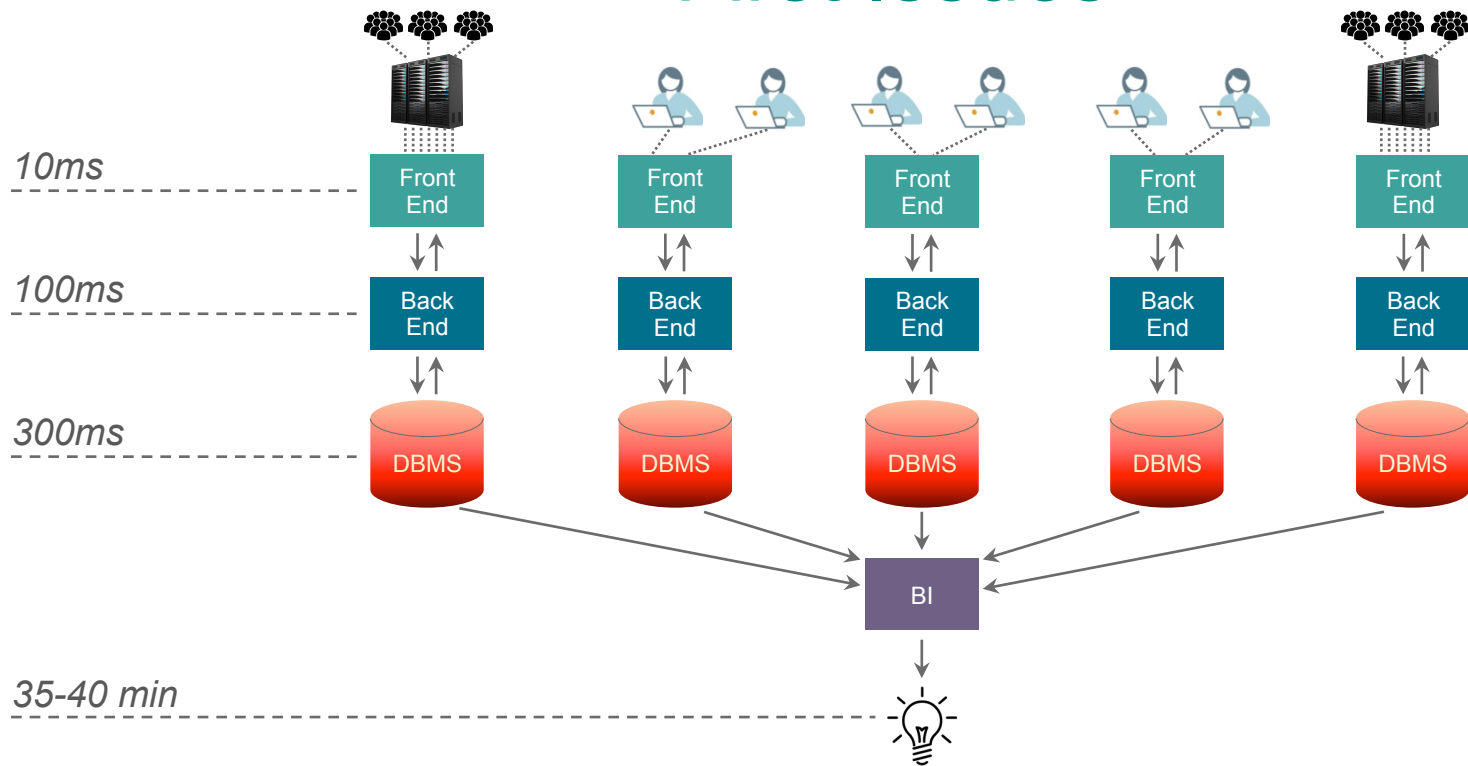
# First Issues



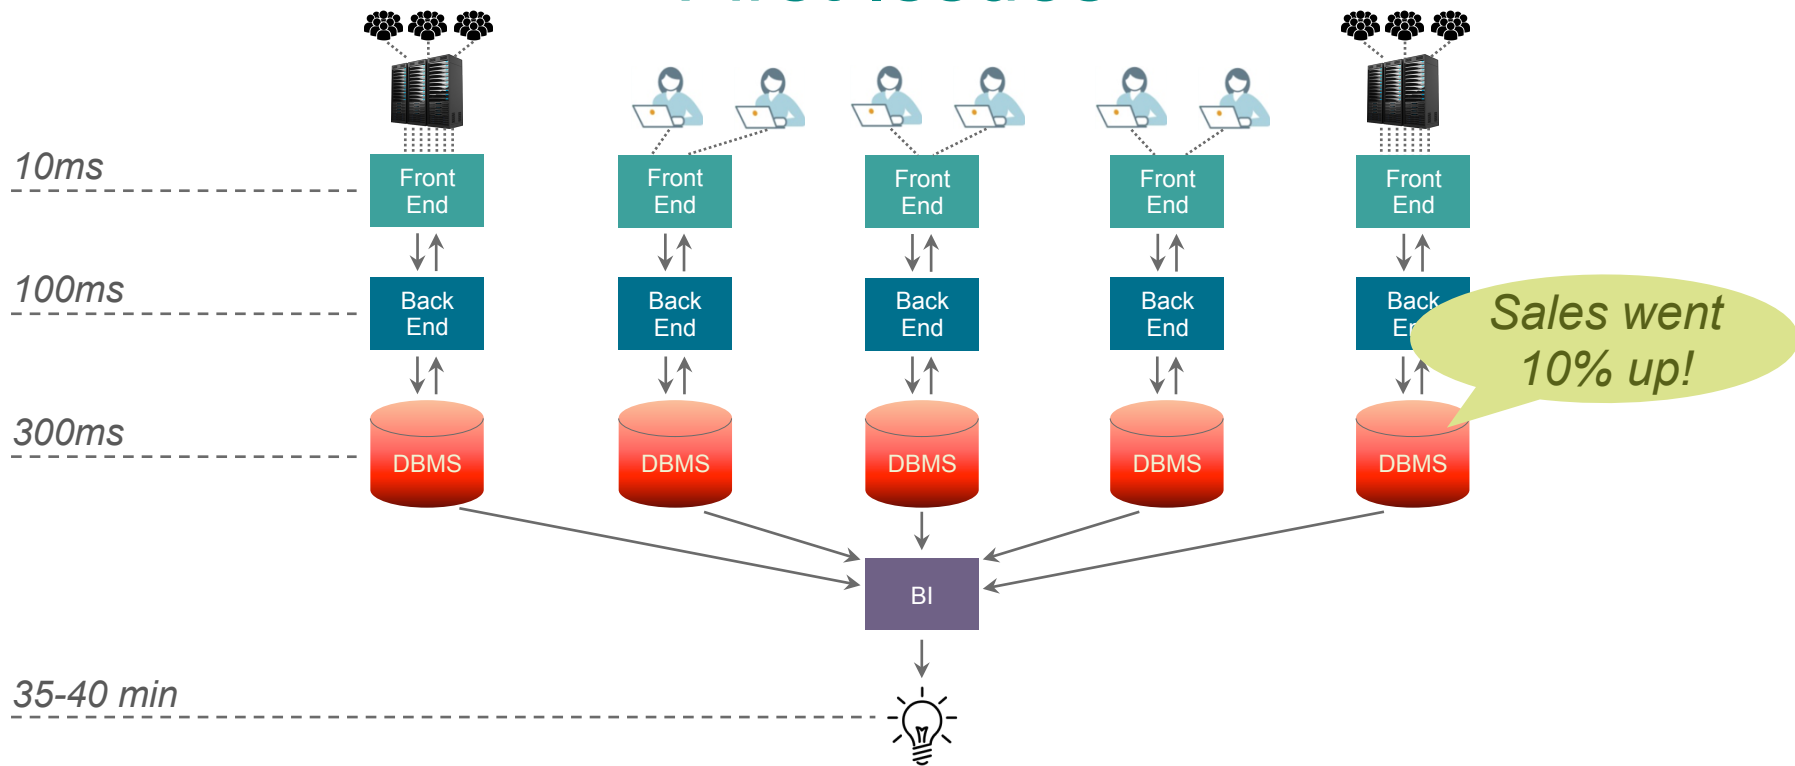
# First Issues



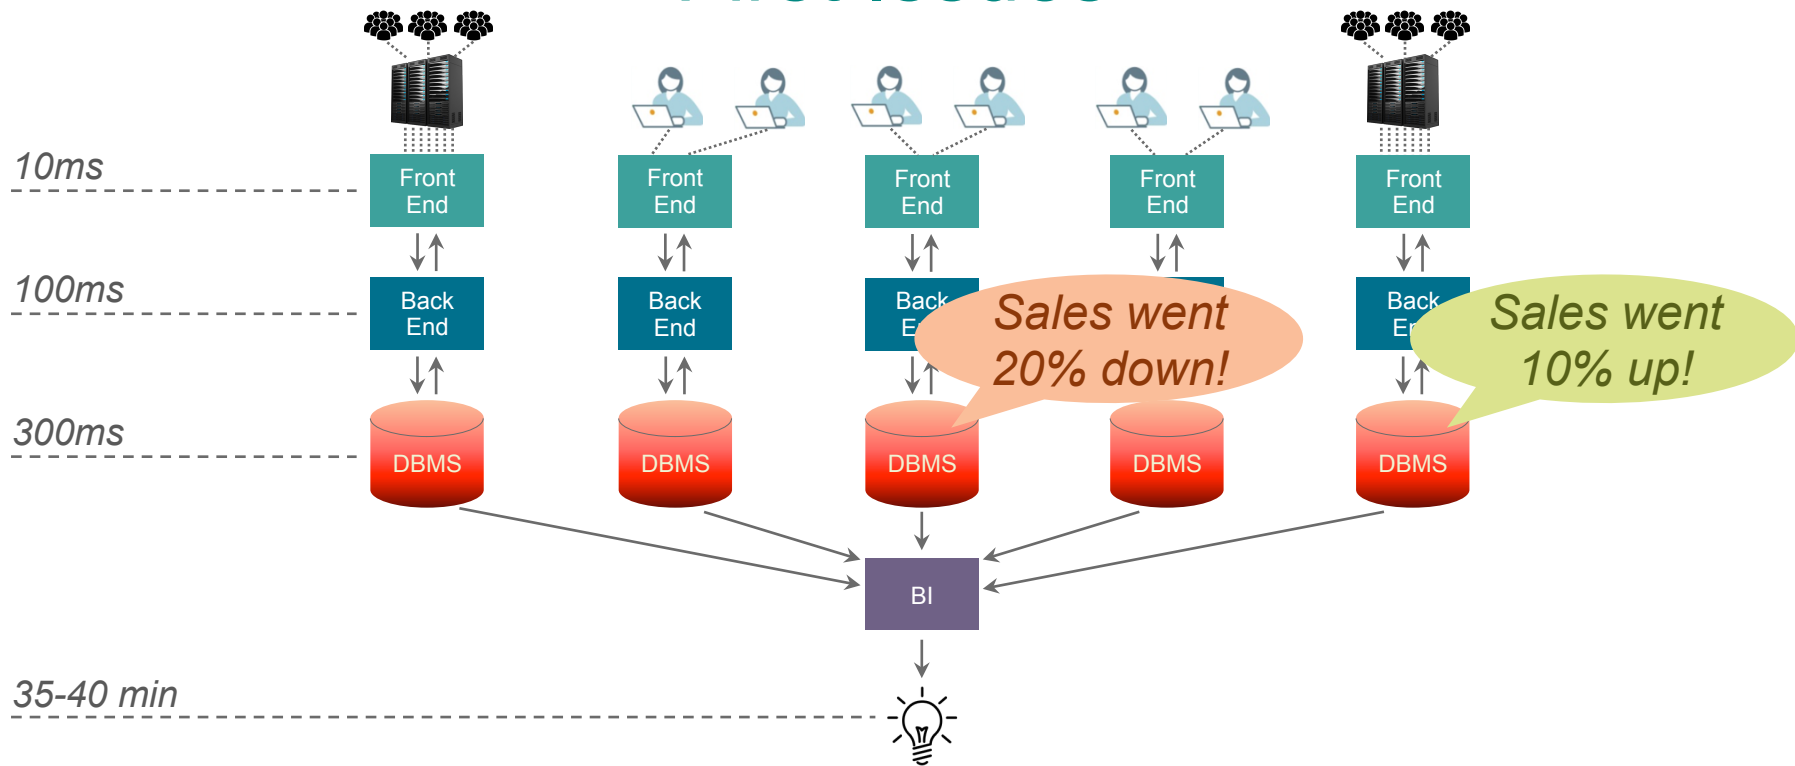
# First Issues



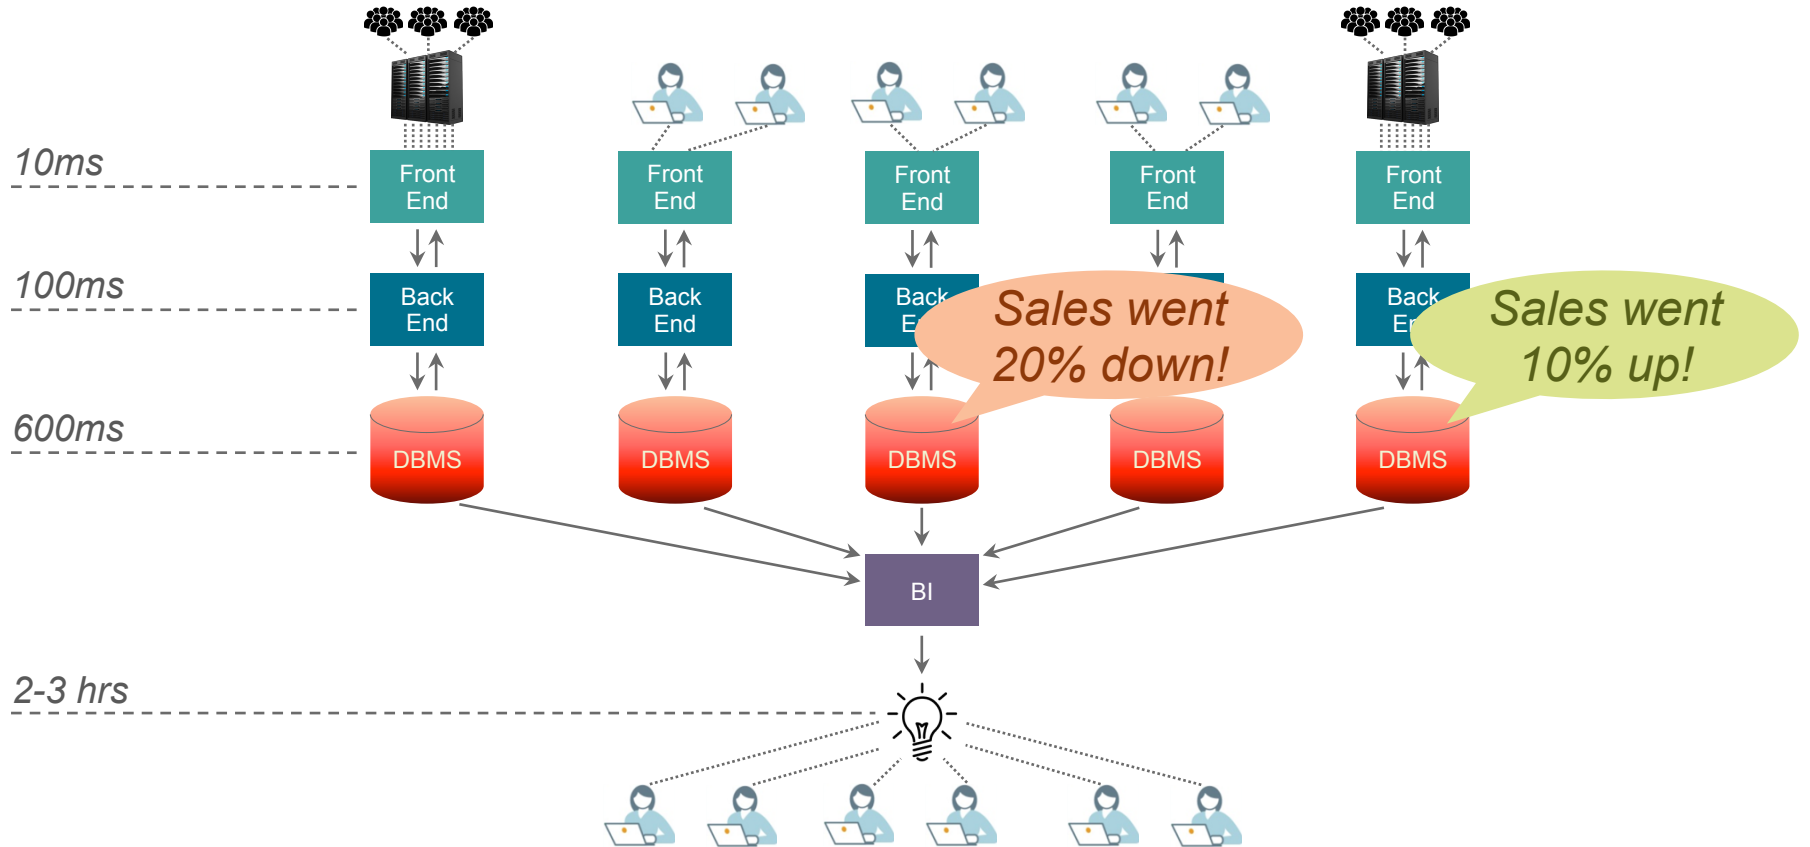
# First Issues



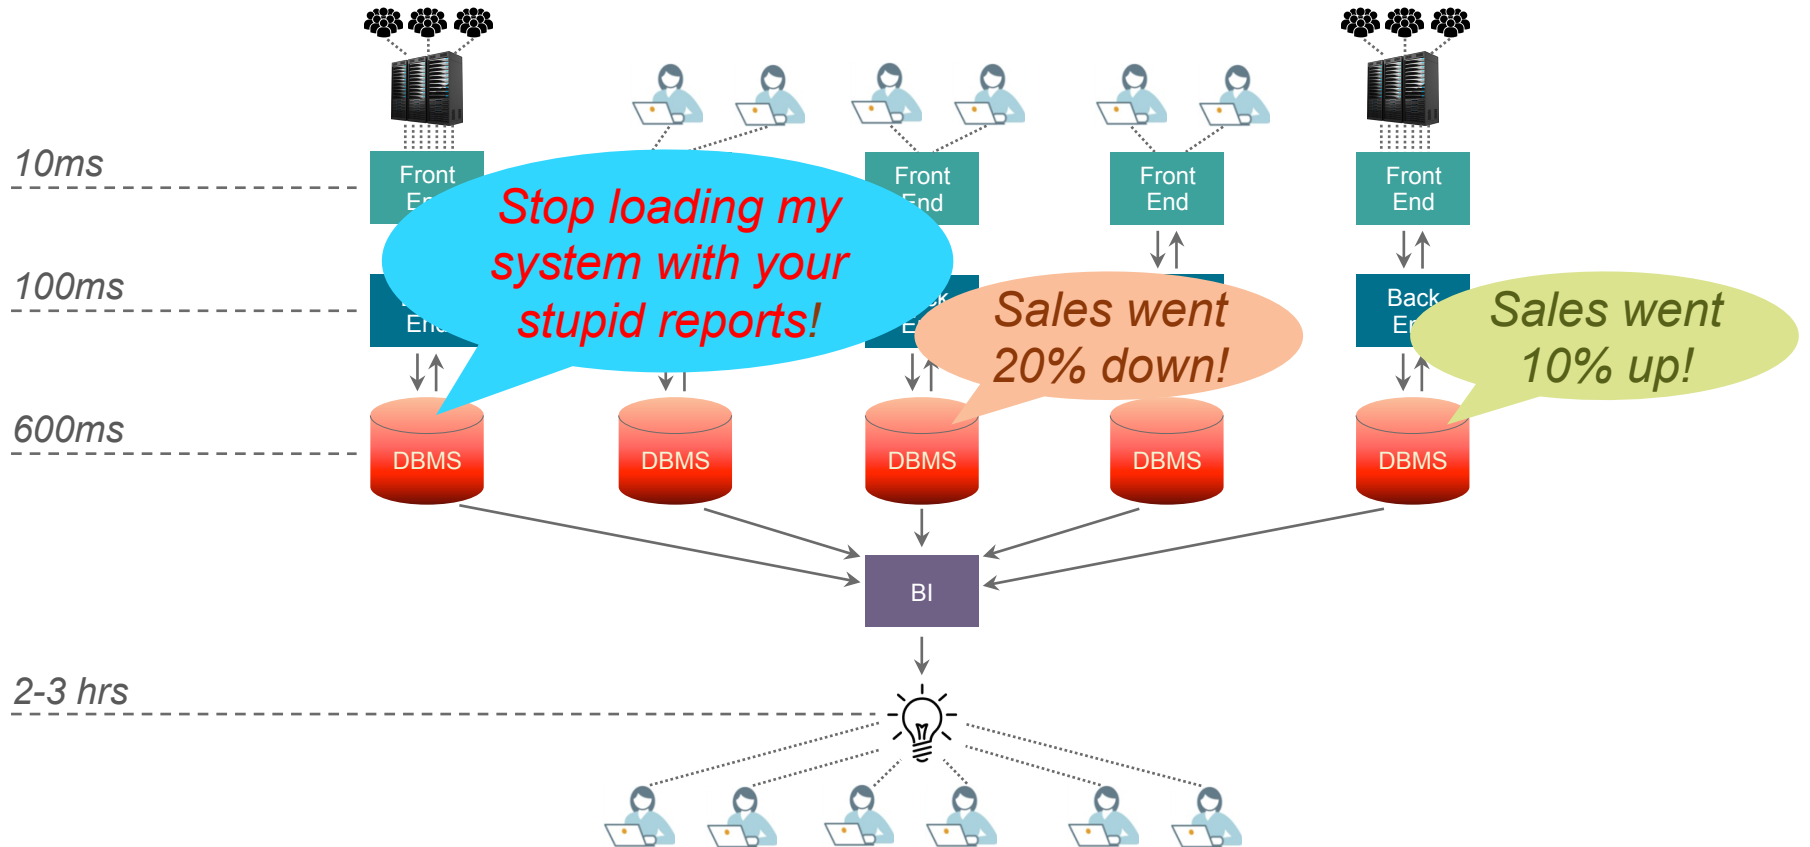
# First Issues




# First Issues



# First Issues

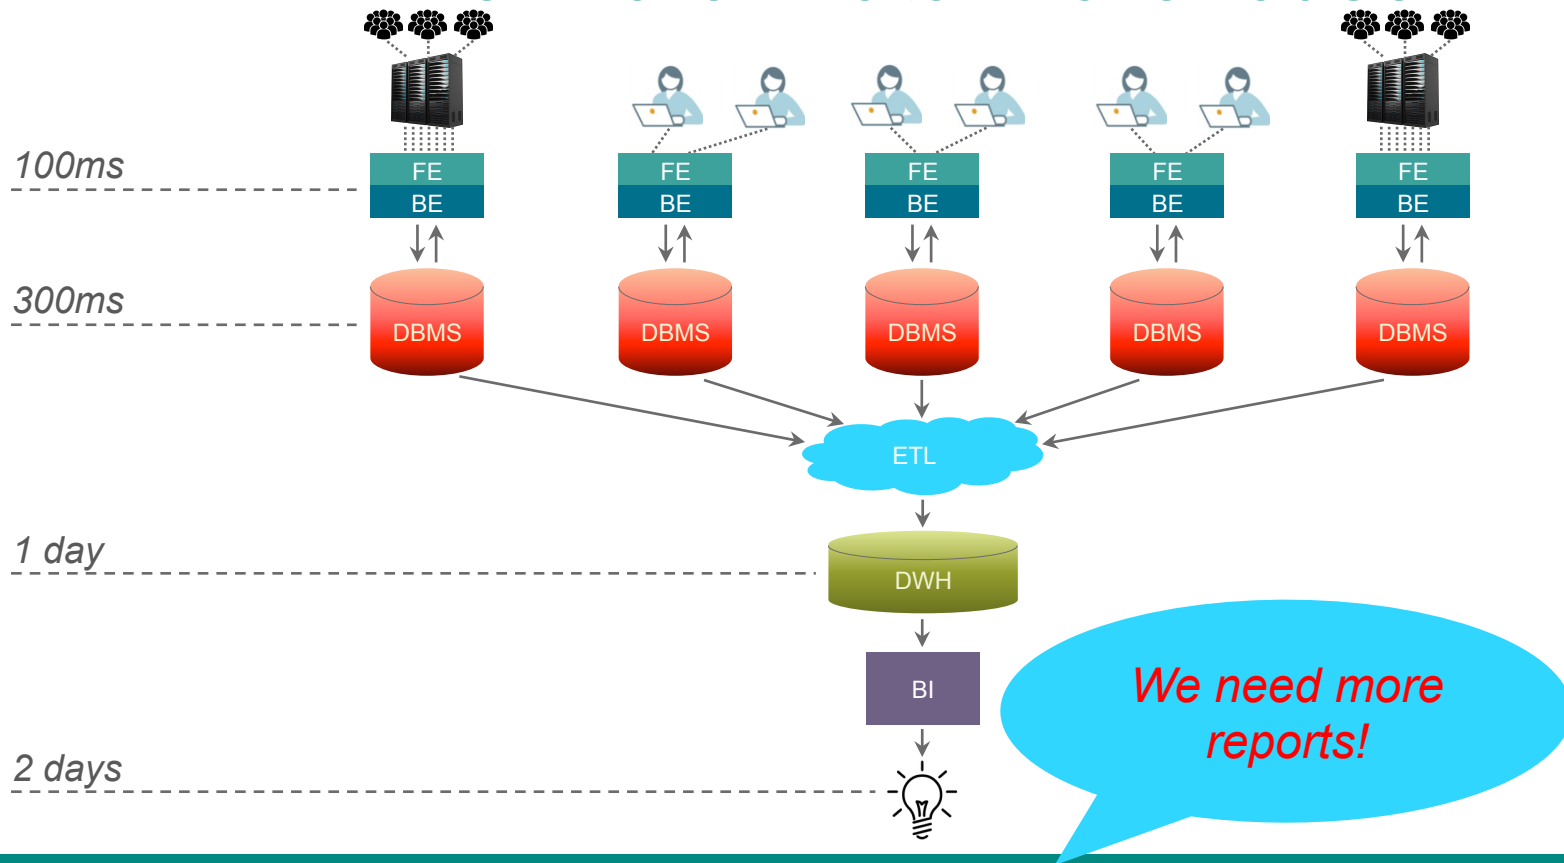


# The Era of Data Warehouse

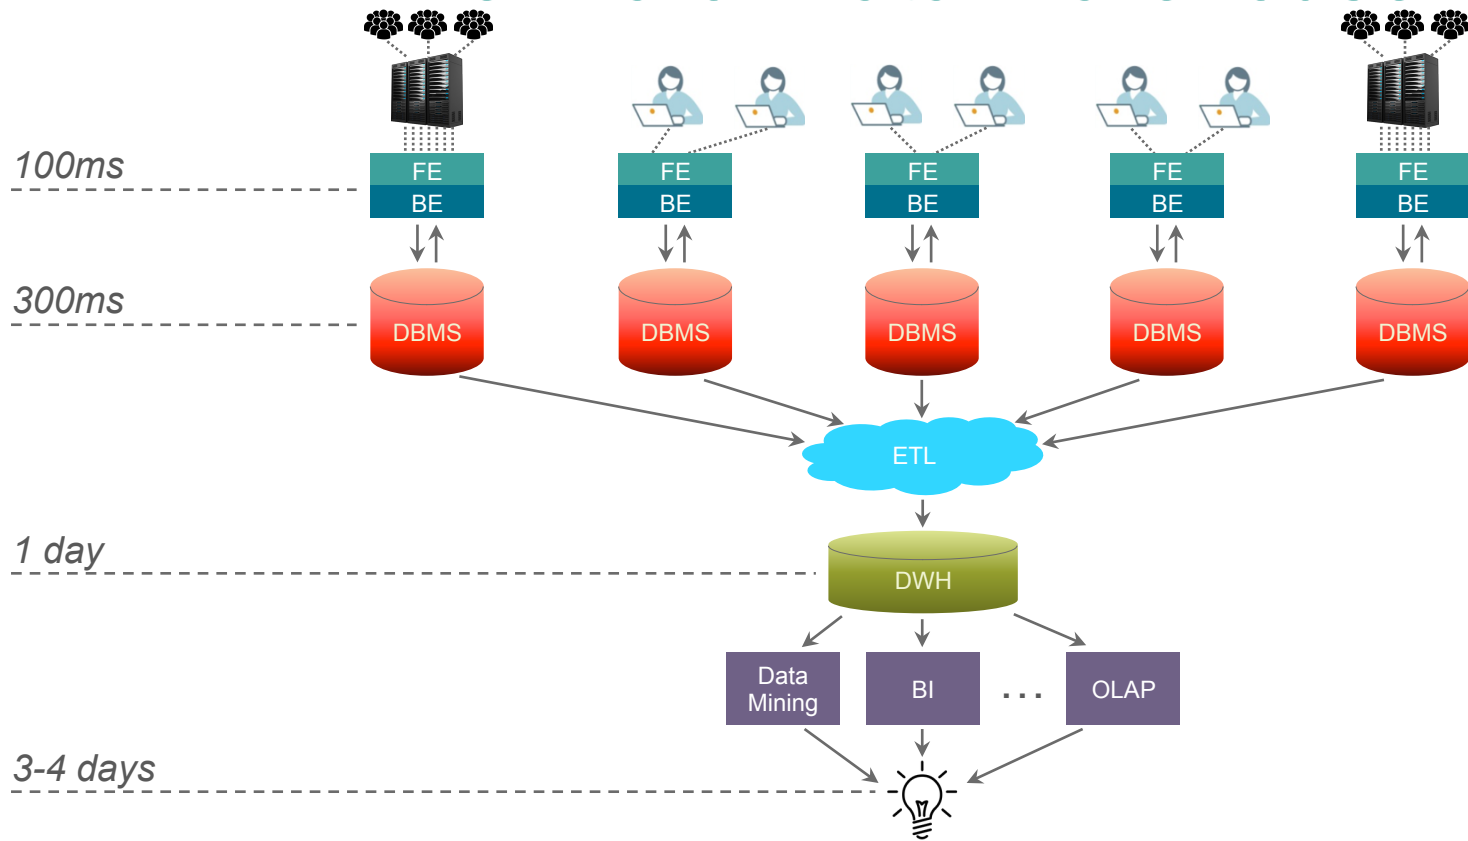




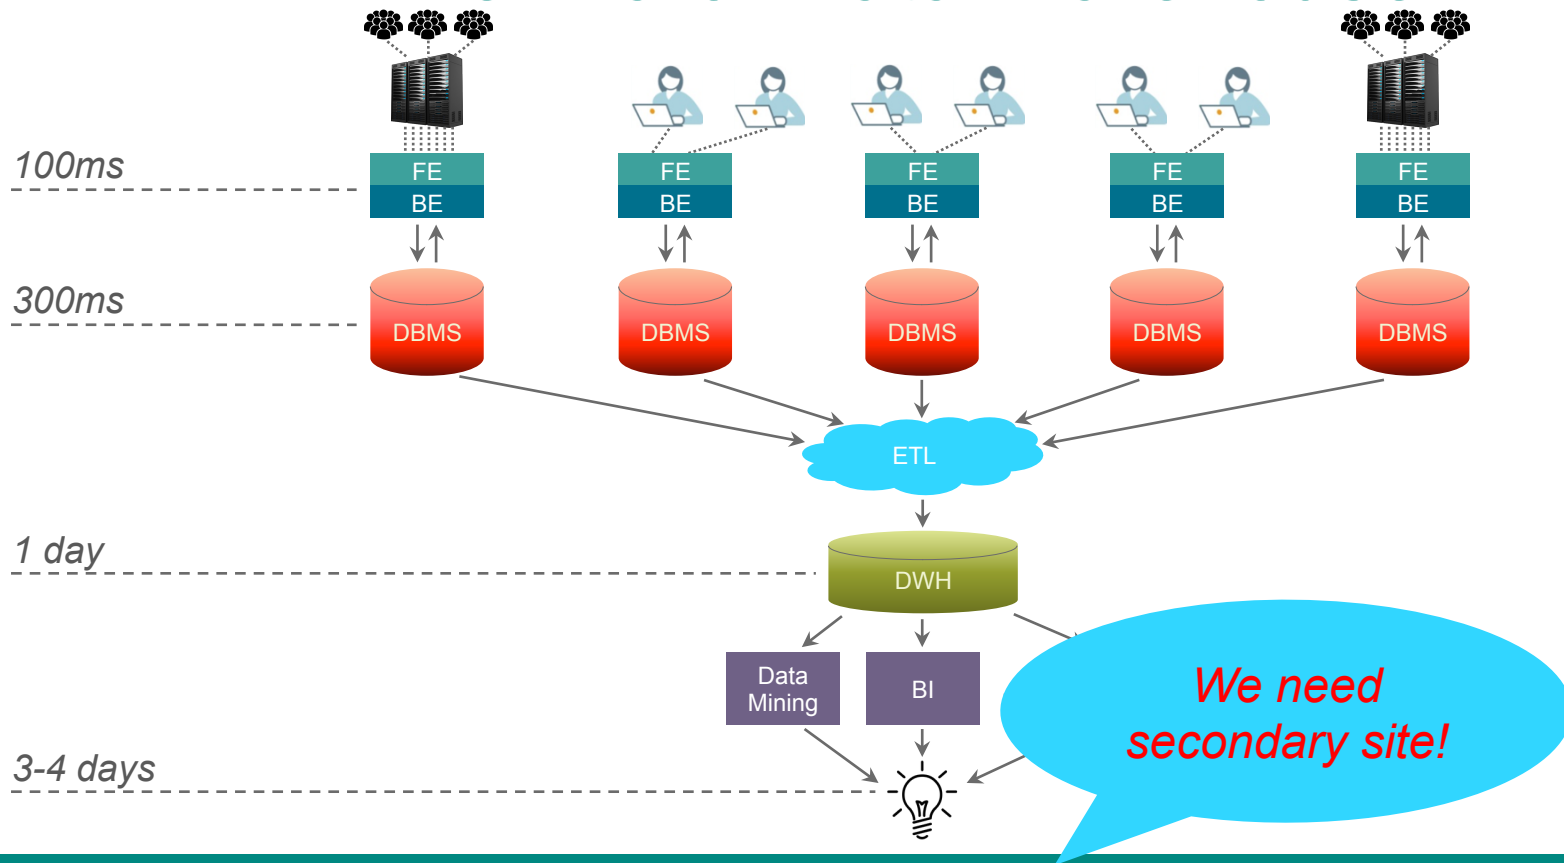
# The Era of Data Warehouse




# The Era of Data Warehouse



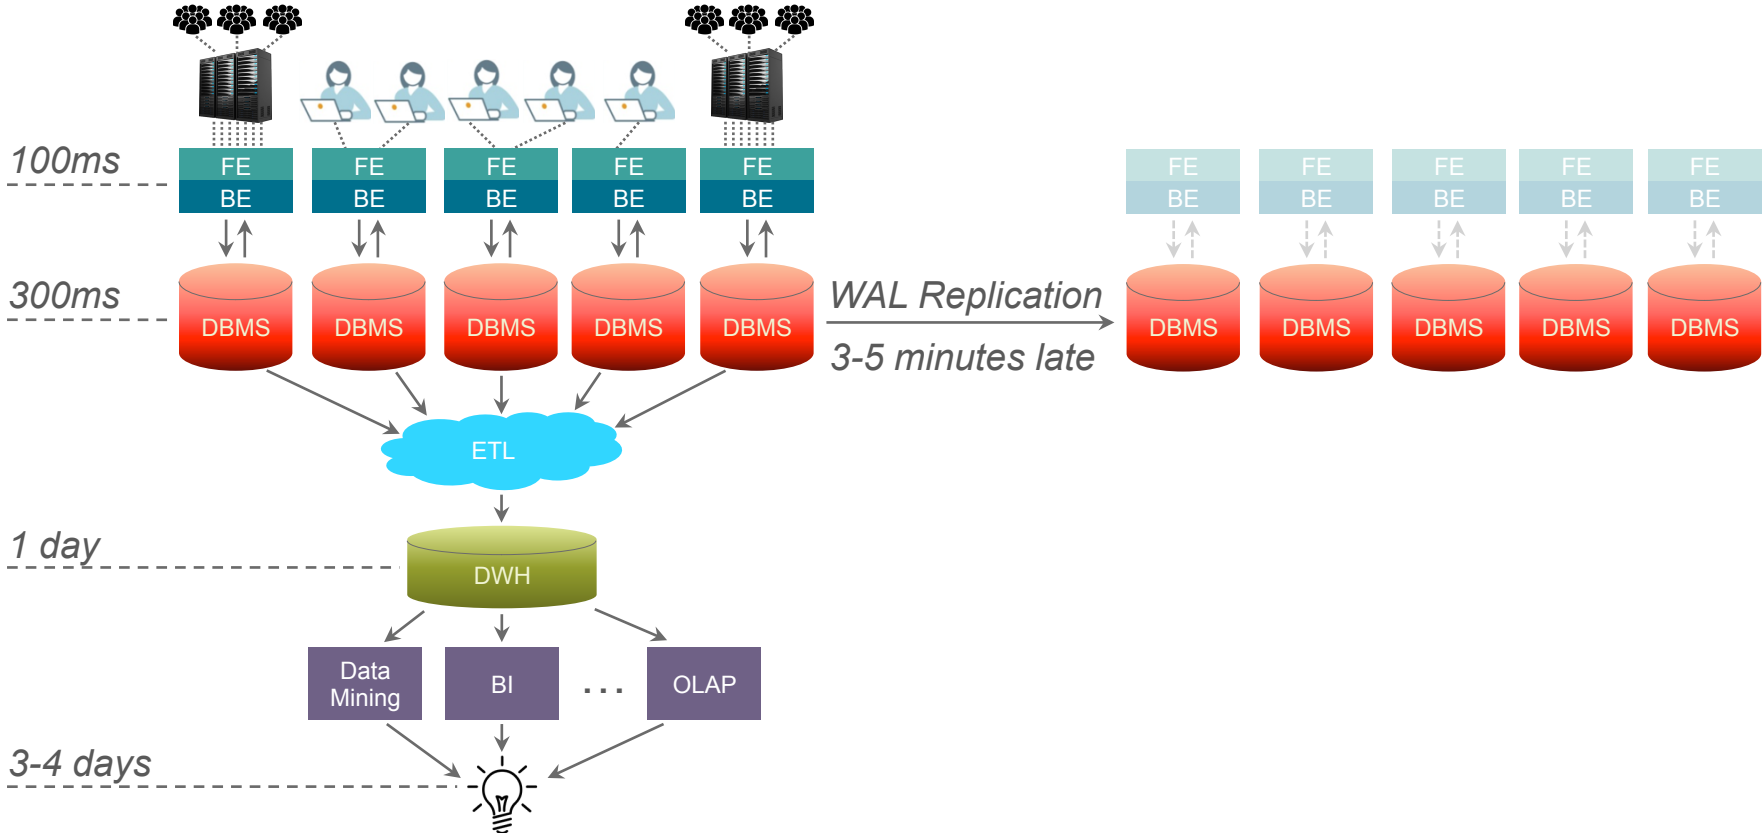
# The Era of Data Warehouse



# The Era of Data Warehouse



# The Era of Data Warehouse



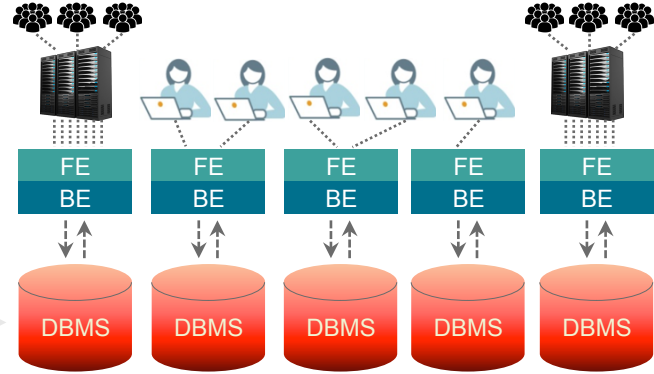


# The Era of Data Warehouse


100ms



OL Replication  
5 minutes late

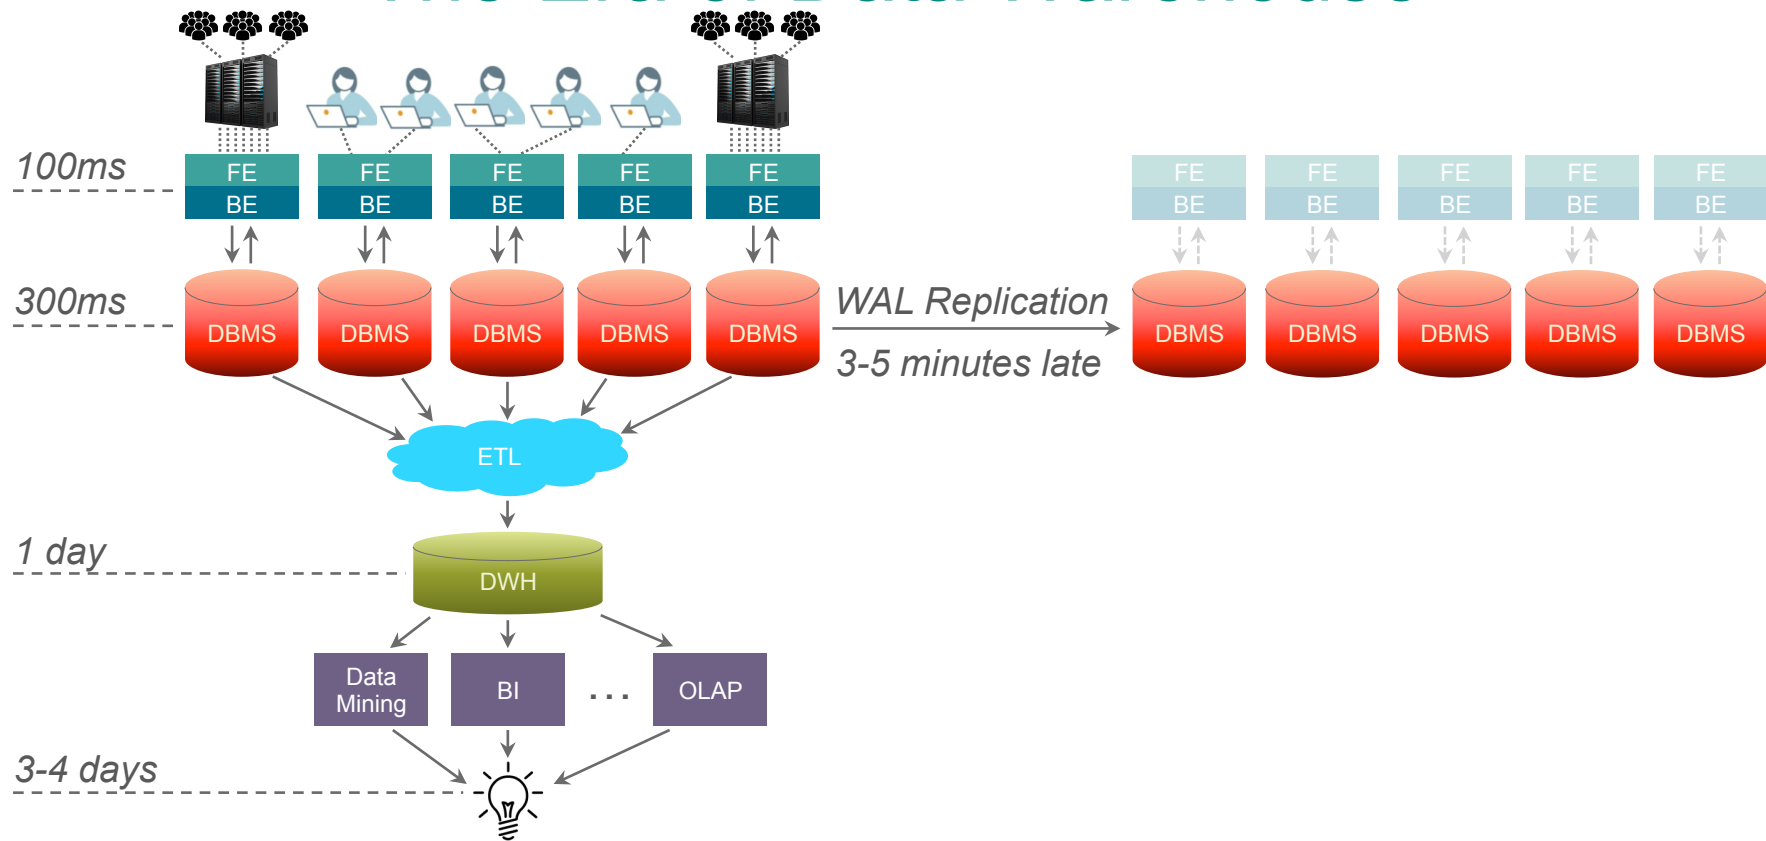


3-4 days



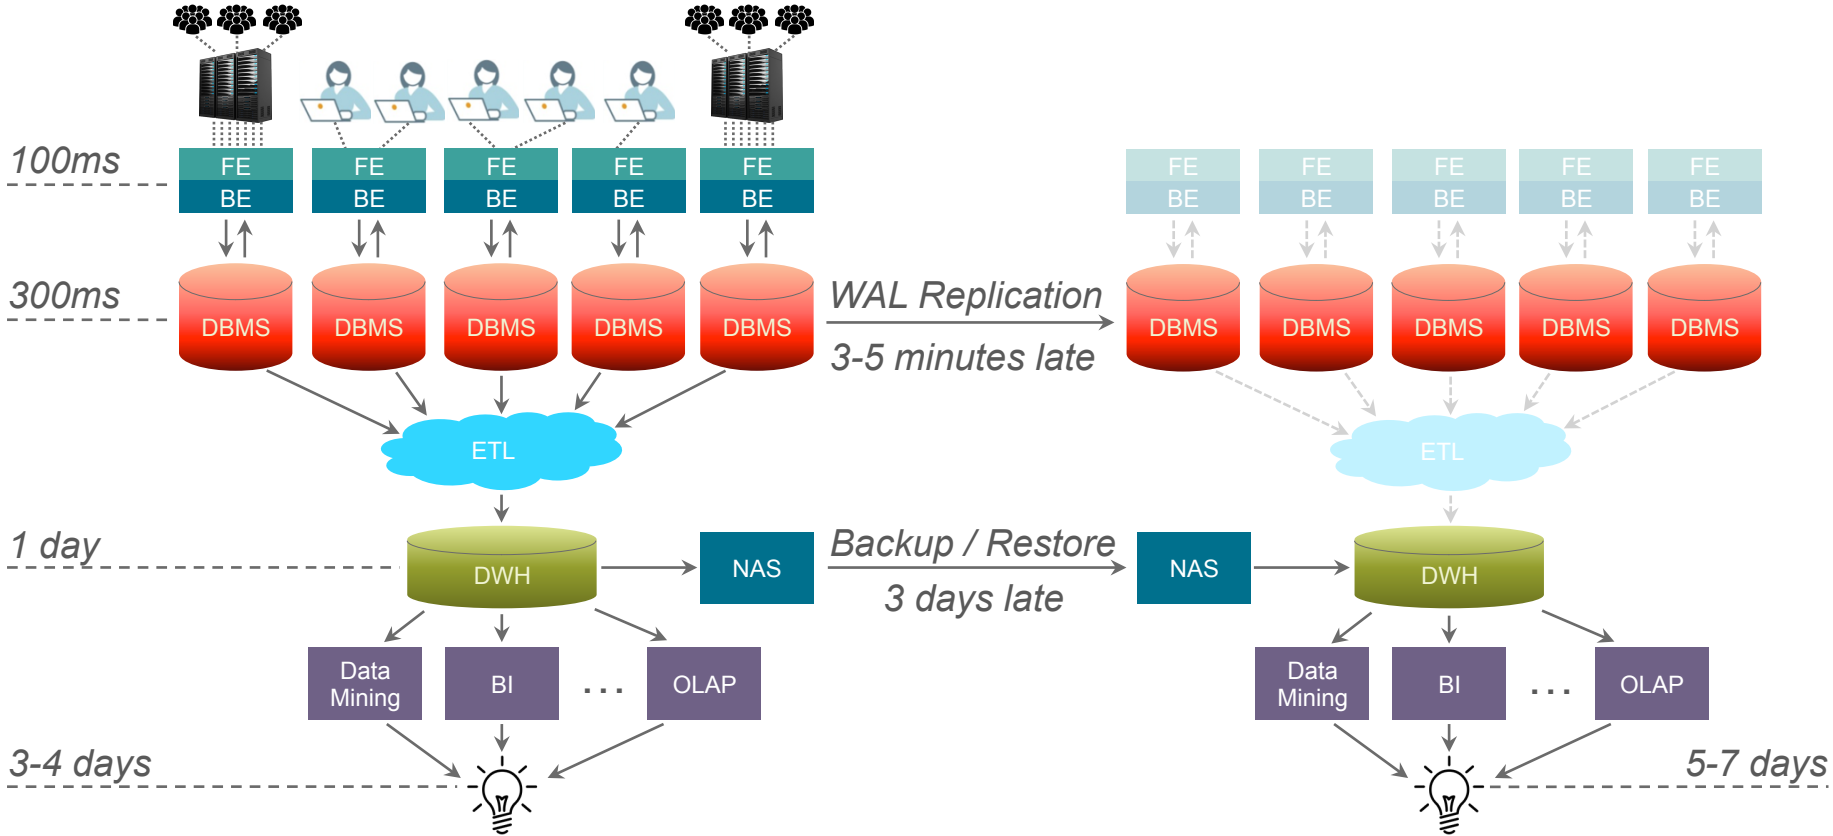
*Where is our DWH? We need this data now!*

# The Era of Data Warehouse

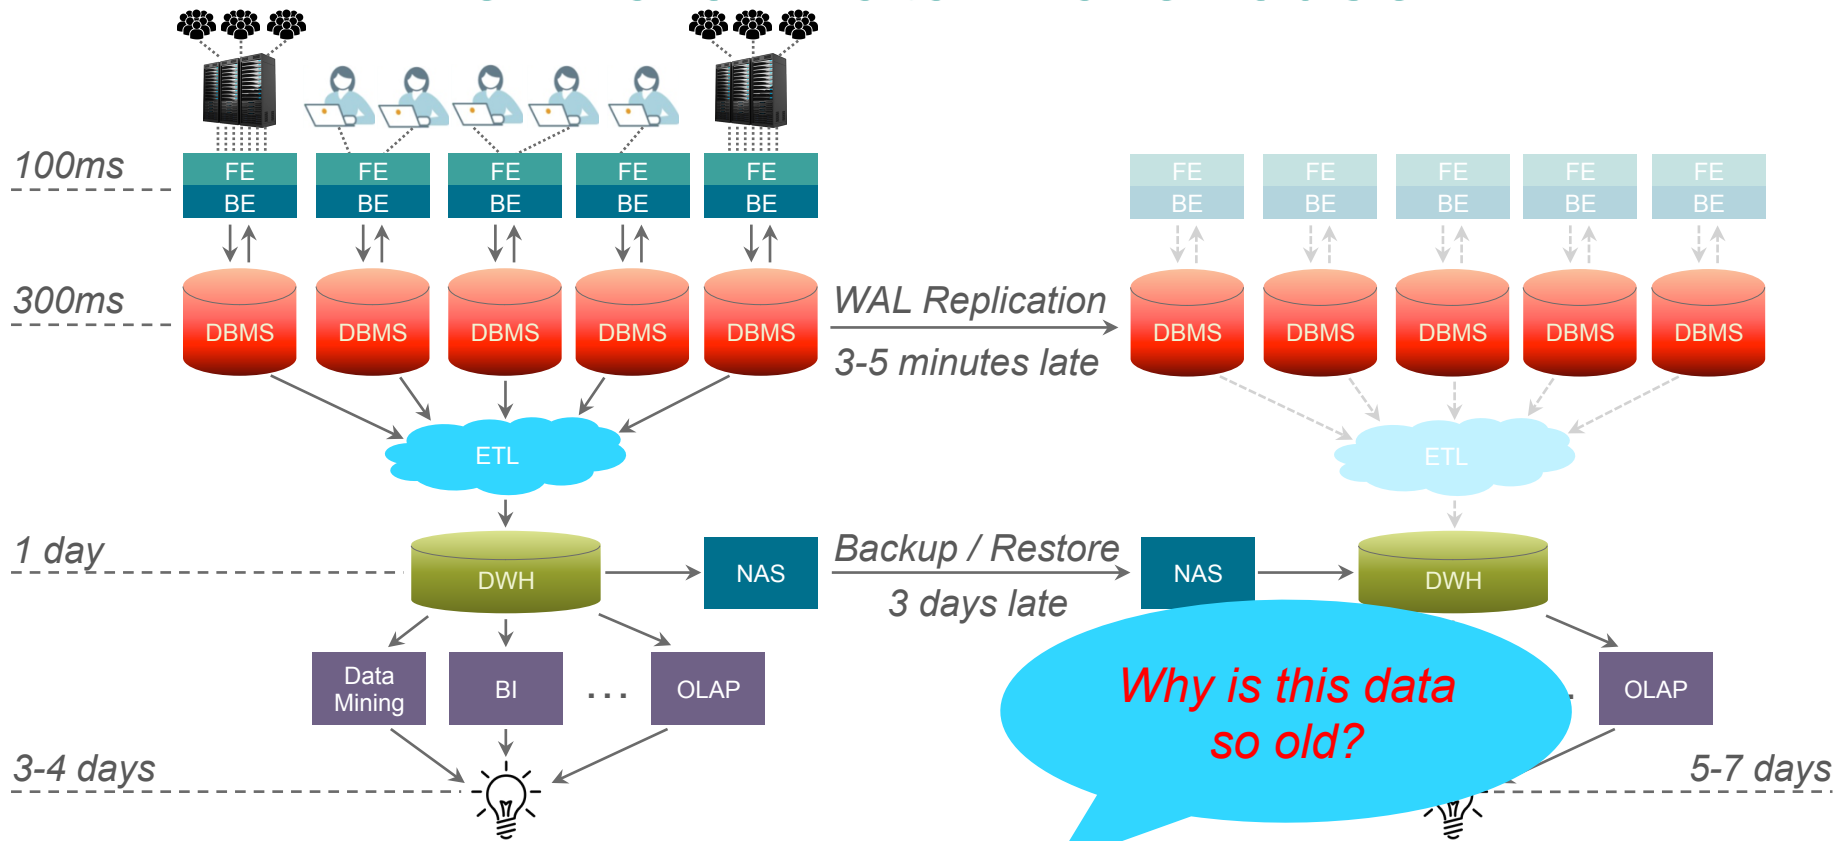





# The Era of Data Warehouse



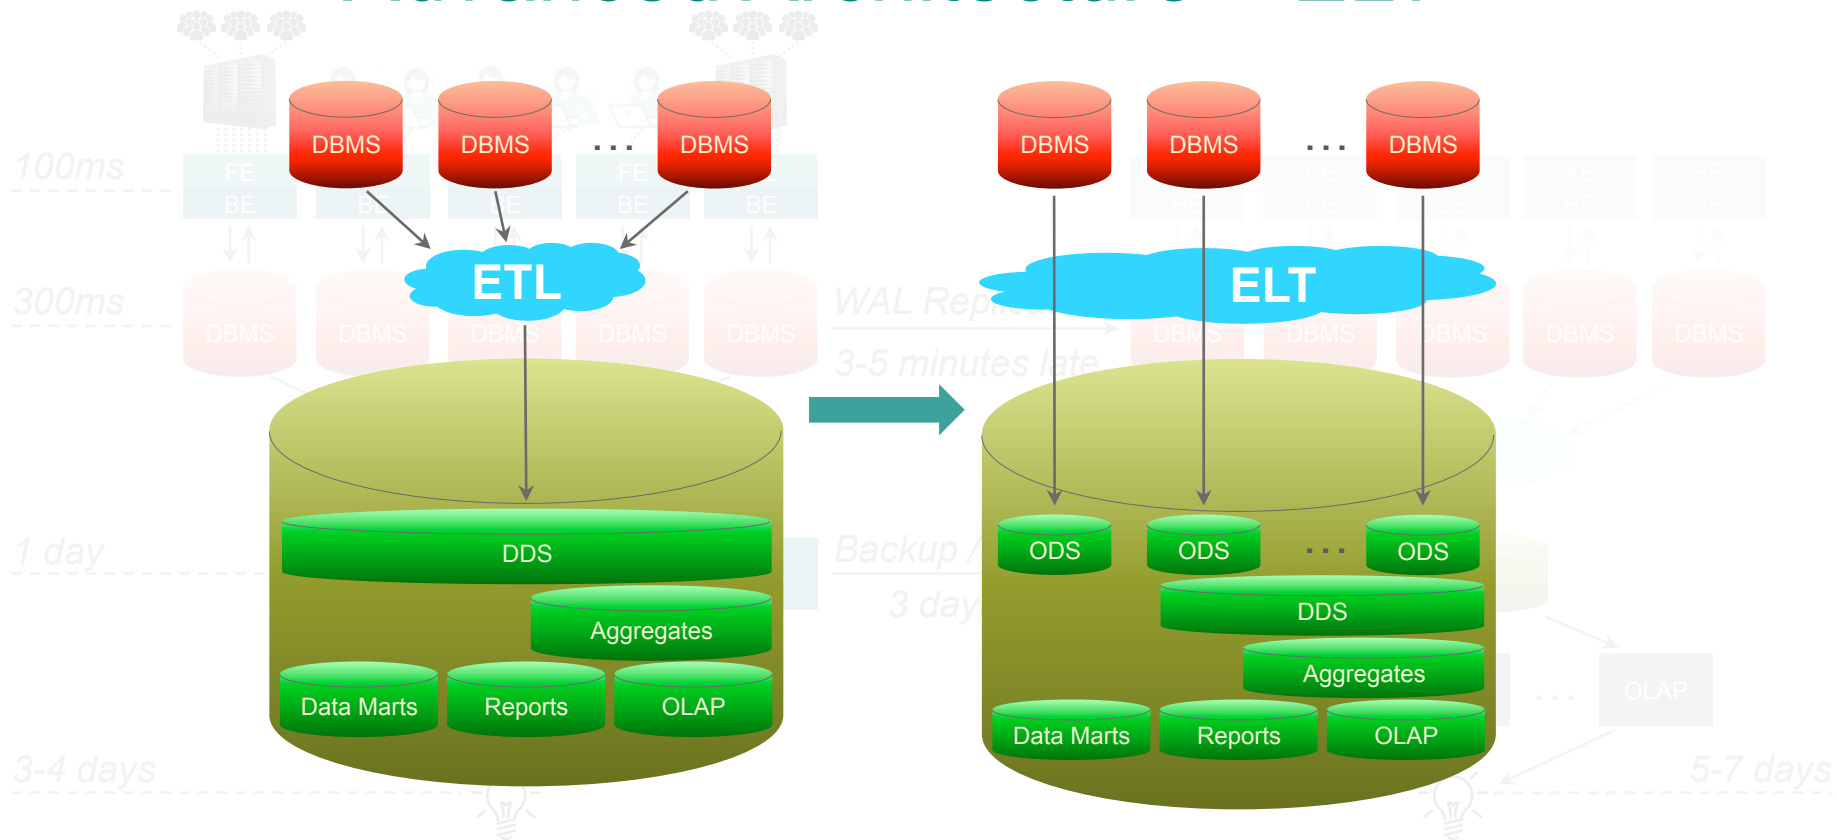
# The Era of Data Warehouse



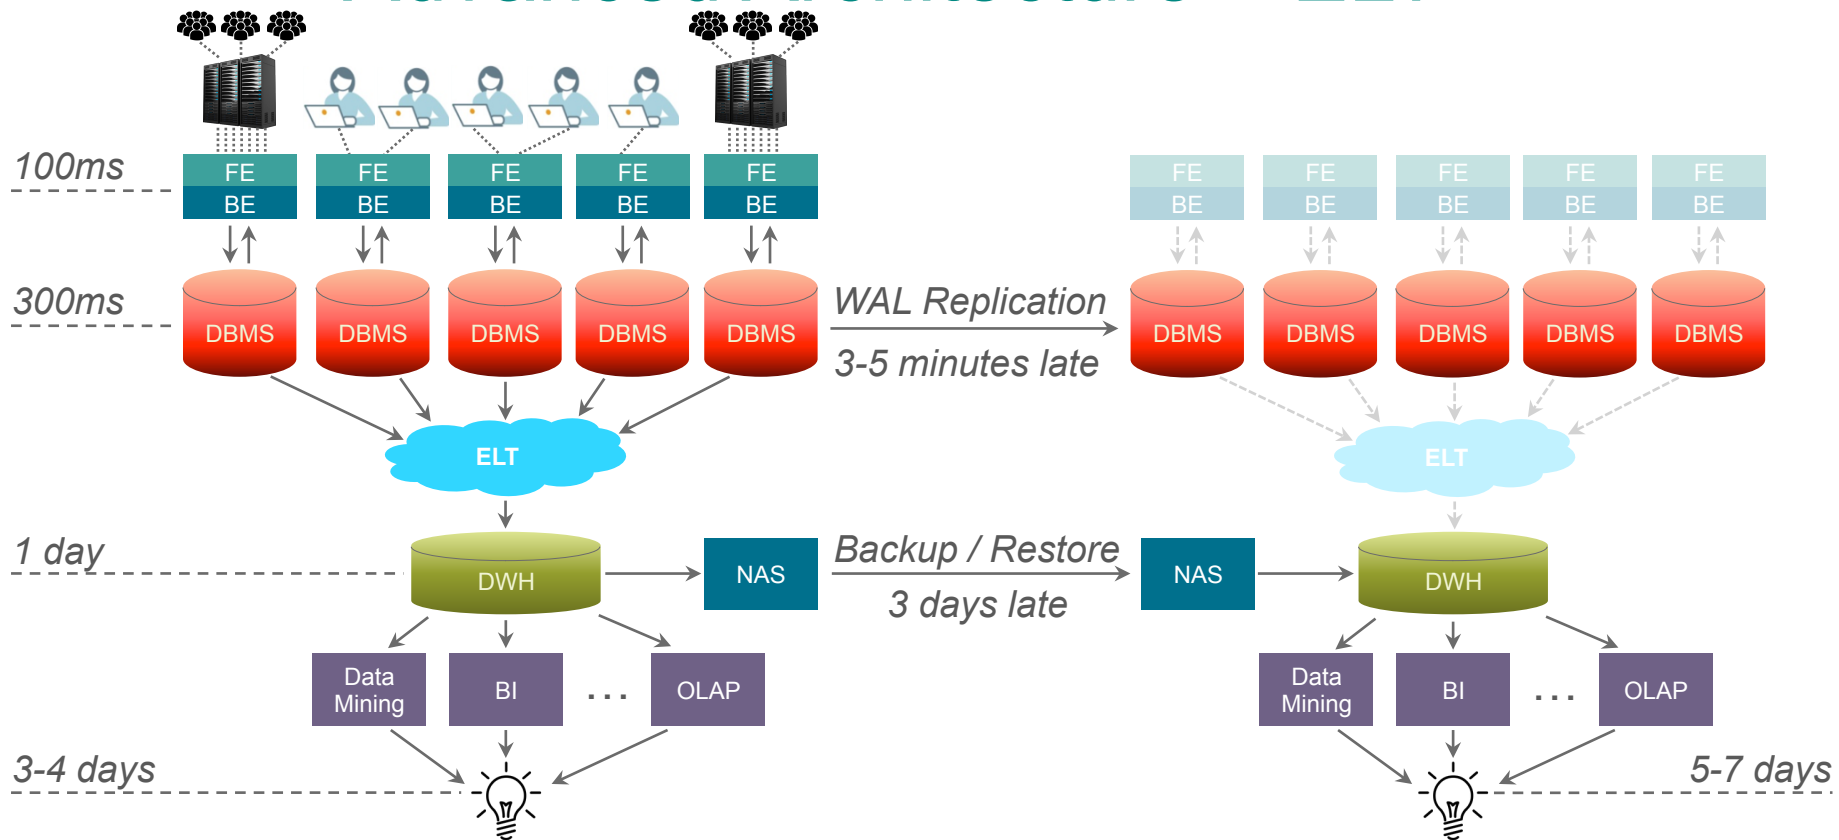
# The Era of Data Warehouse



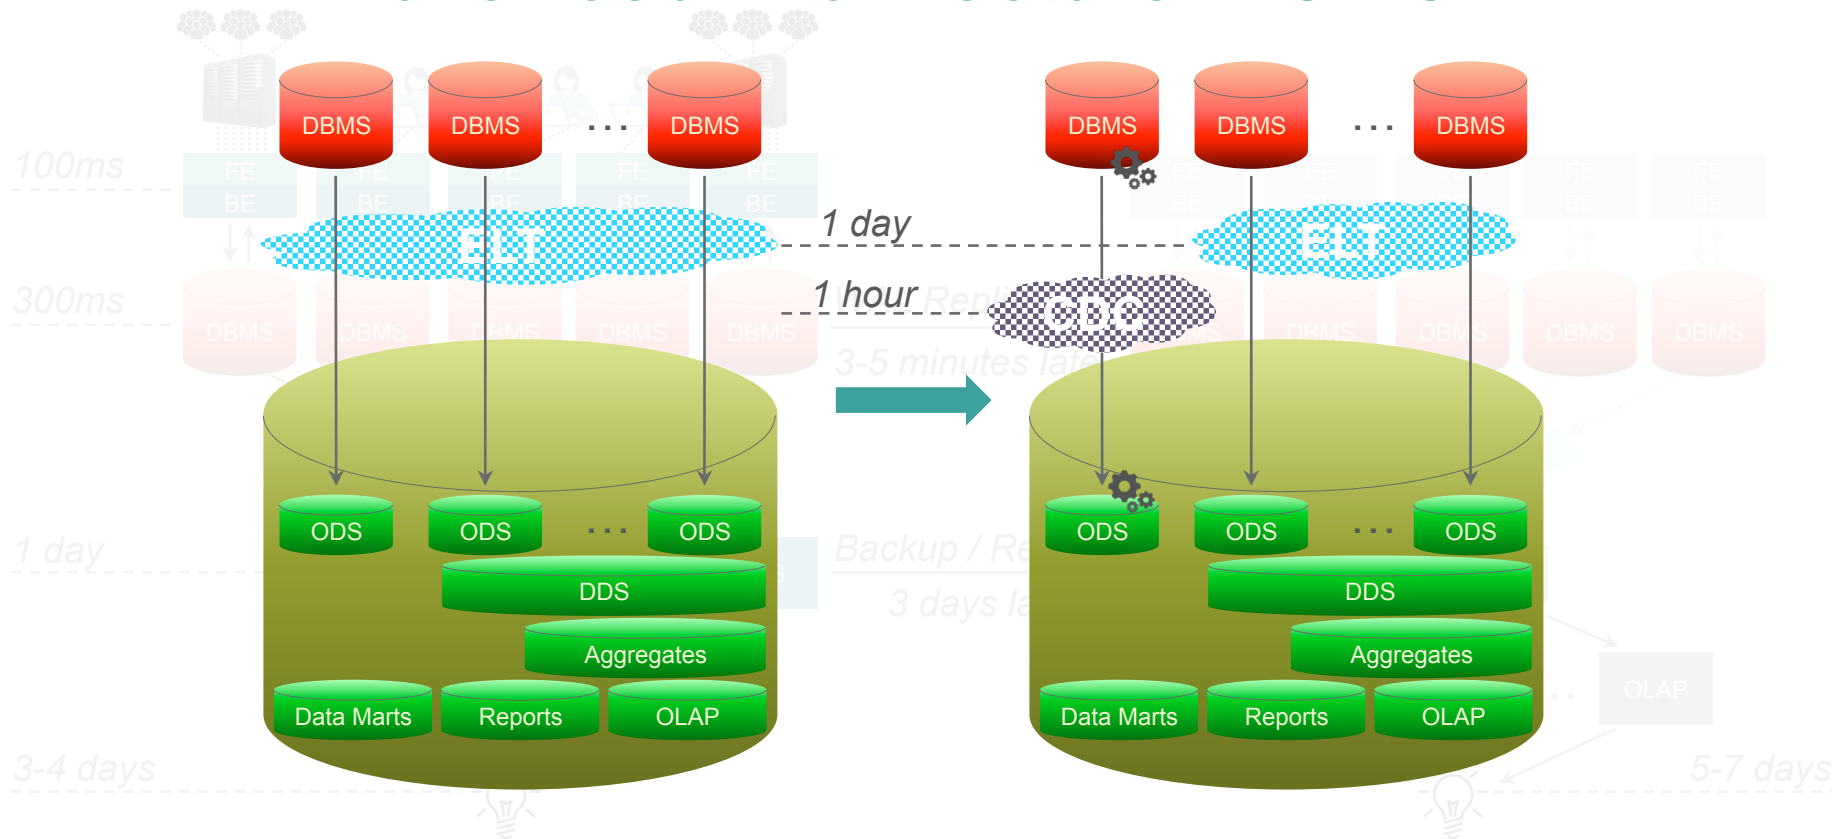
# Advanced Architecture – ELT



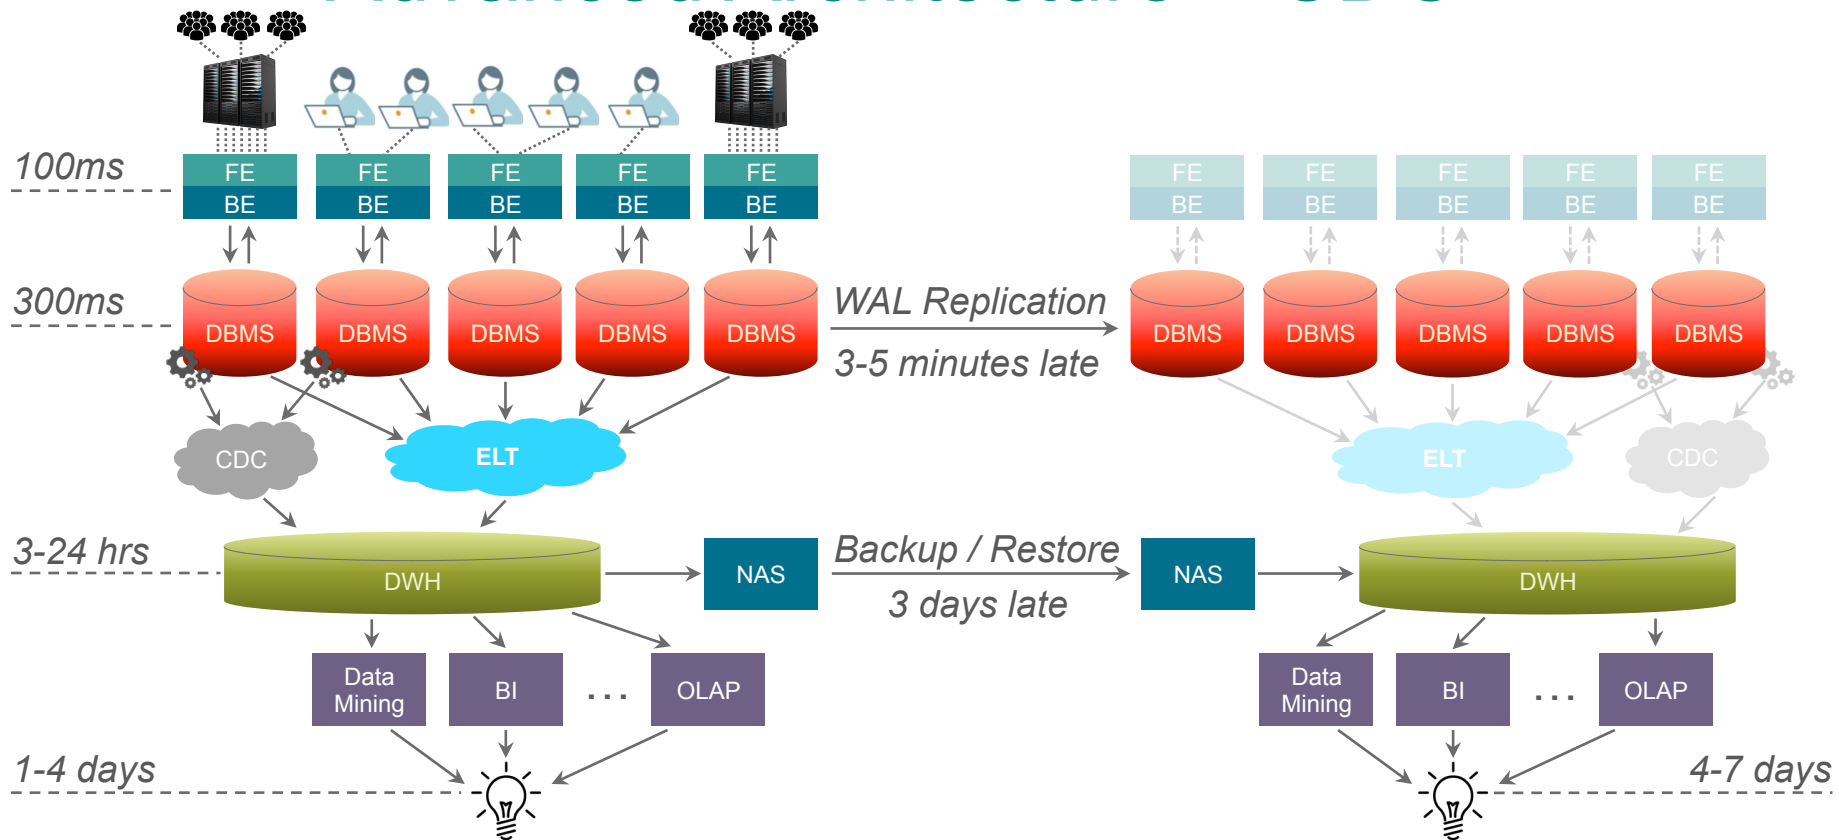
# Advanced Architecture – ELT



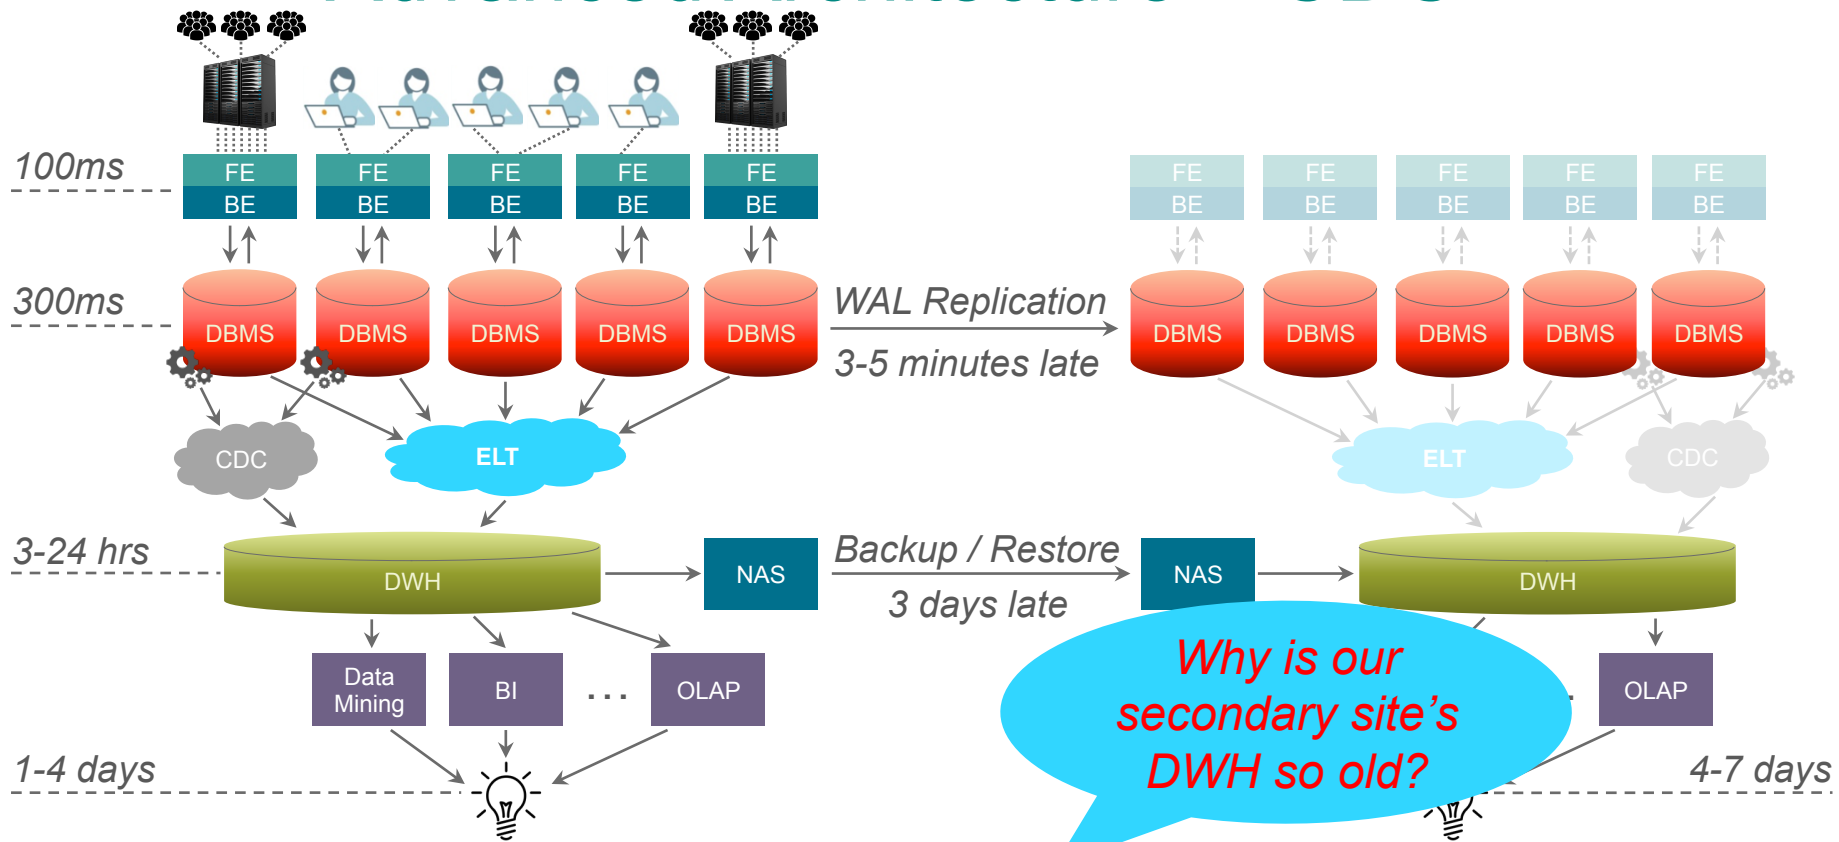
# Advanced Architecture – CDC



# Advanced Architecture – CDC

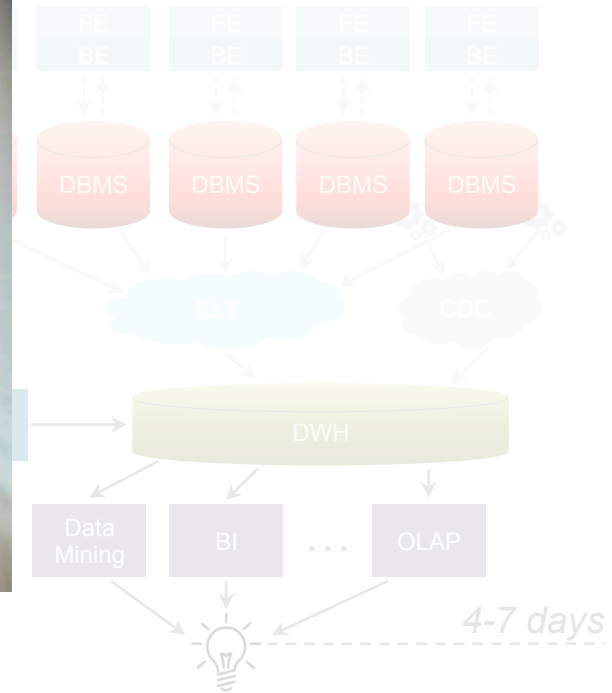
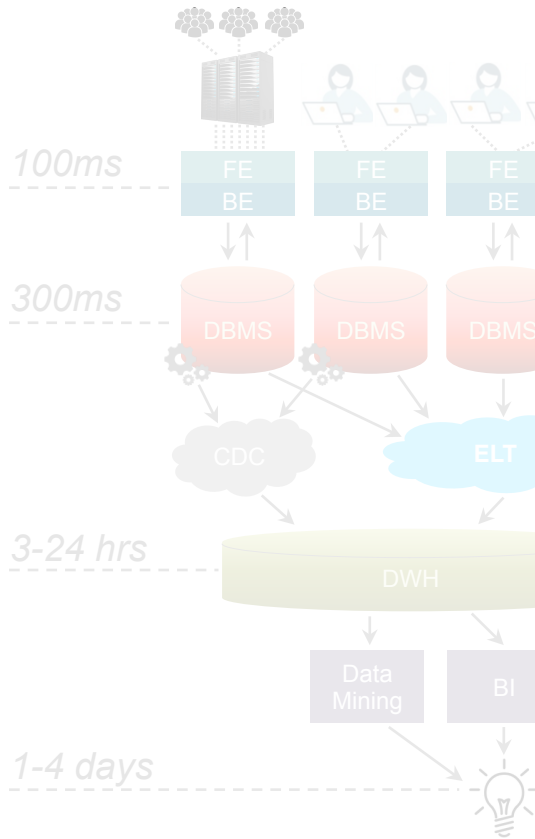


# Advanced Architecture – CDC



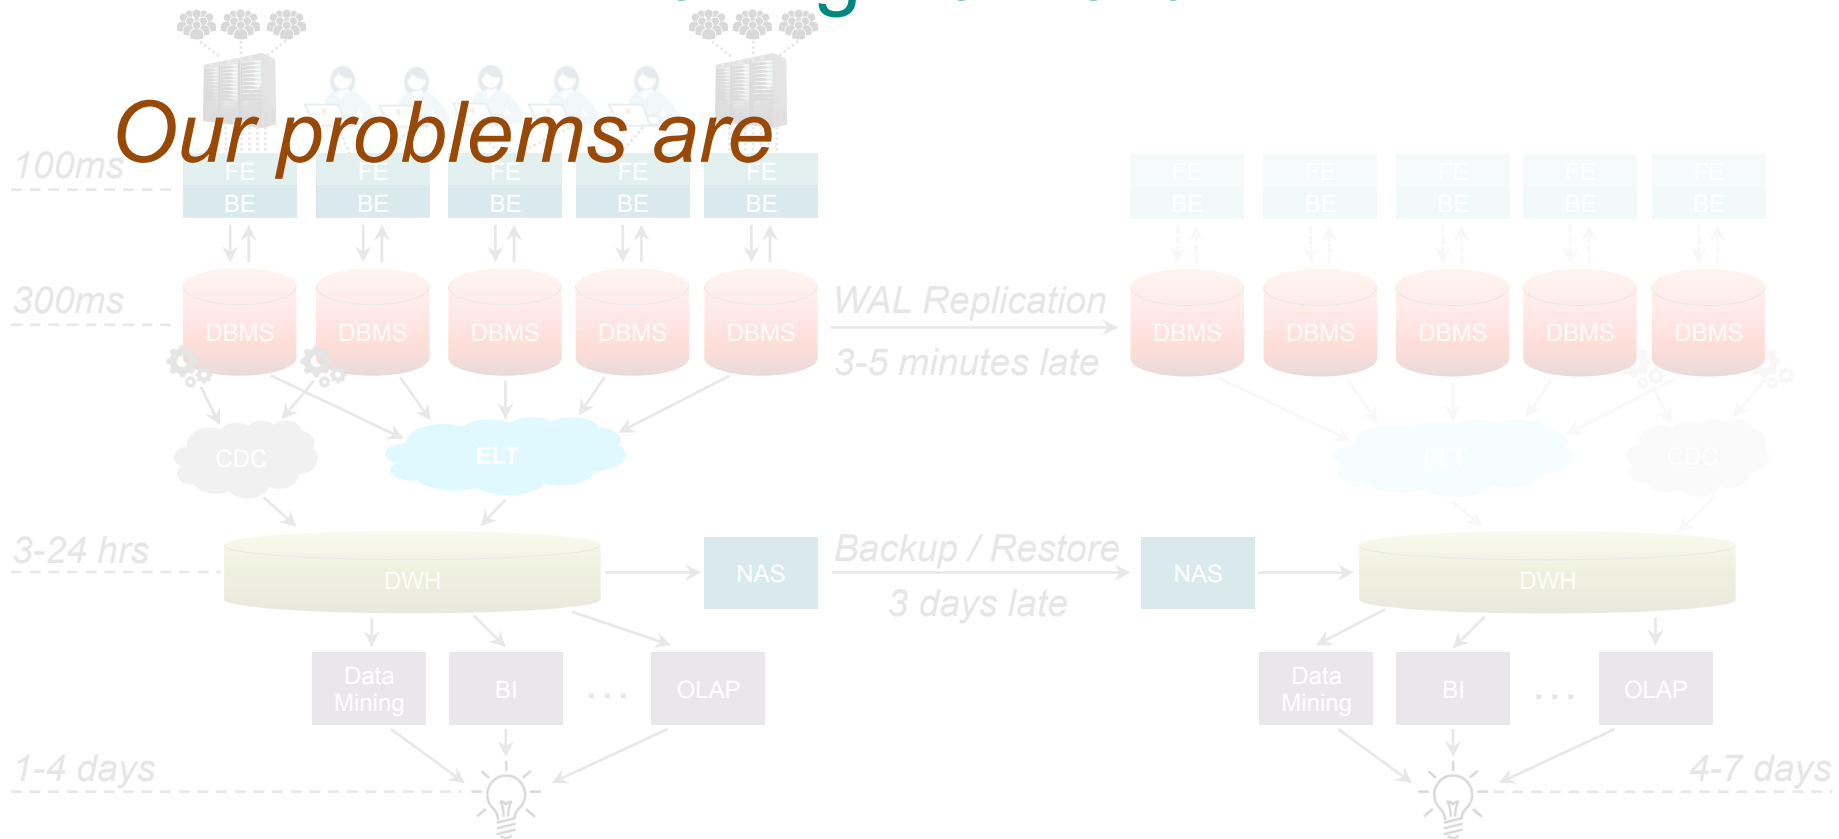


# Moving Forward



# Moving Forward

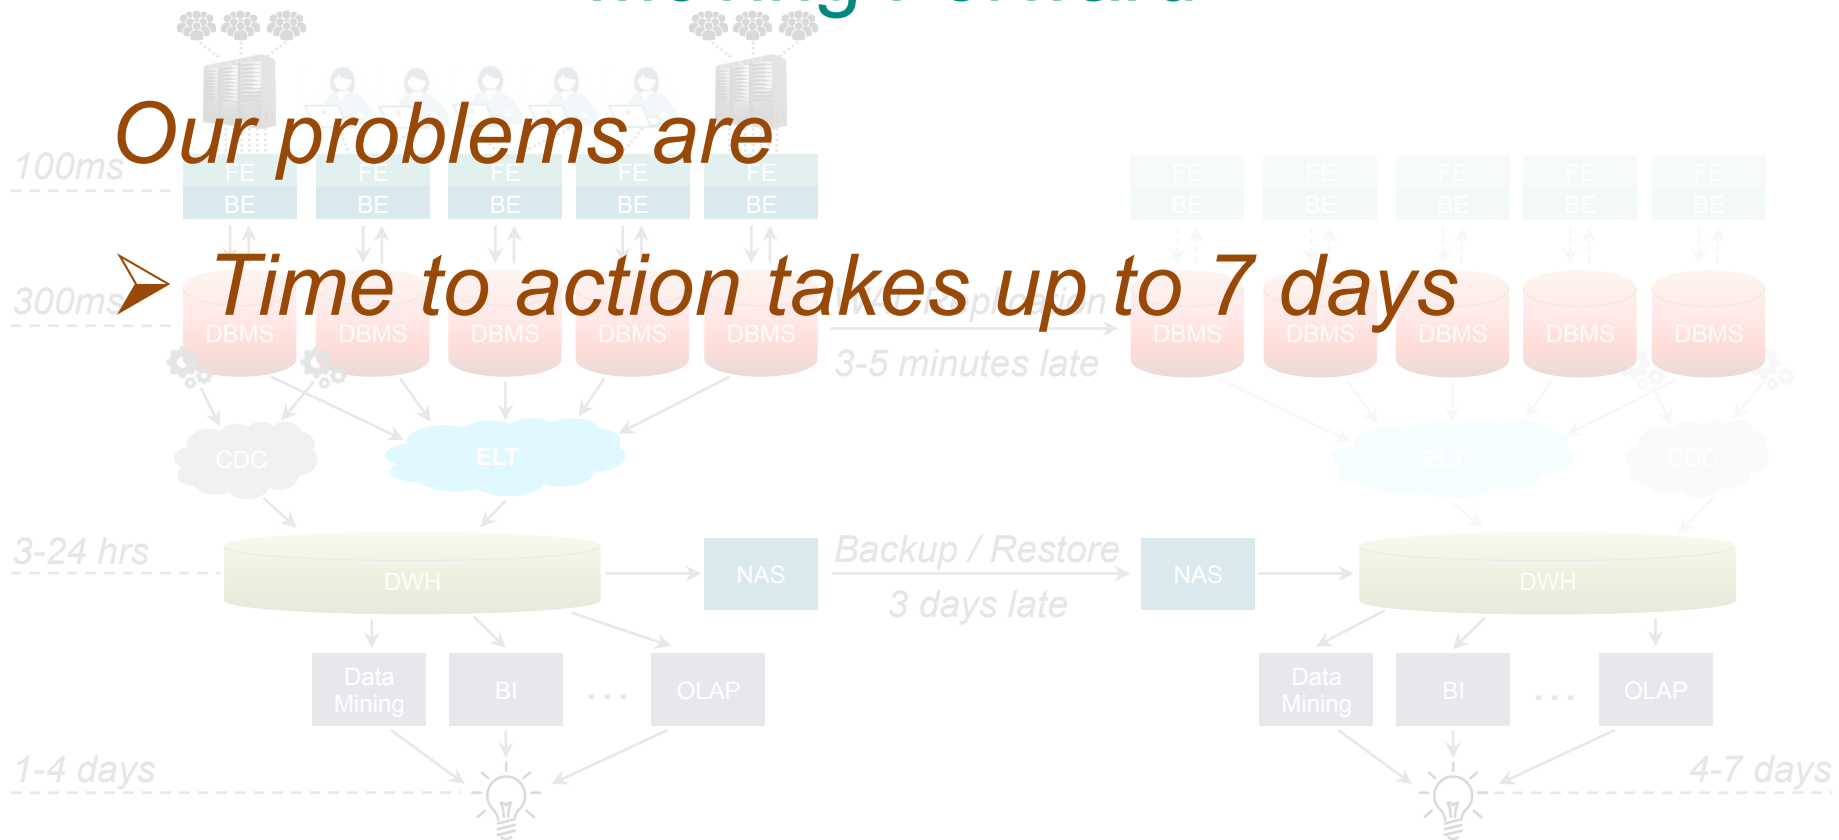
*Our problems are*



# Moving Forward

*Our problems are*

*Time to action takes up to 7 days*



# Moving Forward

*Our problems are*

100ms

300ms

3-24 hrs

1-4 days

➤ *Time to action takes up to 7 days*

➤ *Amount of data is growing*



# Moving Forward

*Our problems are*

100ms

300ms

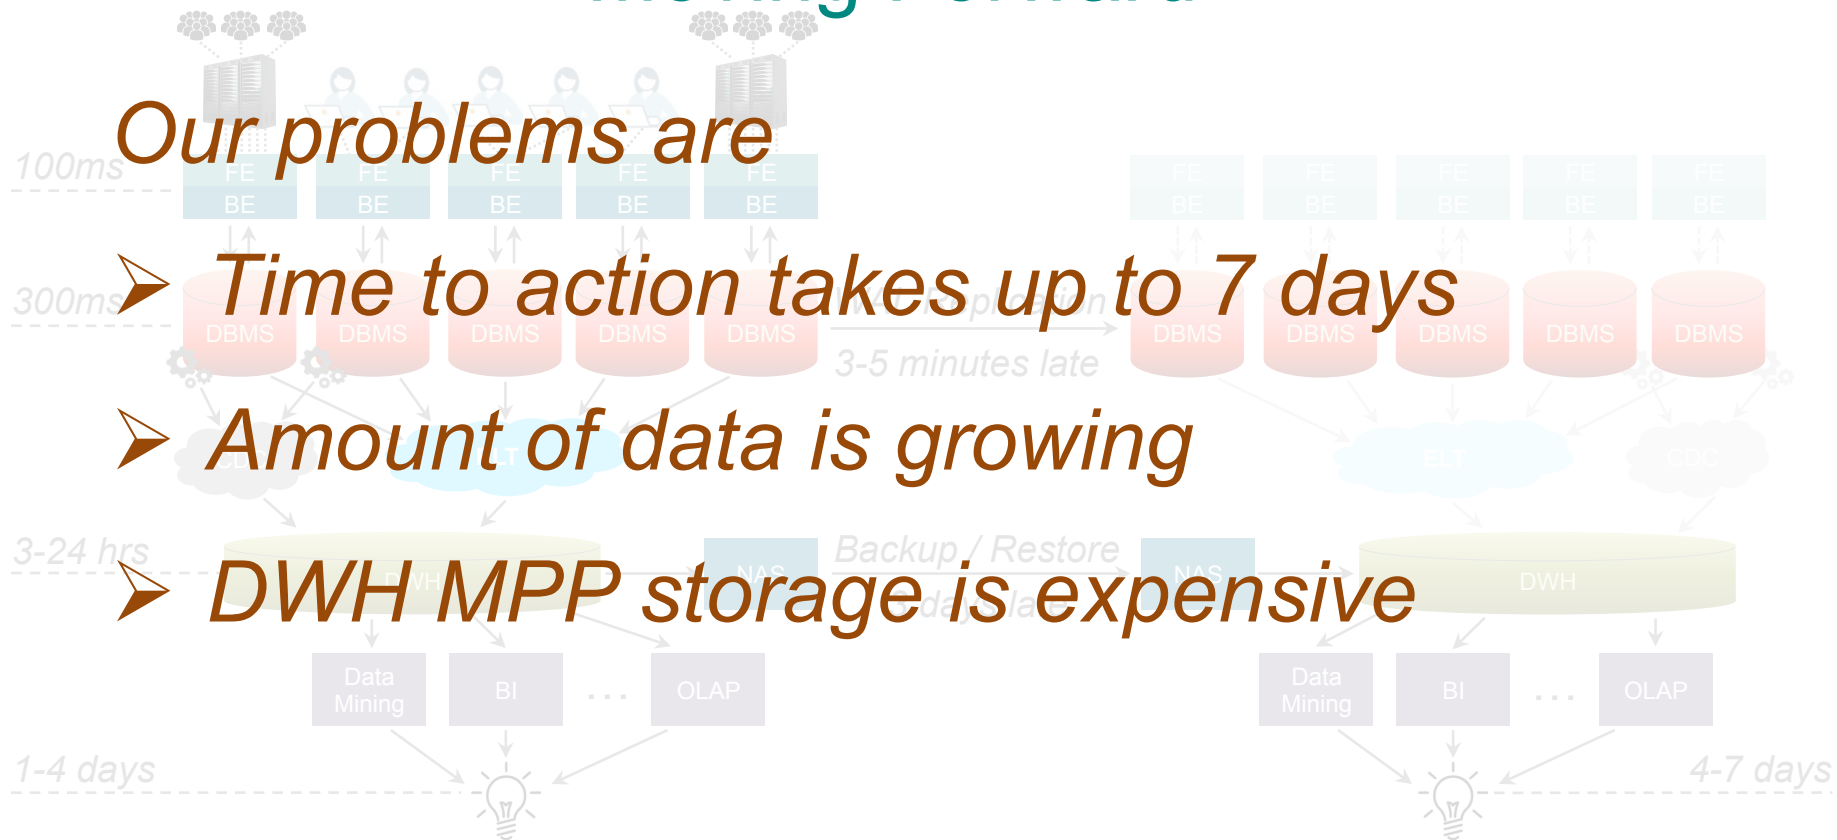
3-24 hrs

1-4 days

➤ *Time to action takes up to 7 days*

➤ *Amount of data is growing*

➤ *DWH MPP storage is expensive*



# Modern Architectures

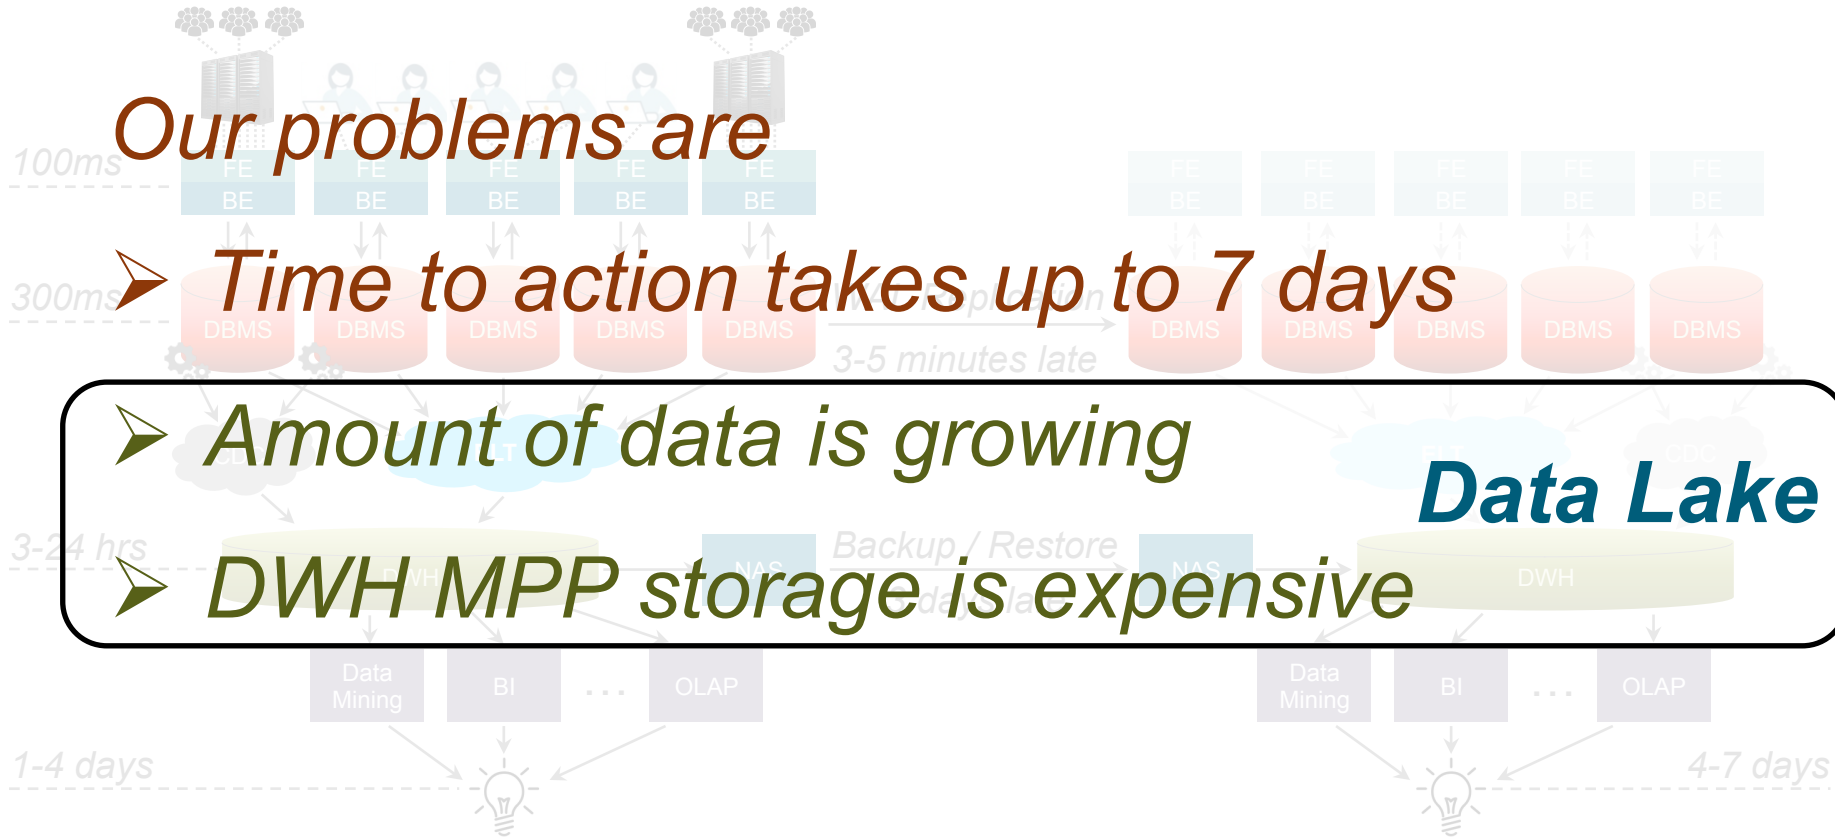
*Our problems are*

➤ *Time to action takes up to 7 days*

➤ *Amount of data is growing*

➤ *DWH MPP storage is expensive*

**Data Lake**



# Modern Architectures

*Our problems are*

➤ *Time to action takes up to 7 days* **Lambda**

➤ *Amount of data is growing*

**Data Lake**

➤ *DWH MPP storage is expensive*

100ms

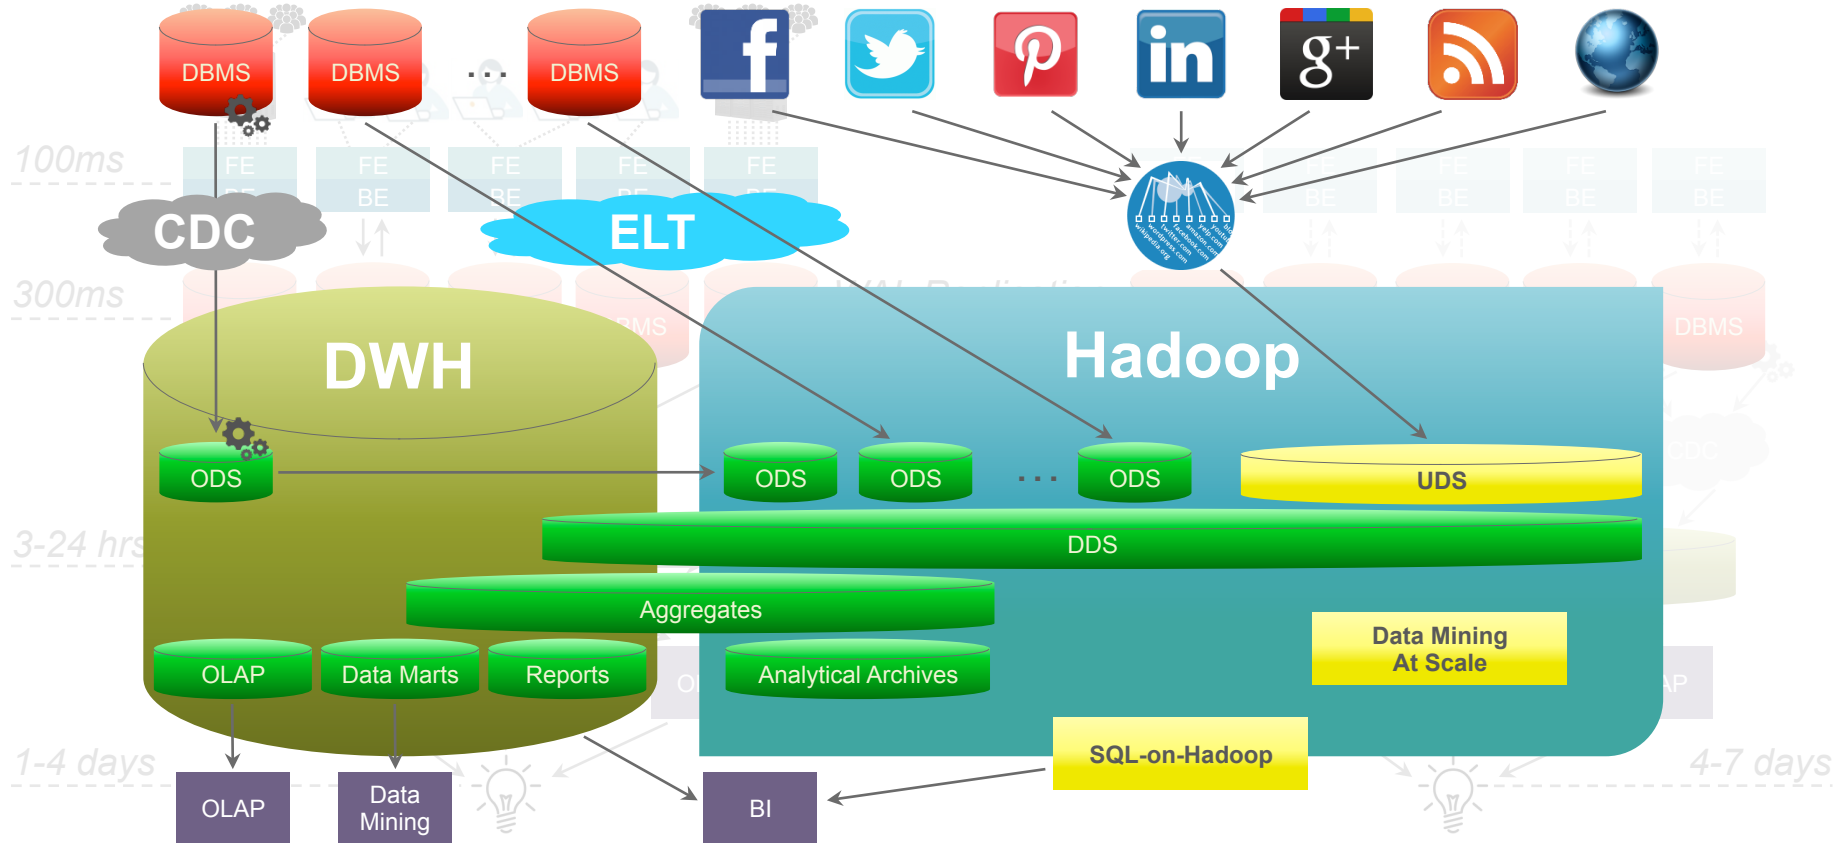
300ms

3-24 hrs

1-4 days

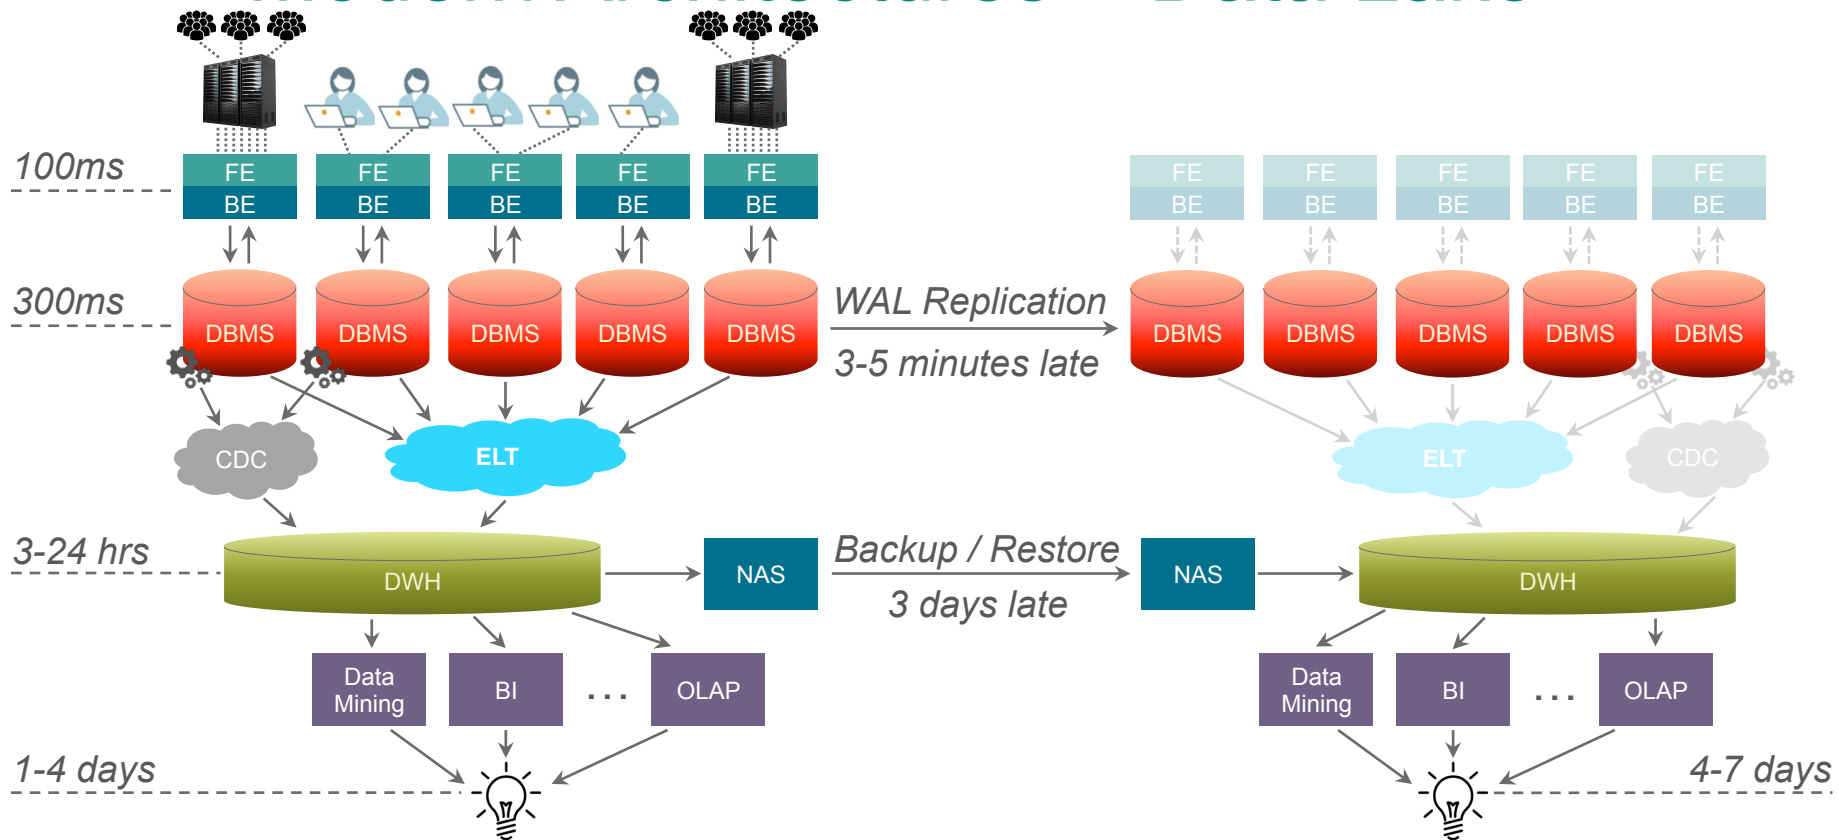
4-7 days

# Modern Architectures – Data Lake

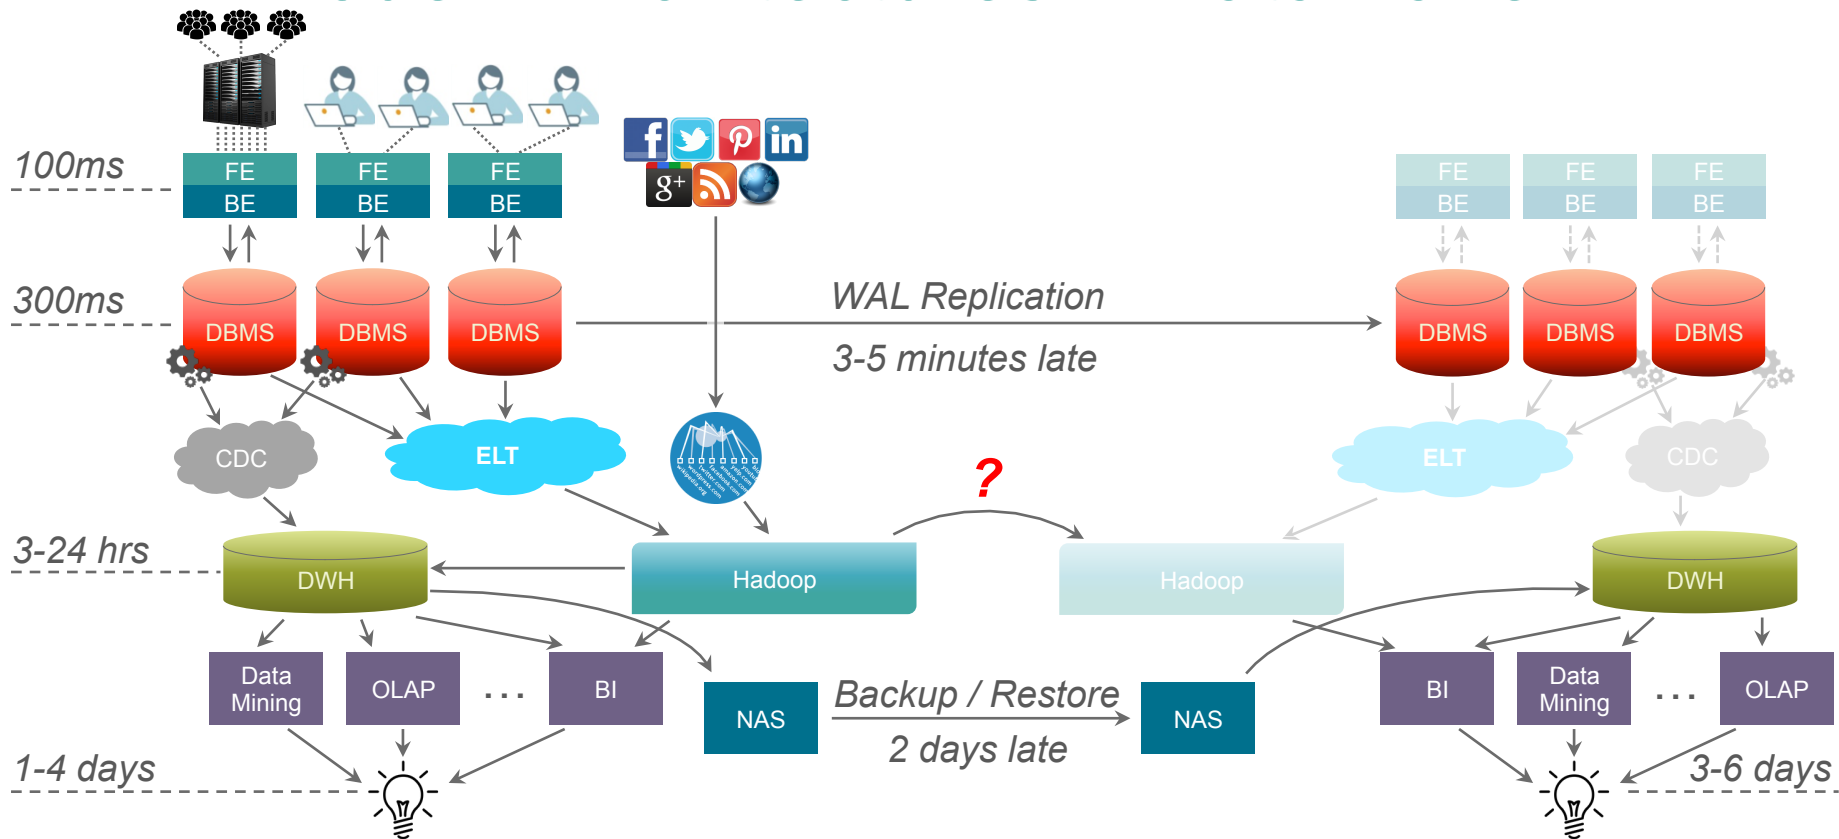




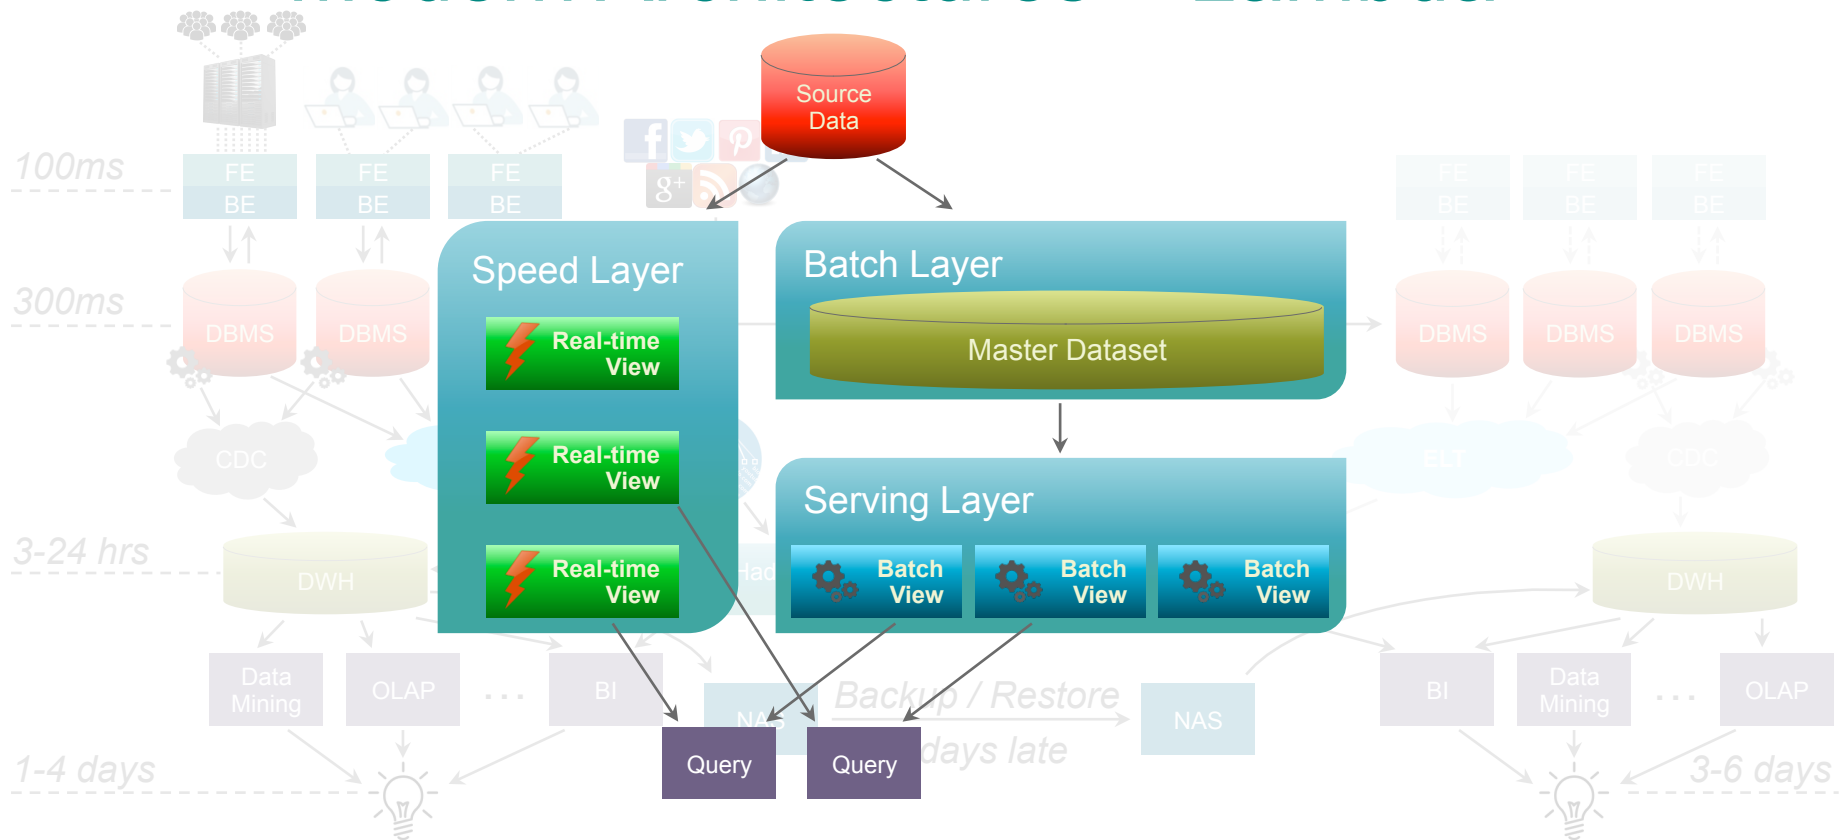
# Modern Architectures – Data Lake



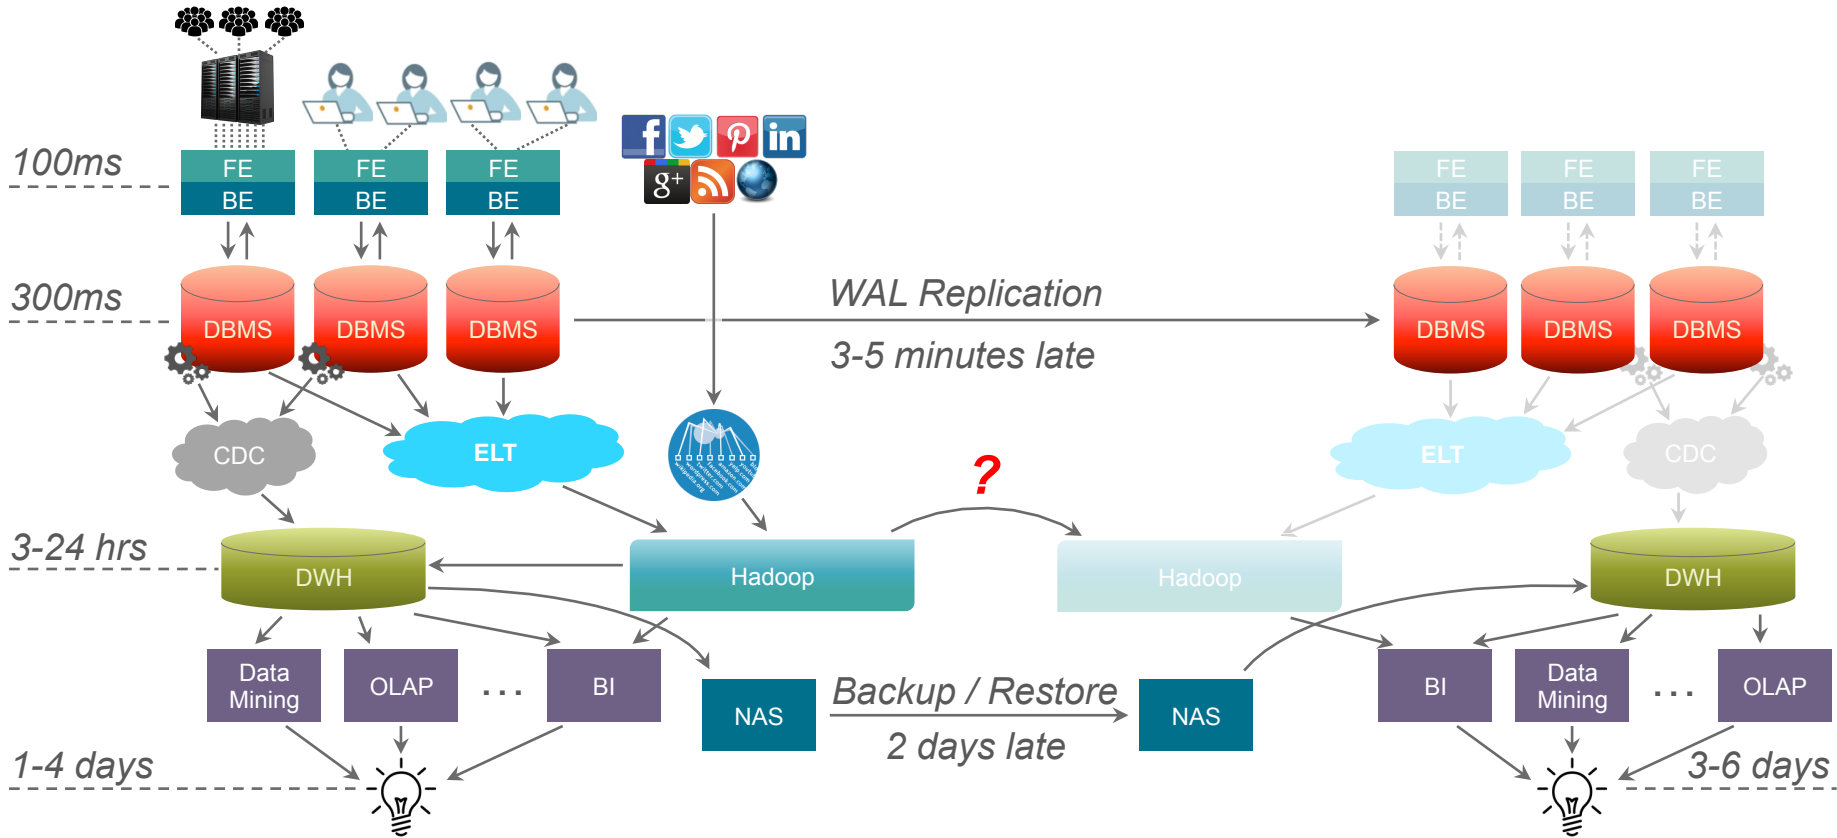
# Modern Architectures – Data Lake



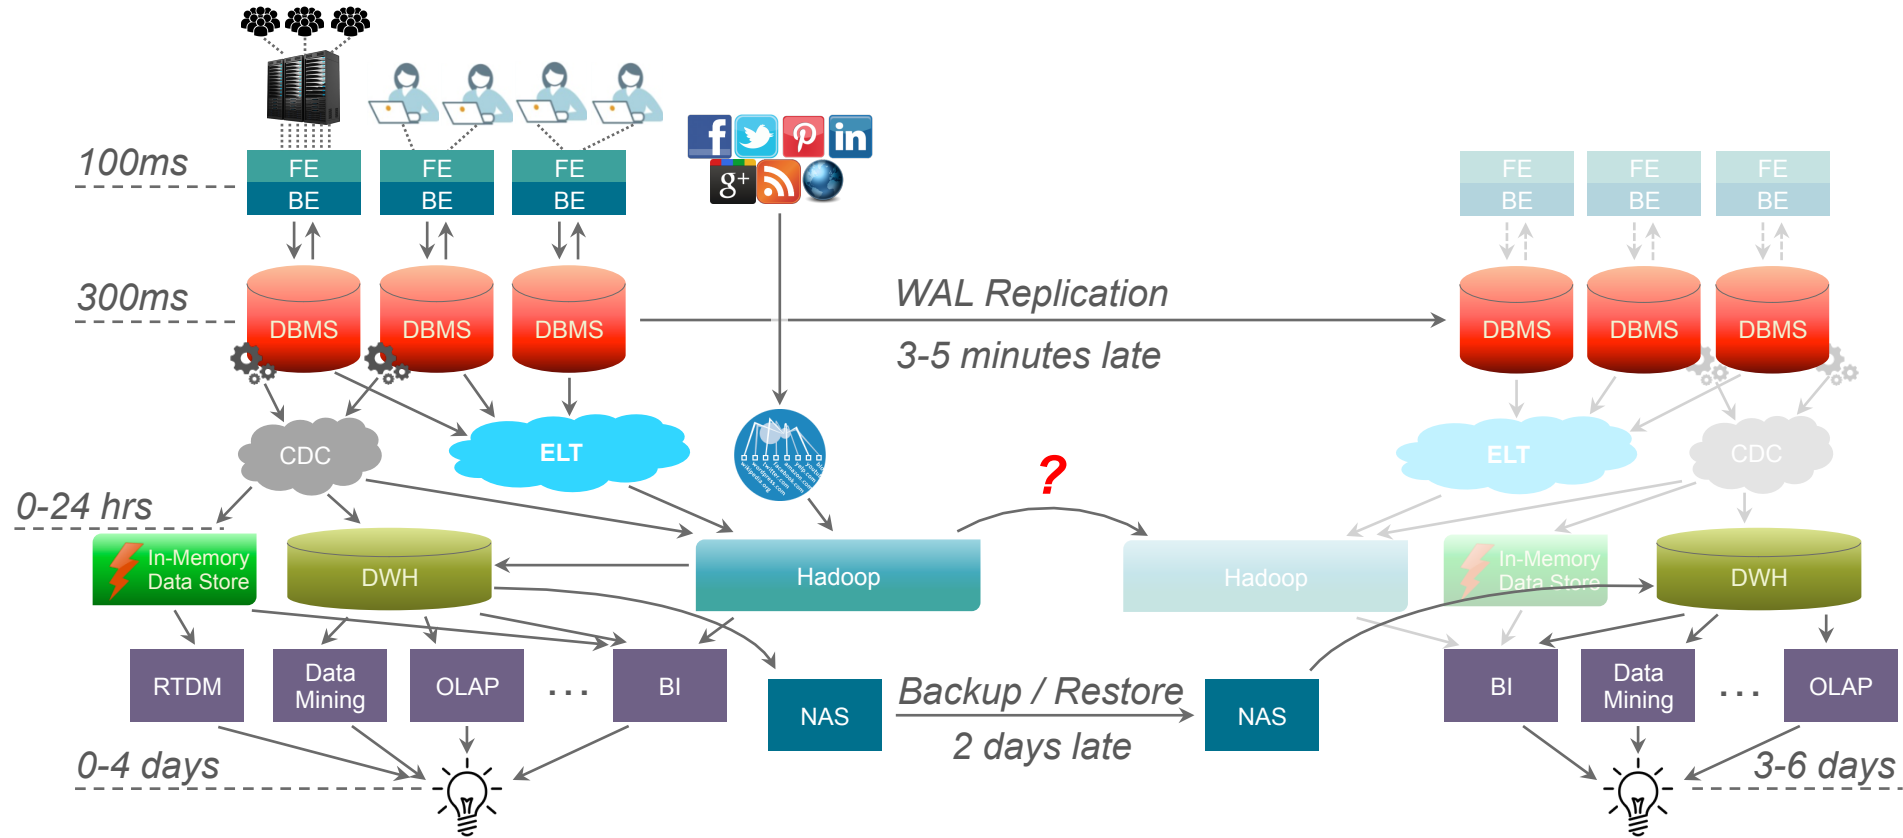
# Modern Architectures – Lambda



# Modern Architectures – Lambda

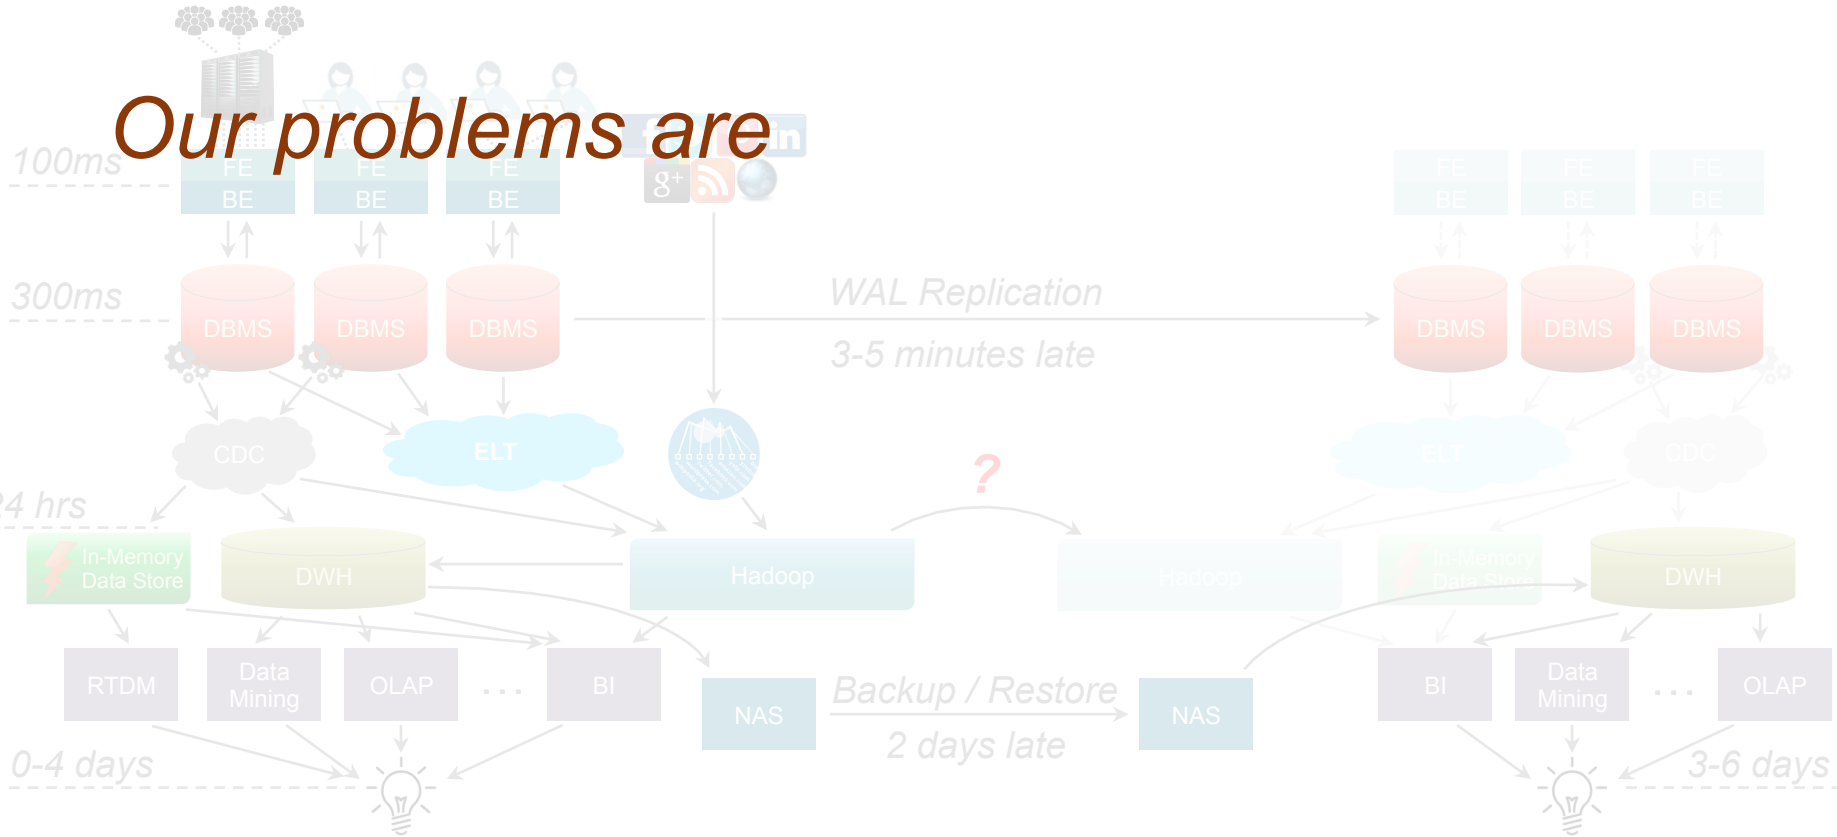


# Modern Architectures – Lambda



# Modern Architectures

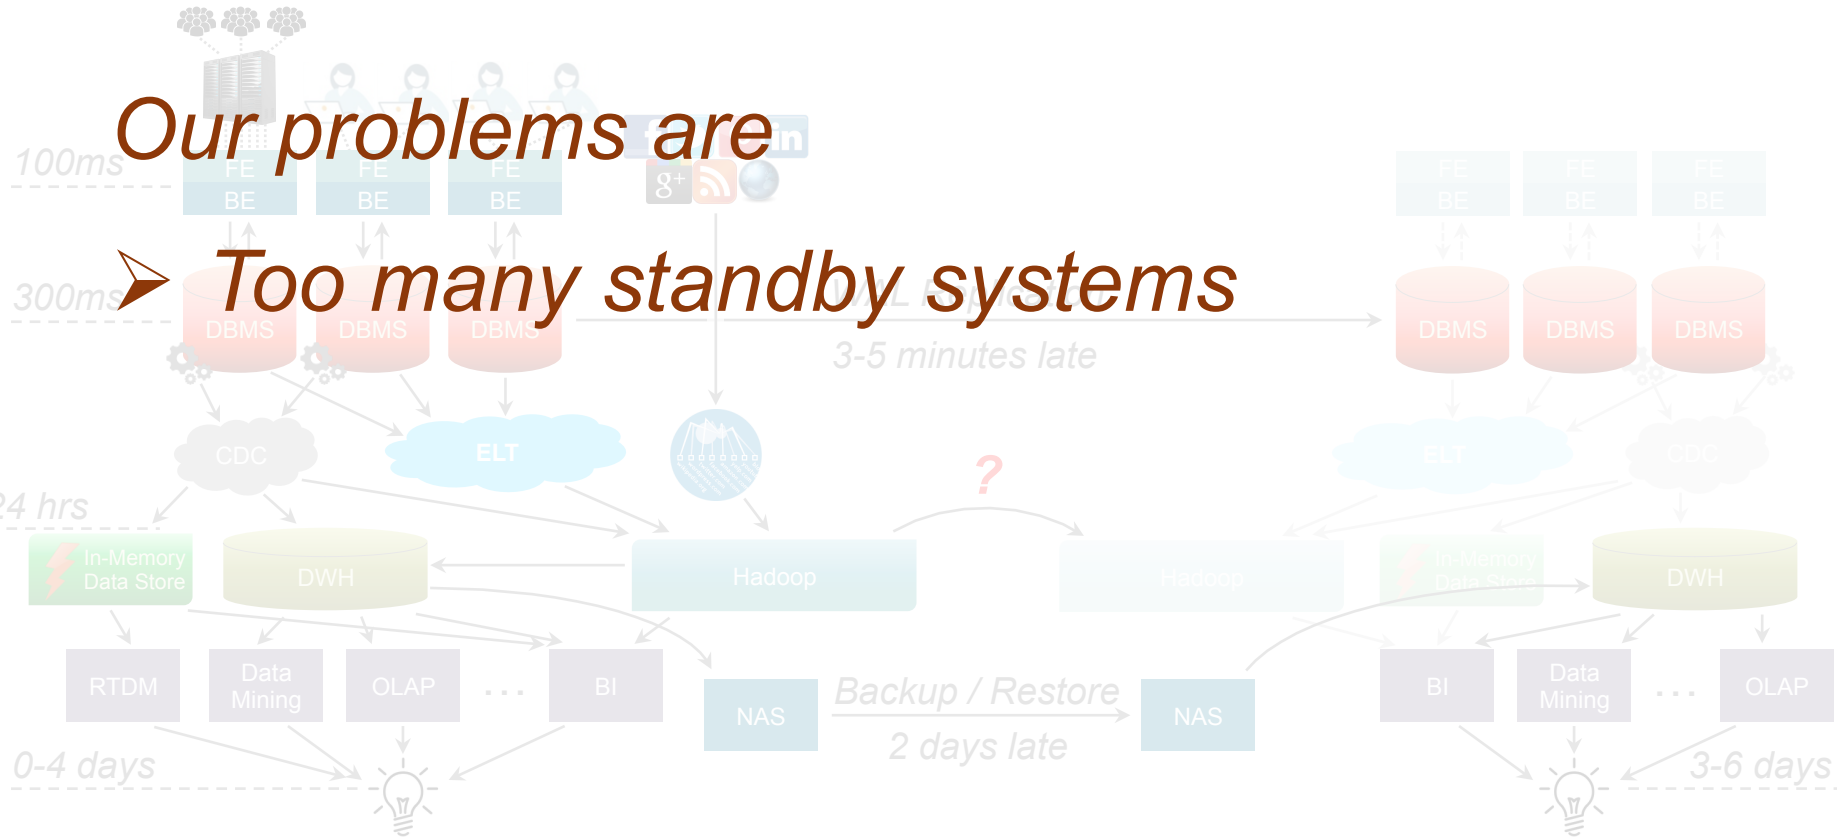
*Our problems are*



# Modern Architectures

*Our problems are*

*Too many standby systems*



# Modern Architectures

*Our problems are*


100ms



300ms

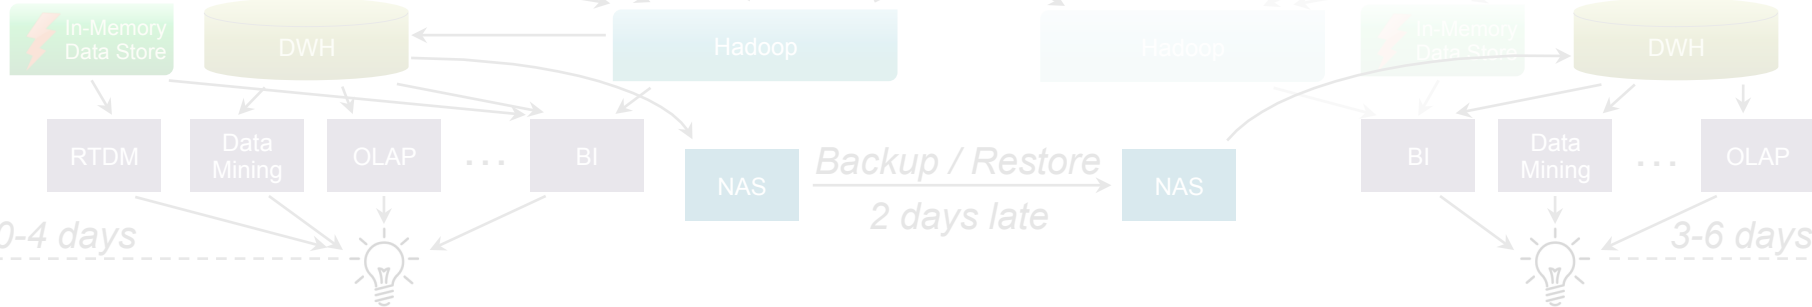
➤ *Too many standby systems*

3-5 minutes late



➤ *How to replicate Hadoop cluster?*

0-24 hrs



0-4 days

3-6 days



# Modern Architectures

*Our problems are*

100ms

300ms

3-5 minutes late

Backup / Restore  
2 days late

➤ *Too many standby systems*

➤ *How to replicate Hadoop cluster?*

➤ *How to sync data in real-time systems?*

0-4 days

3-6 days

# Modern Architectures

*Our problems are*

100ms

300ms

3-5 minutes late

0-24 hrs

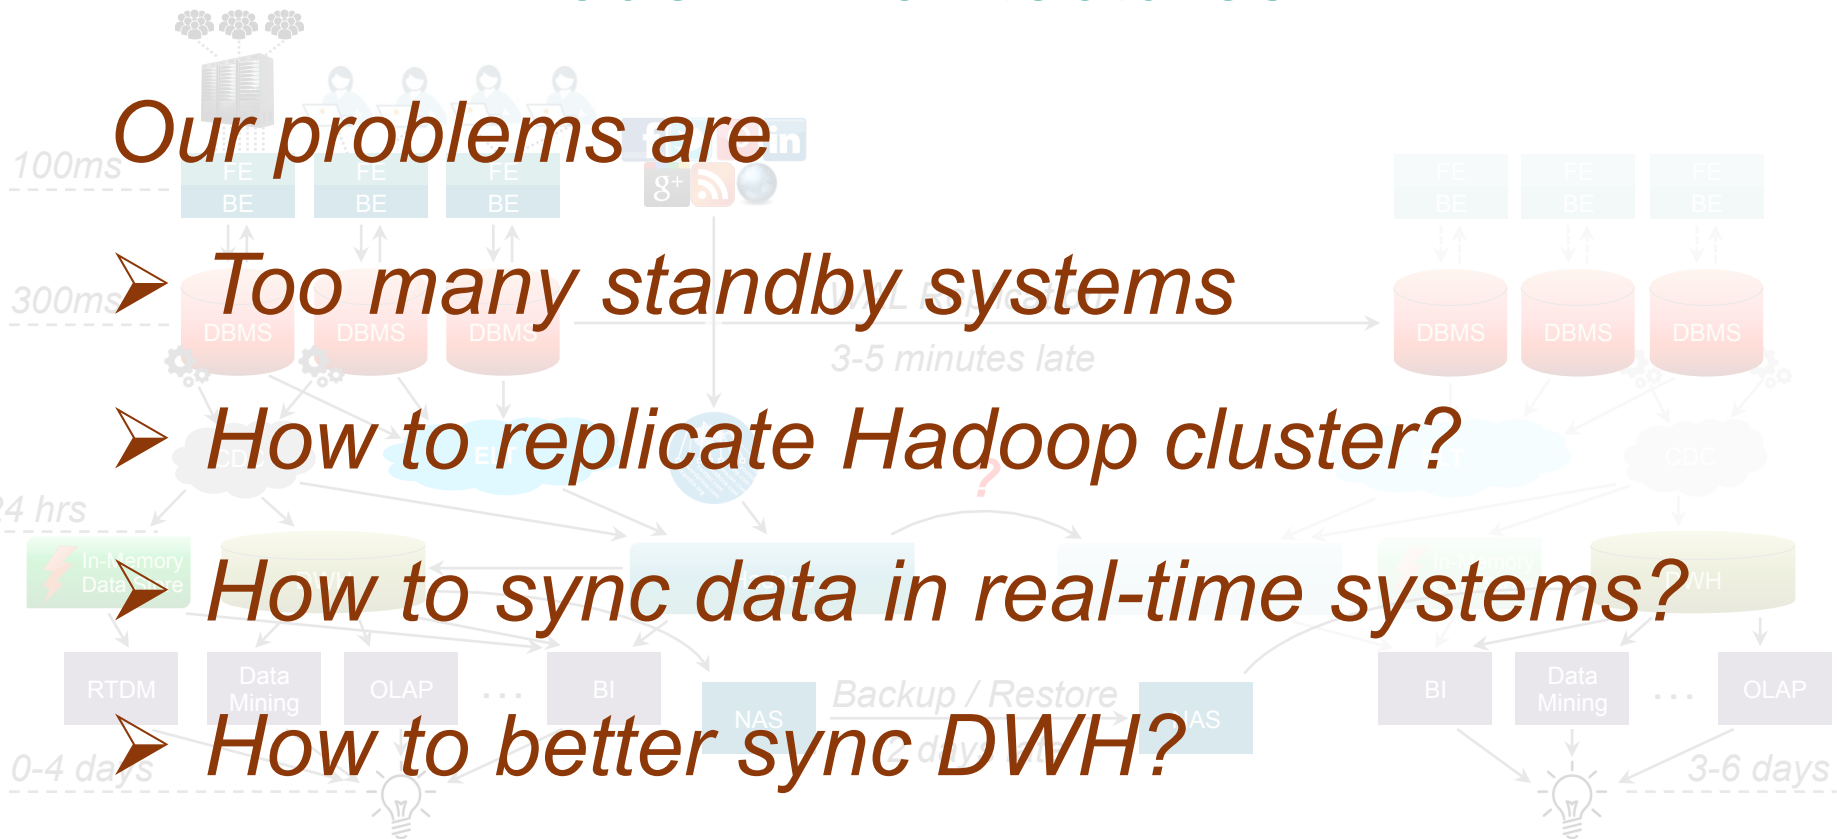
0-4 days

➤ *Too many standby systems*

➤ *How to replicate Hadoop cluster?*

➤ *How to sync data in real-time systems?*

➤ *How to better sync DWH?*



# Modern Architectures

*Our problems are*

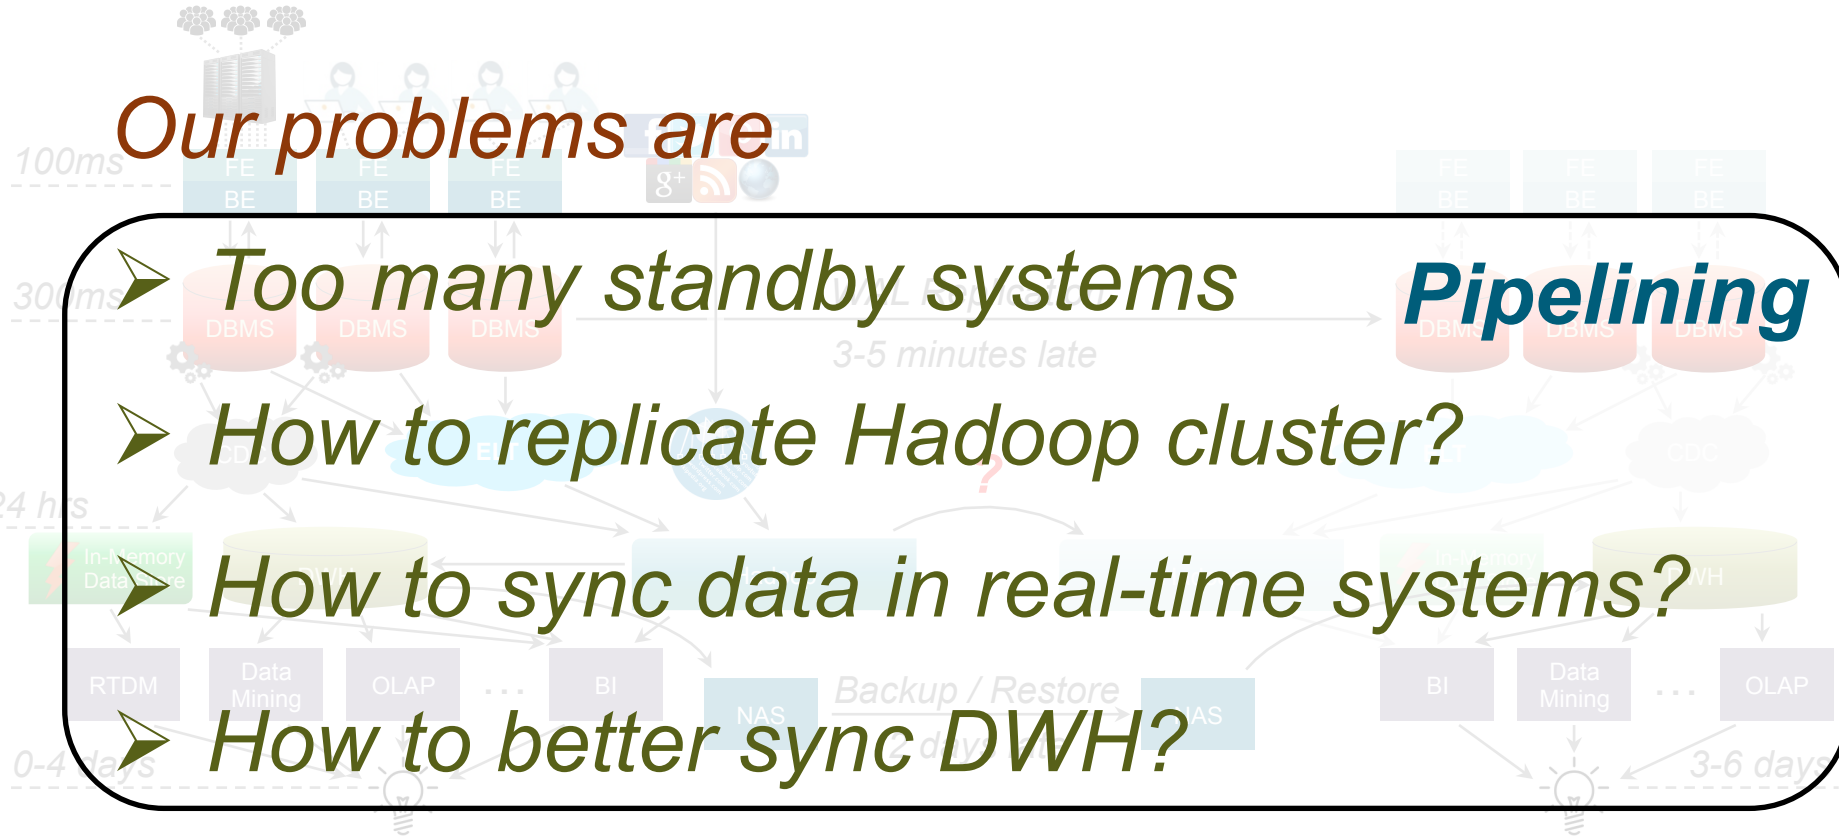
➤ *Too many standby systems*

**Pipelining**

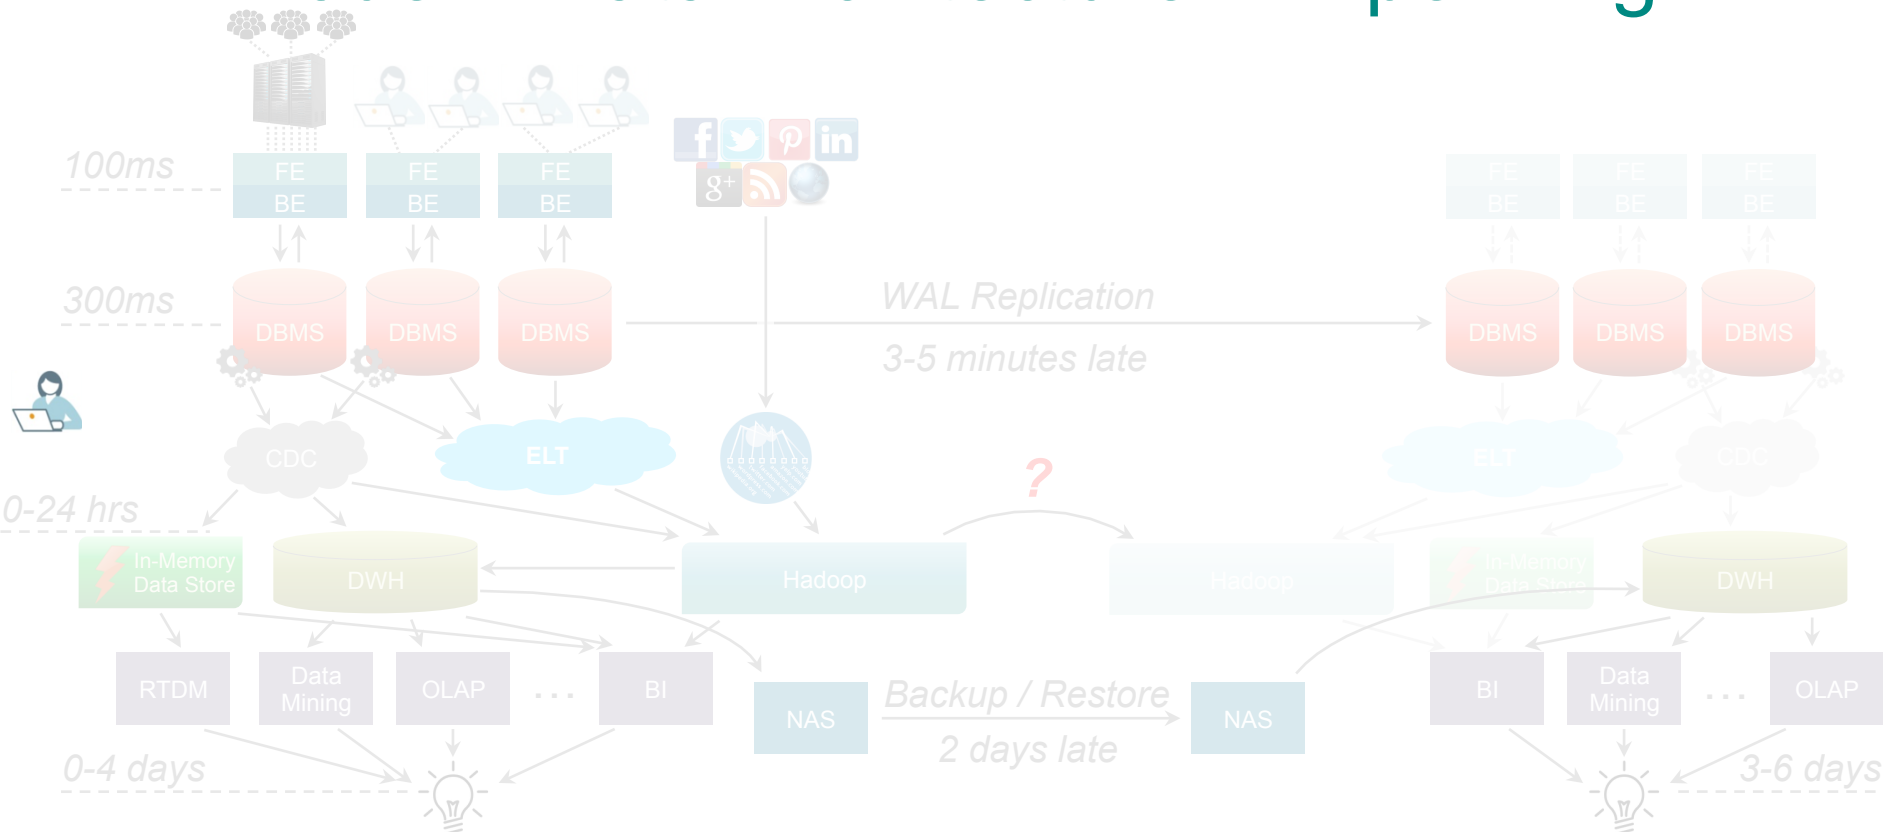
➤ *How to replicate Hadoop cluster?*

➤ *How to sync data in real-time systems?*

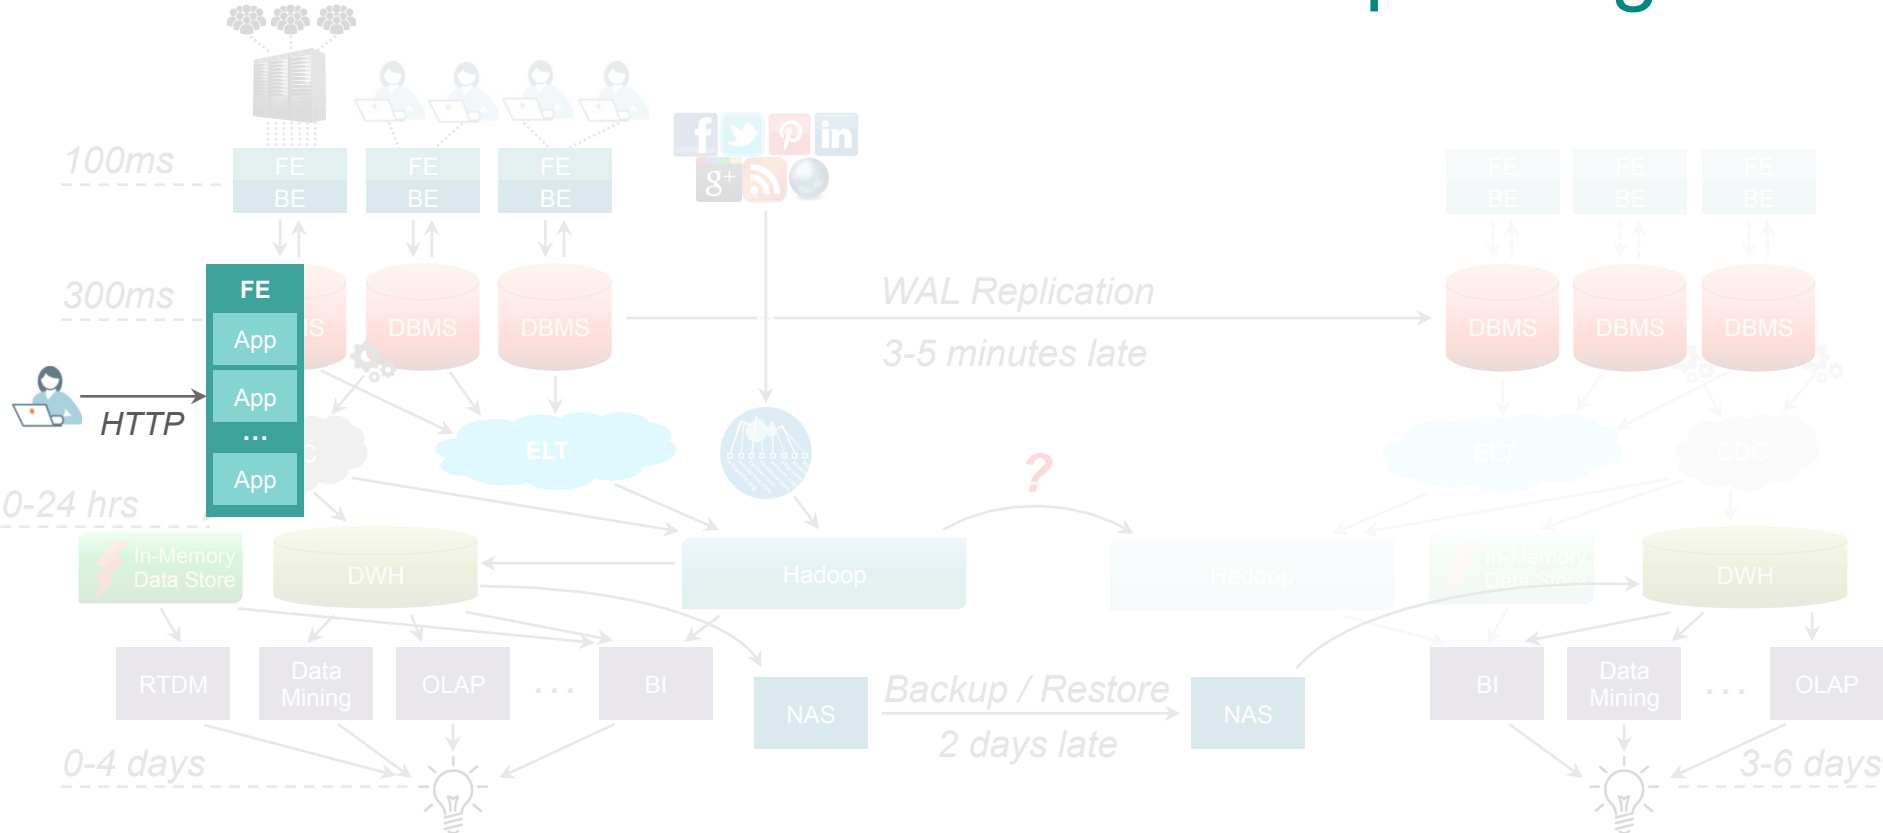
➤ *How to better sync DWH?*



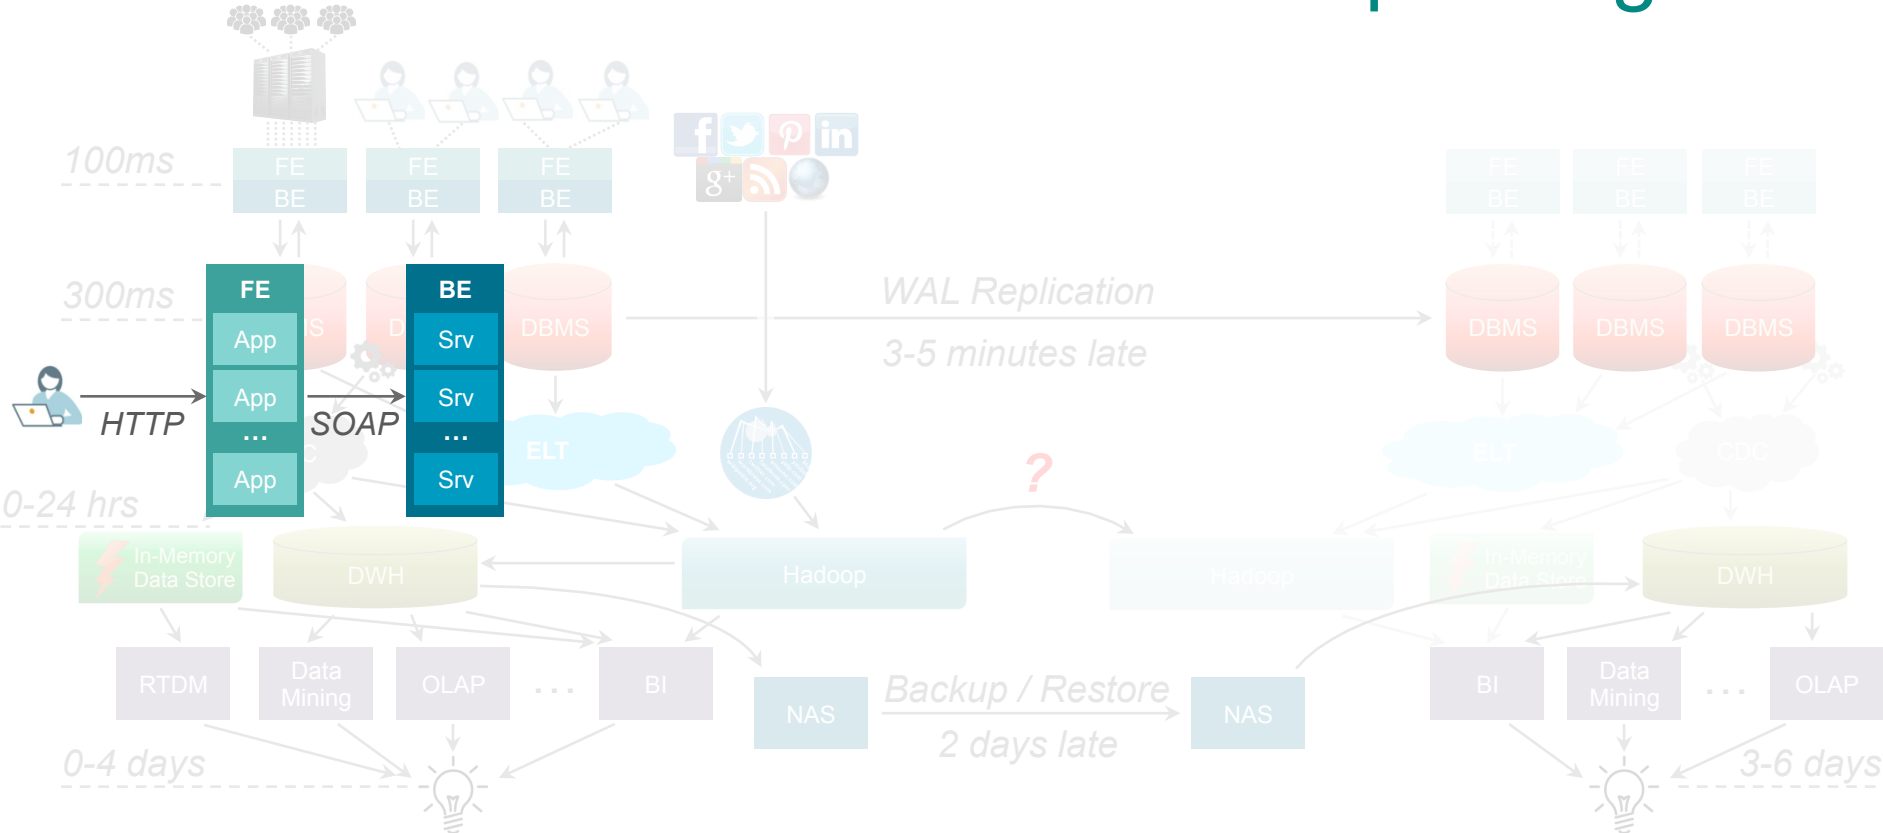
# Modern Data Architecture – Pipelining



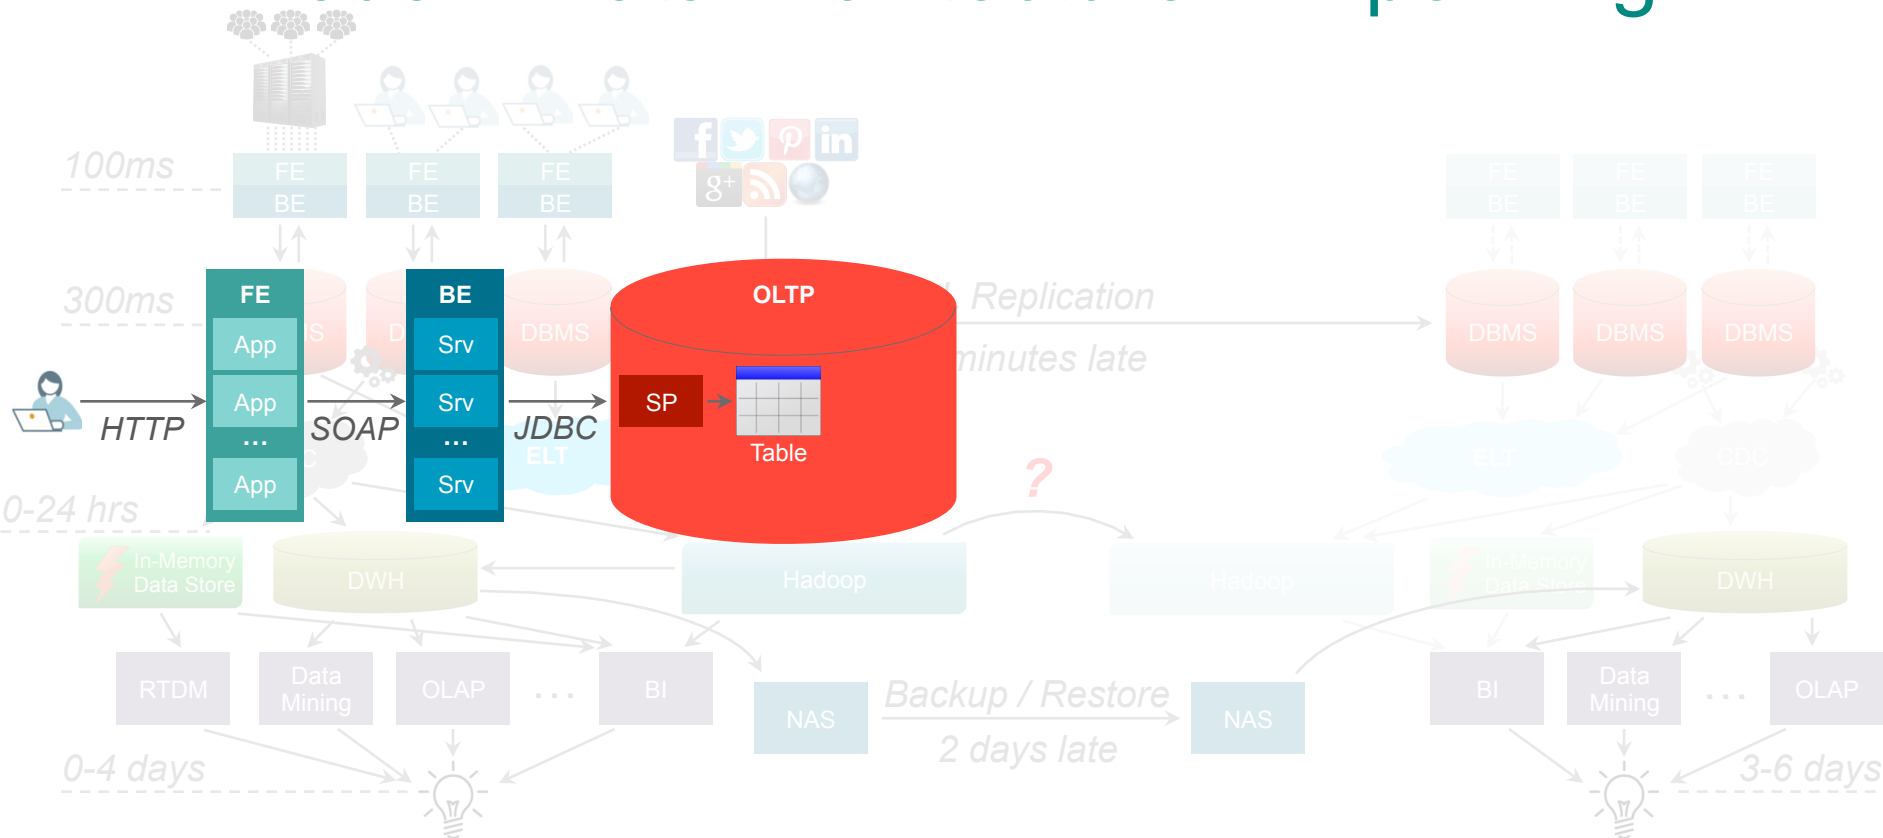
# Modern Data Architecture – Pipelining



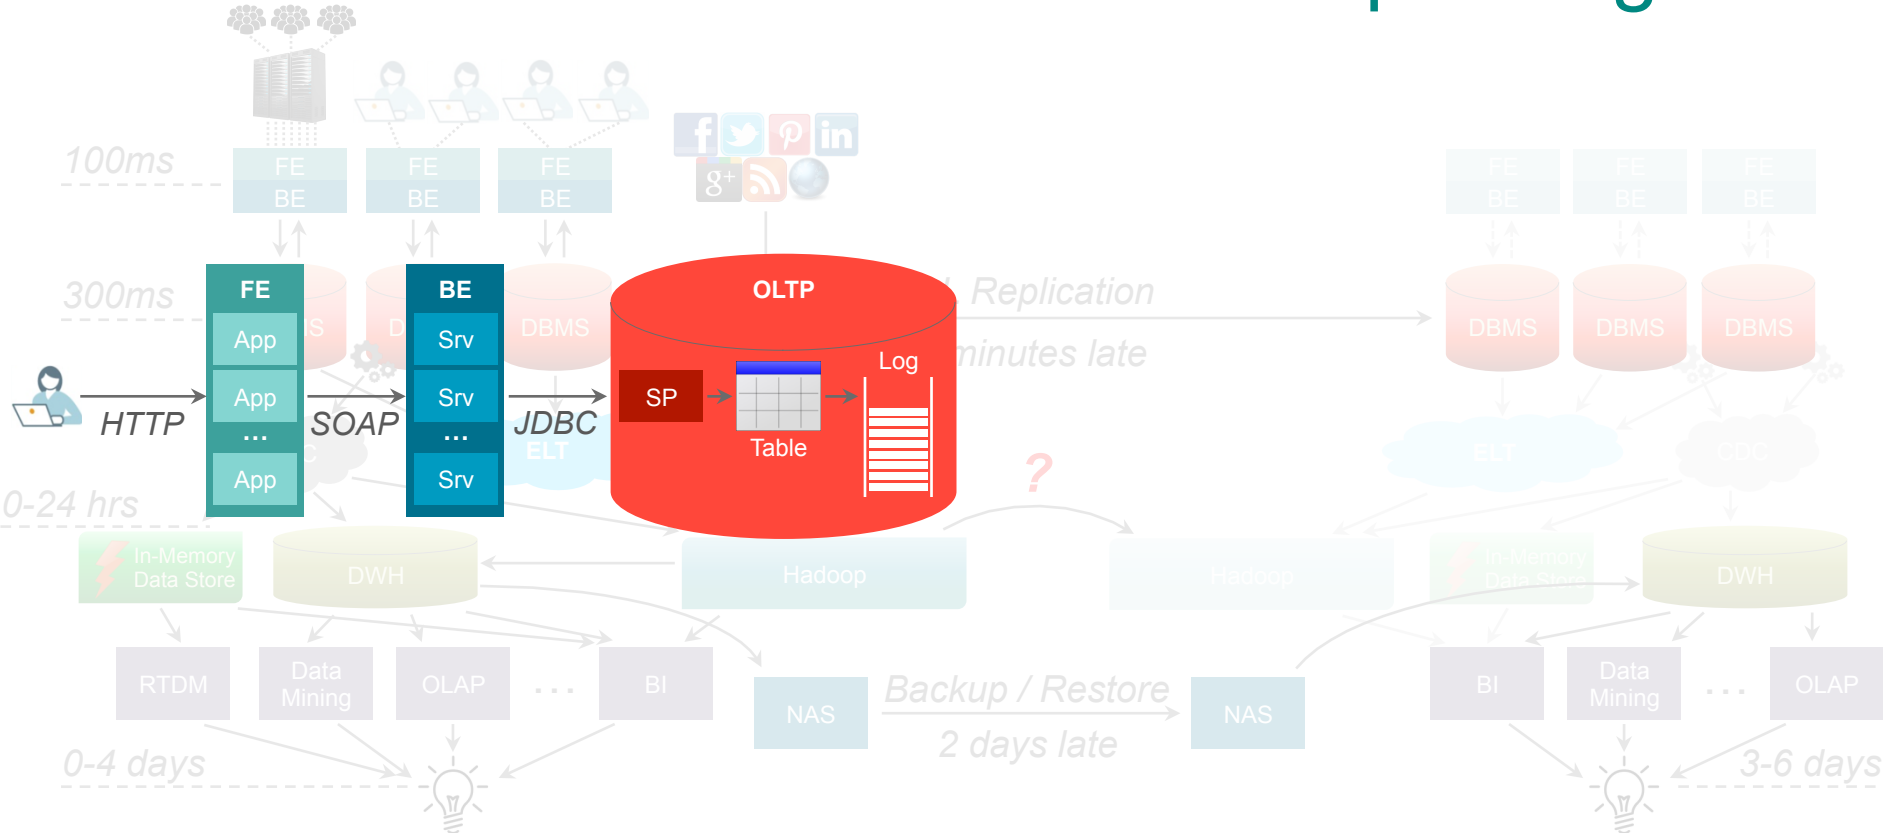
# Modern Data Architecture – Pipelining



# Modern Data Architecture – Pipelining

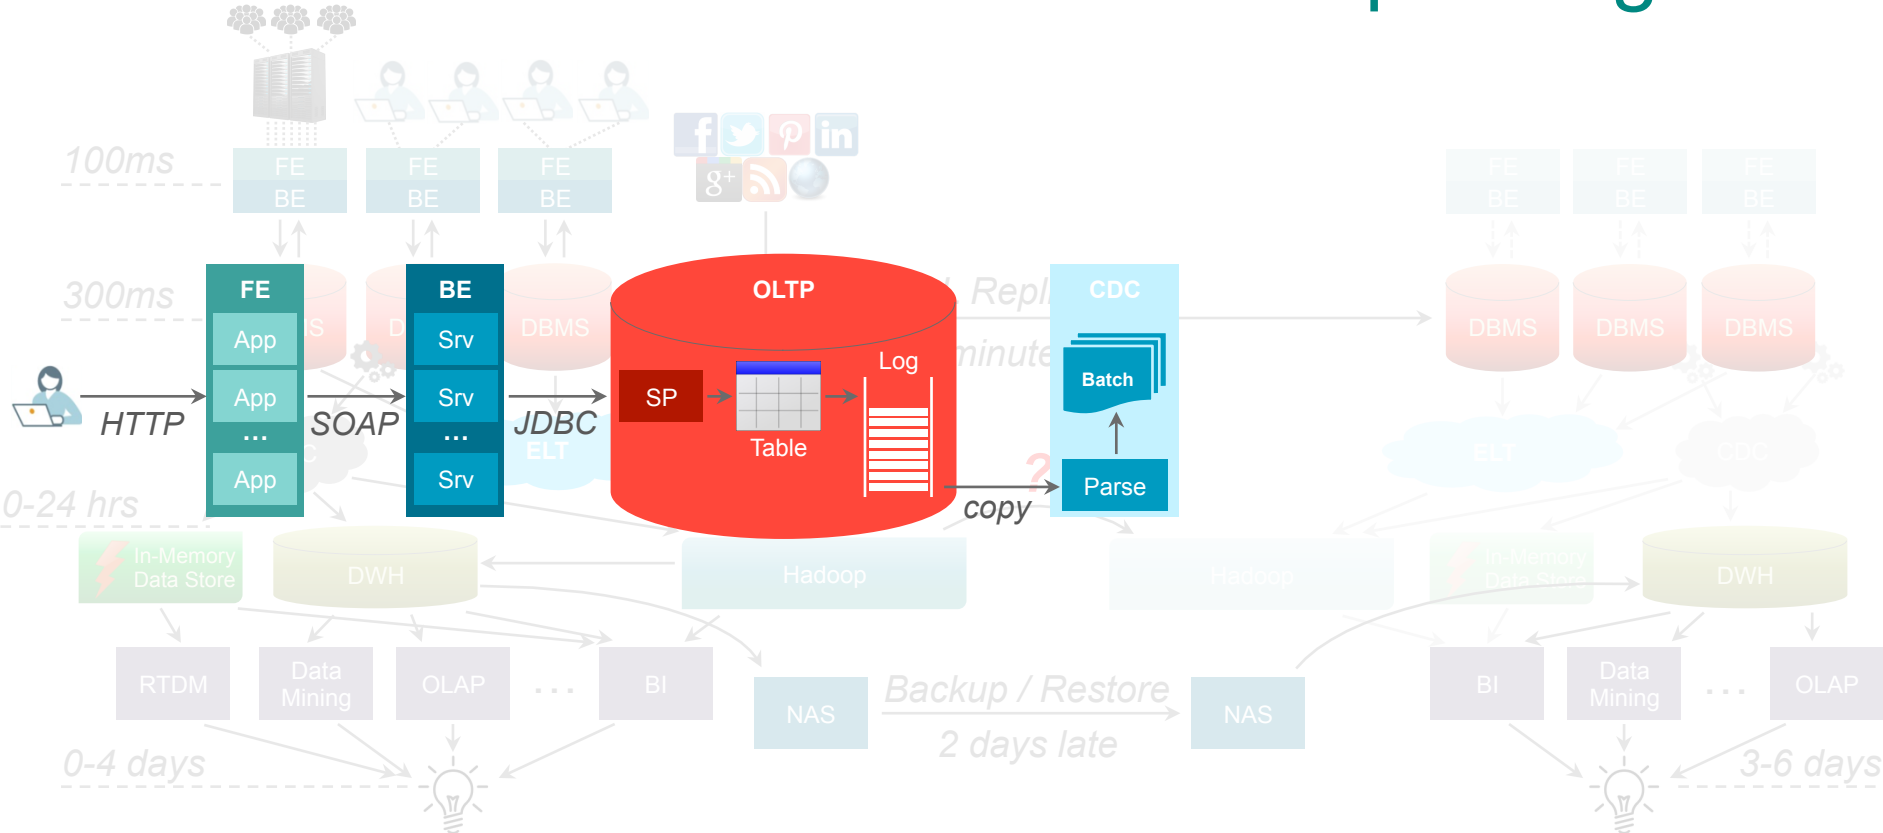


# Modern Data Architecture – Pipelining

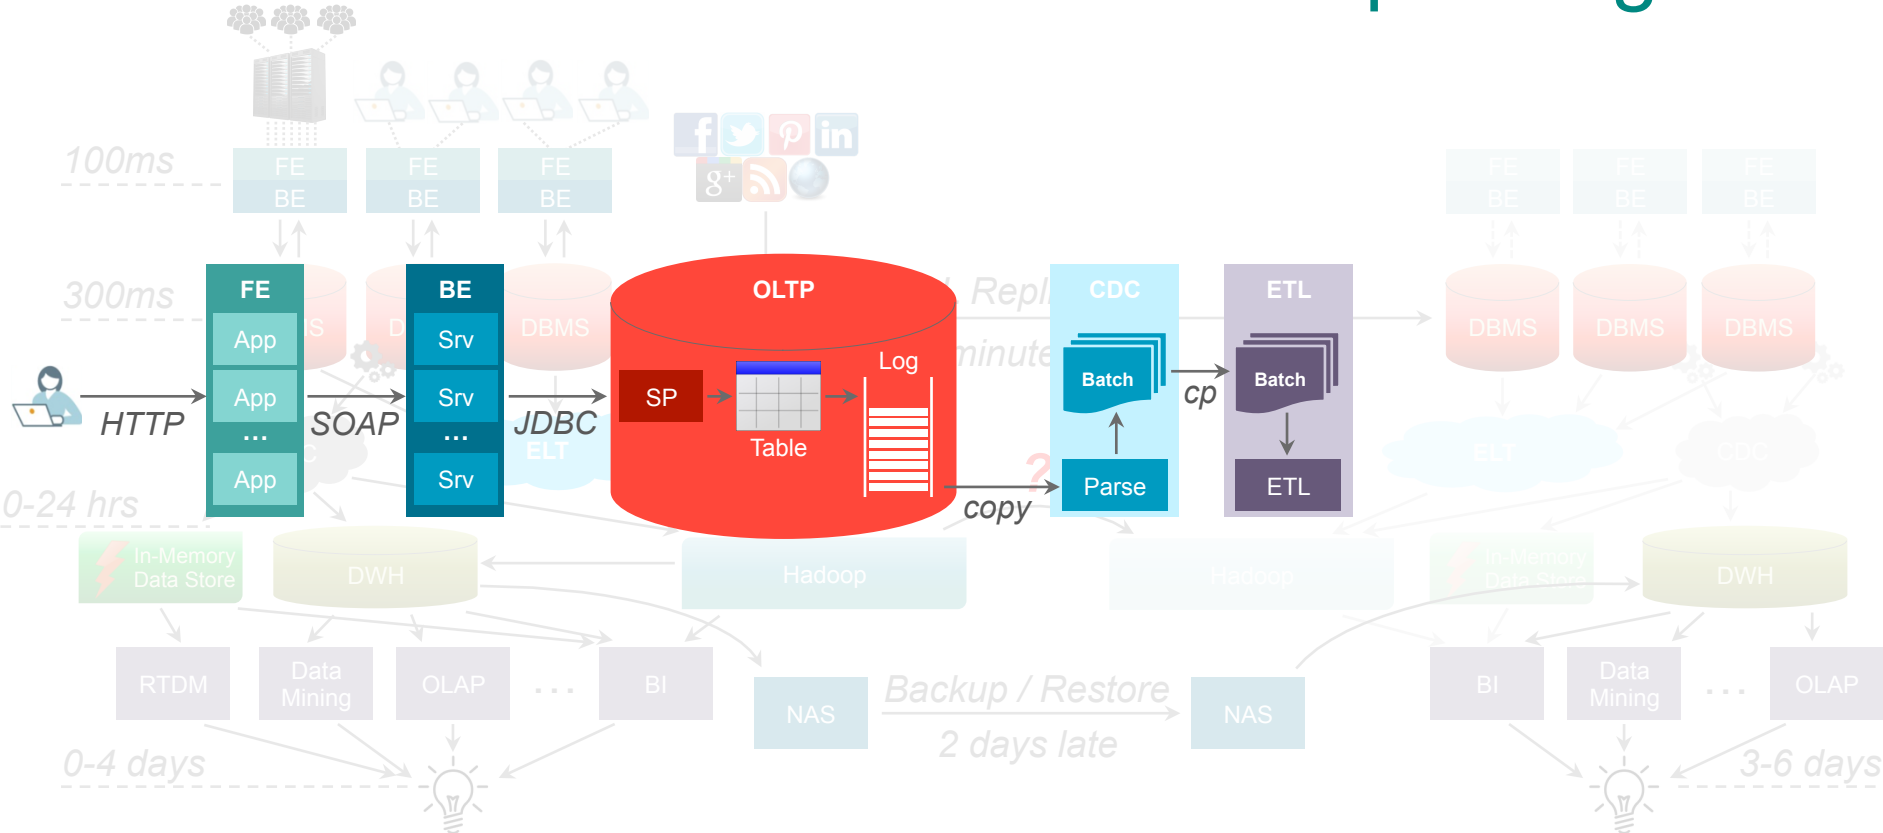




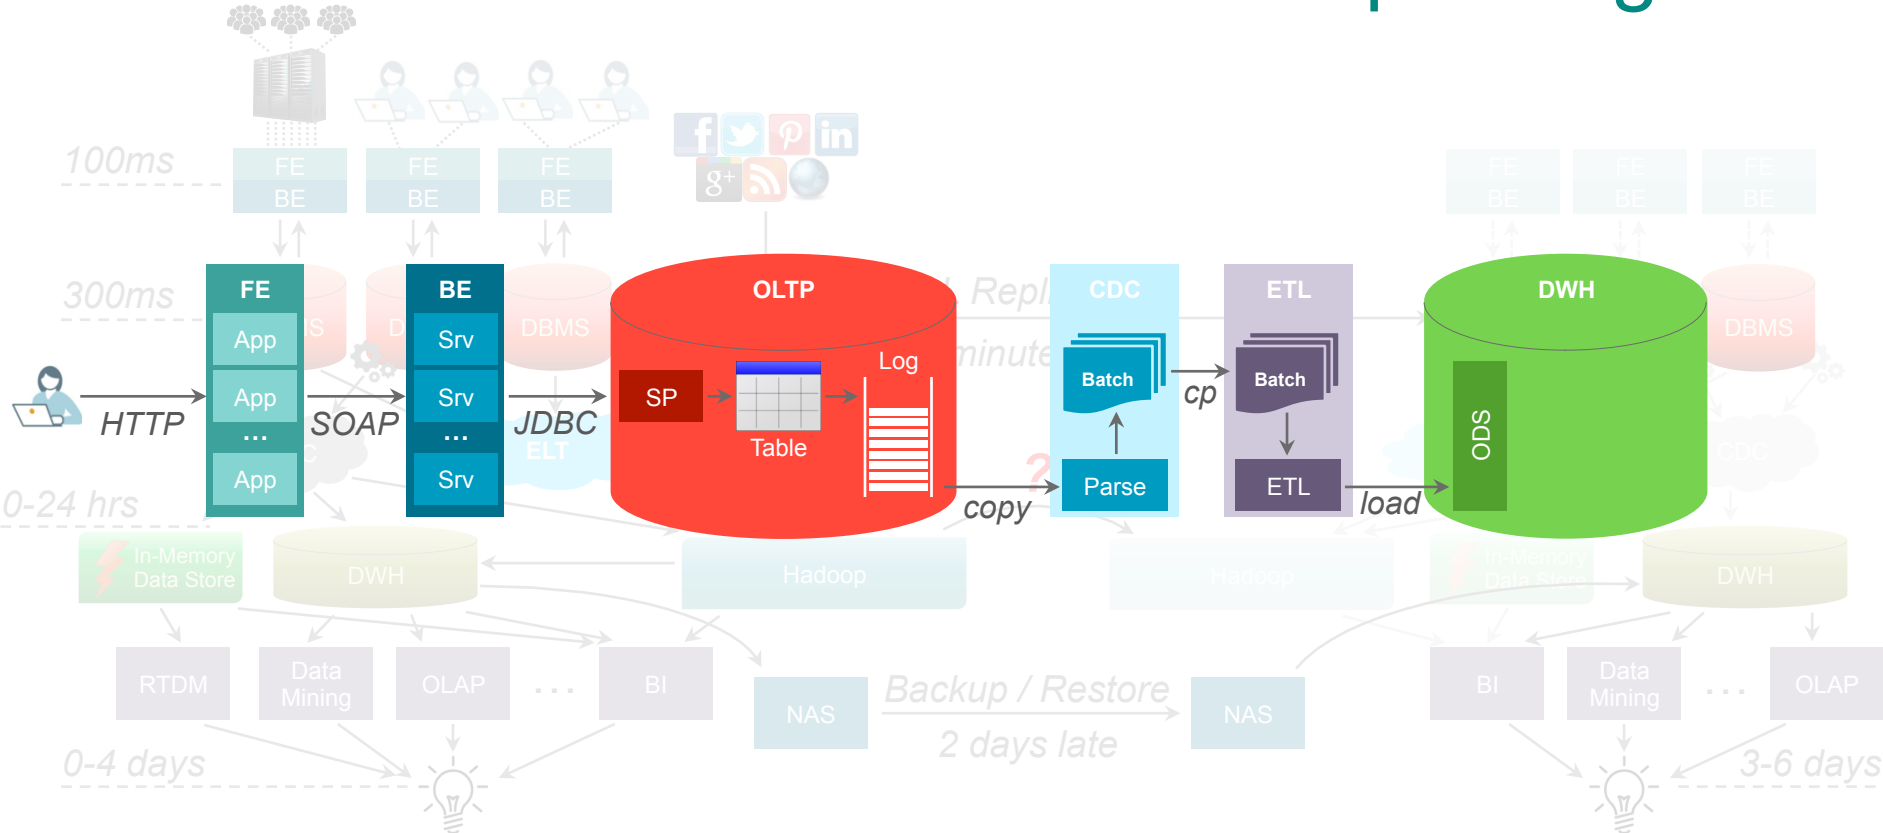
# Modern Data Architecture – Pipelining



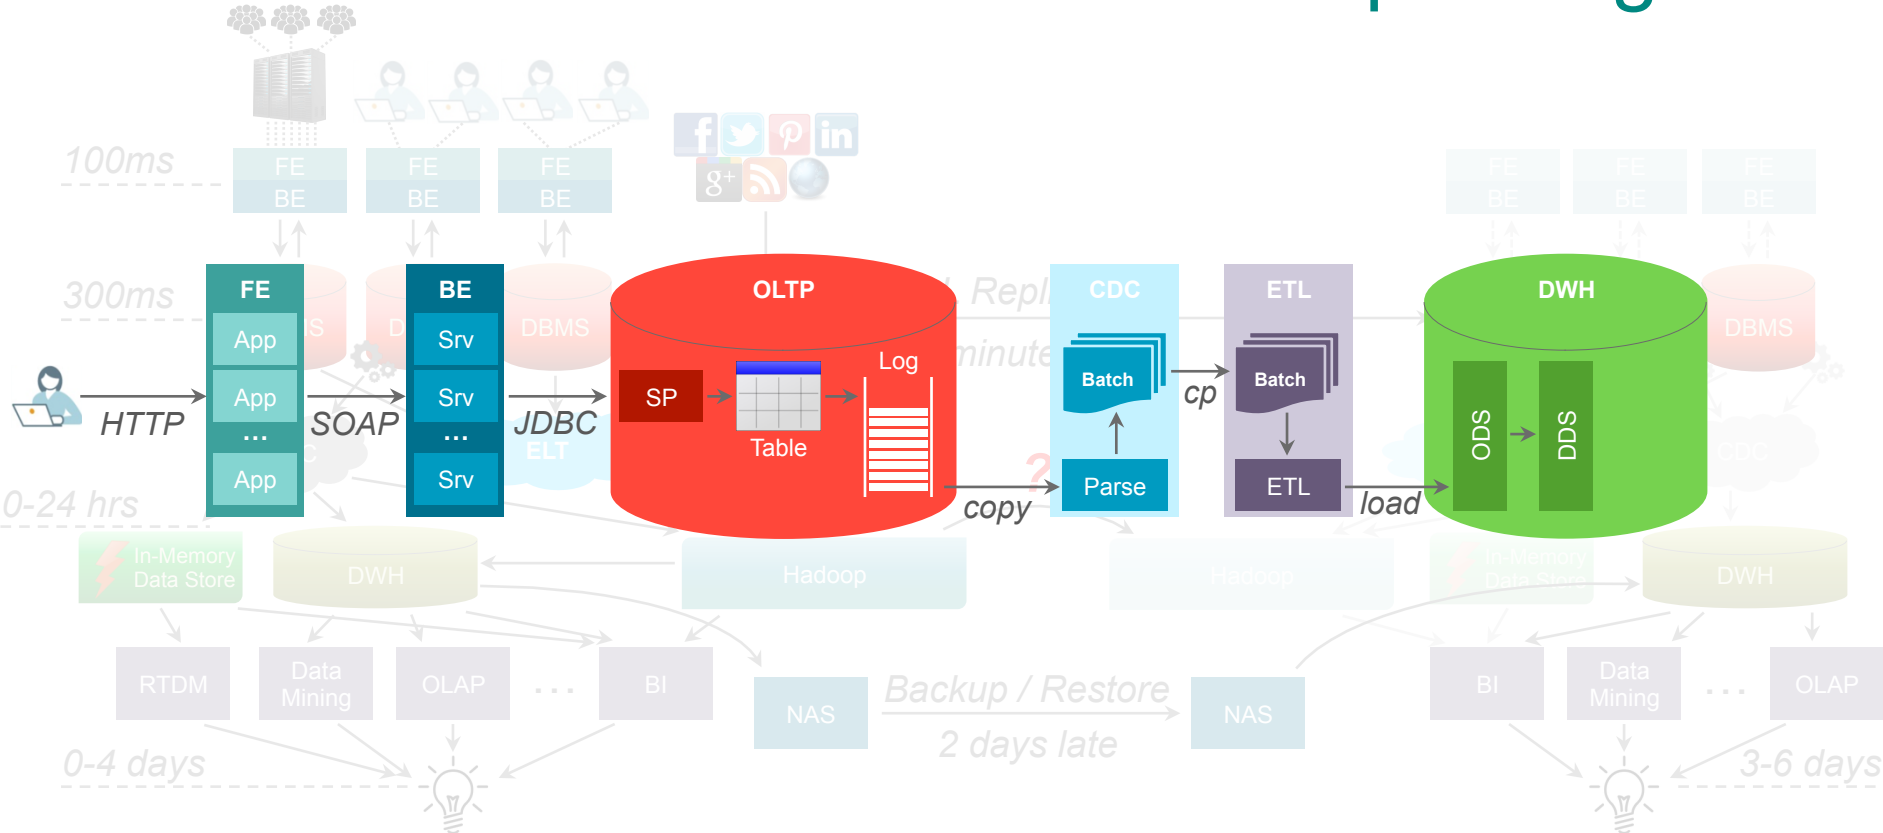
# Modern Data Architecture – Pipelining



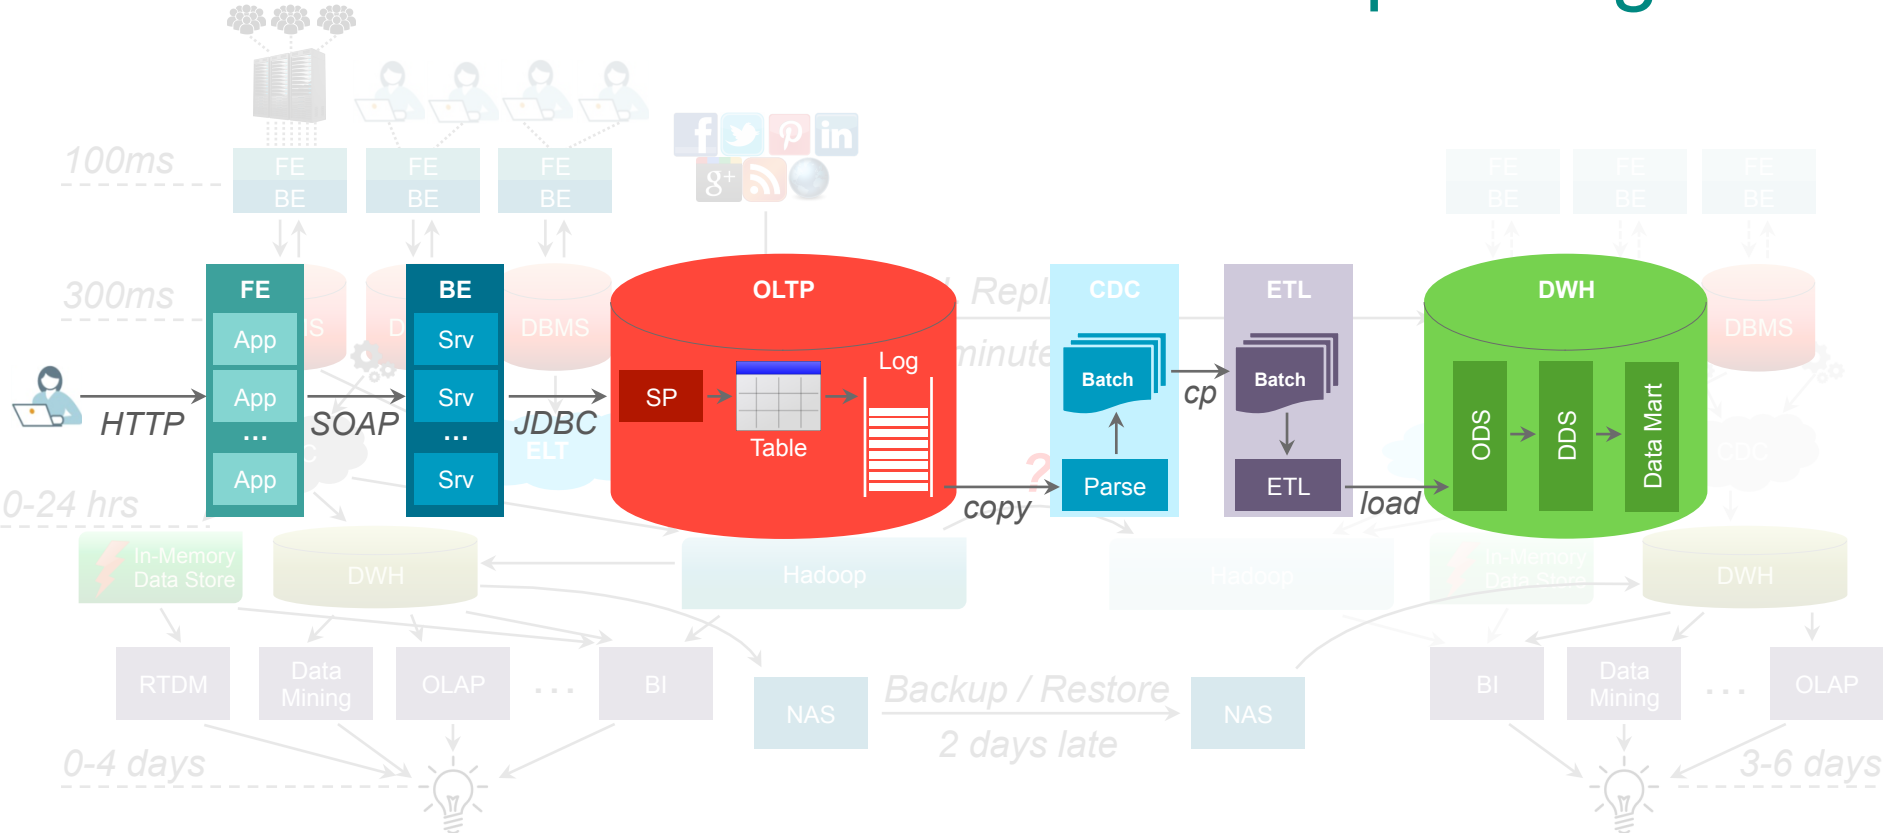
# Modern Data Architecture – Pipelining



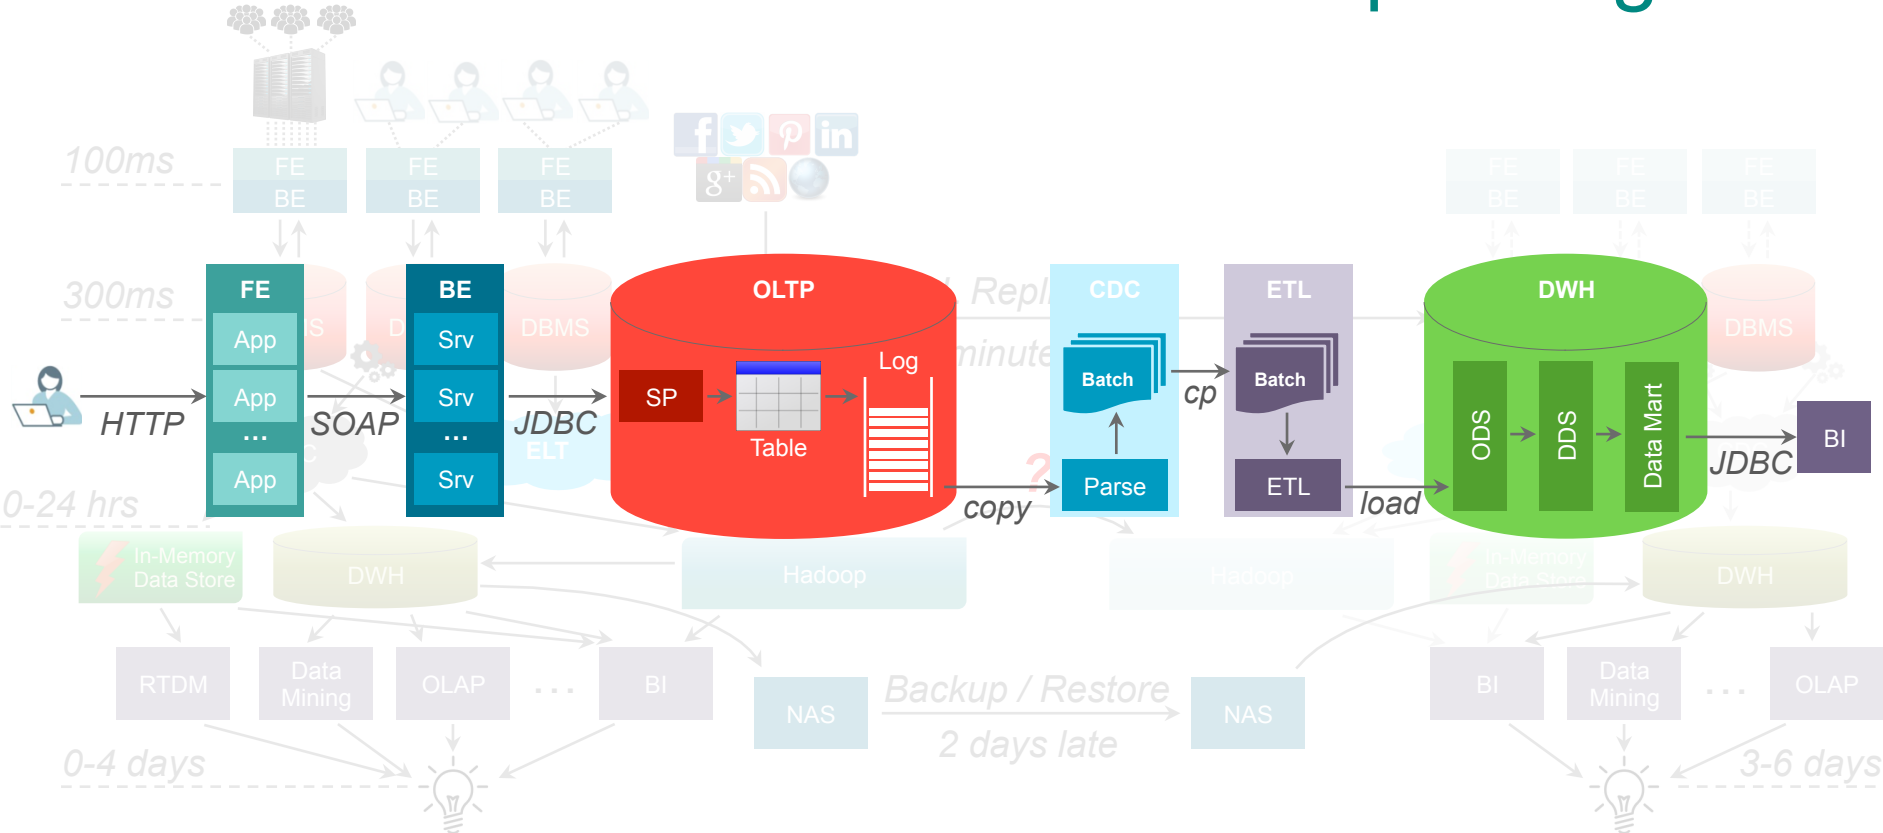
# Modern Data Architecture – Pipelining



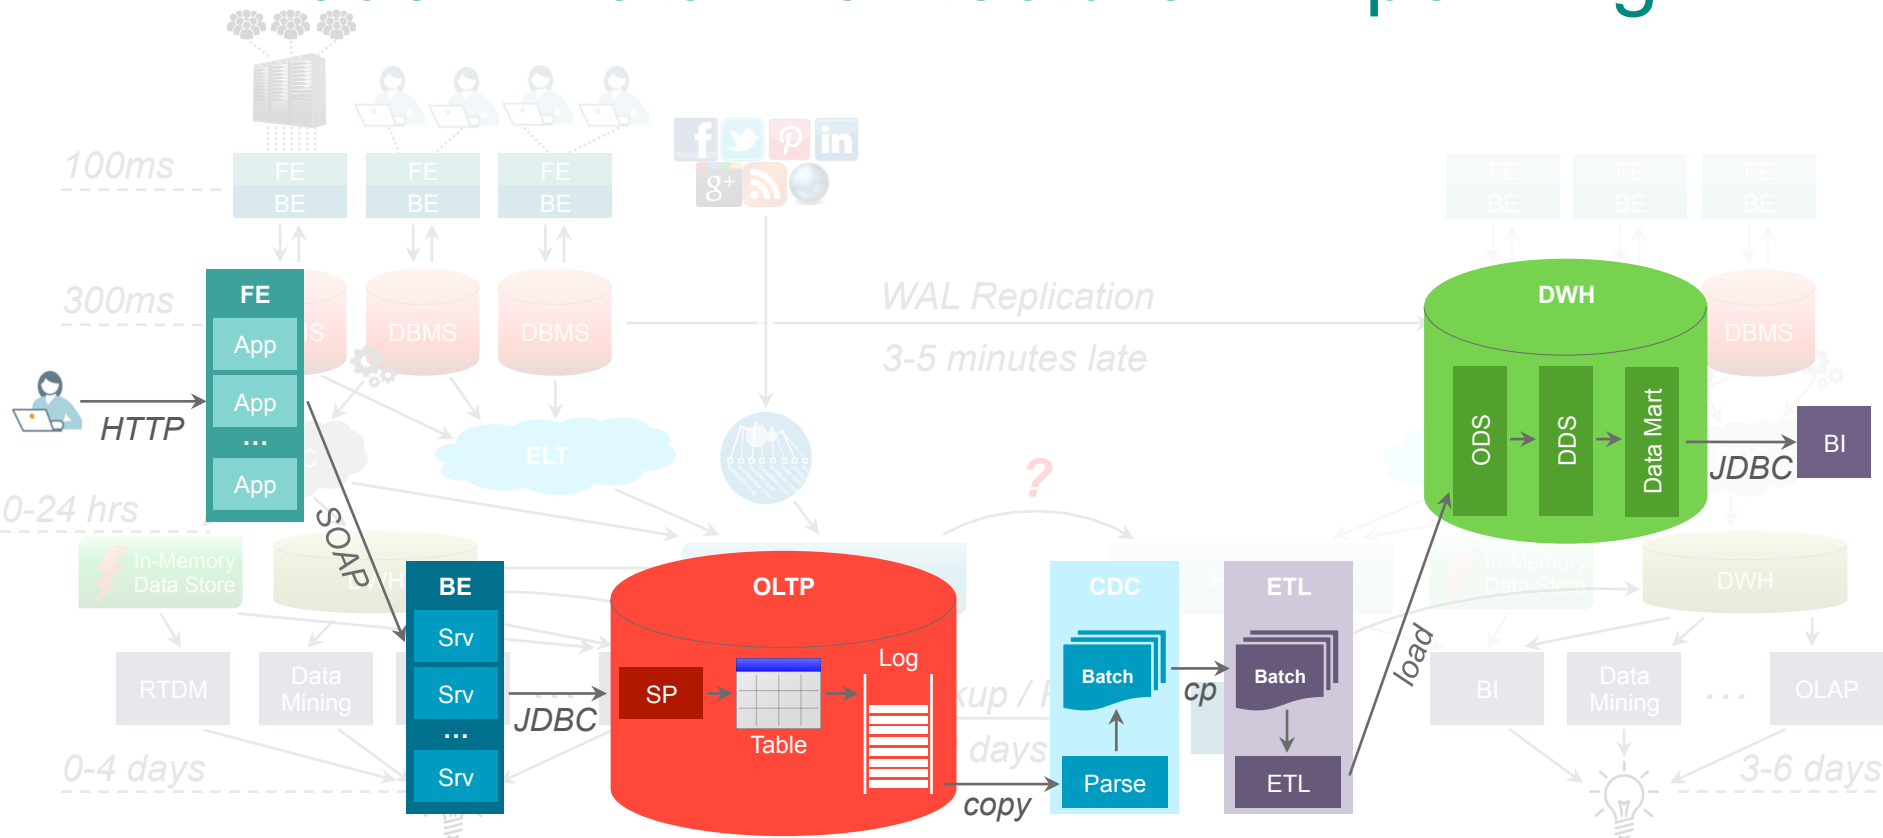
# Modern Data Architecture – Pipelining



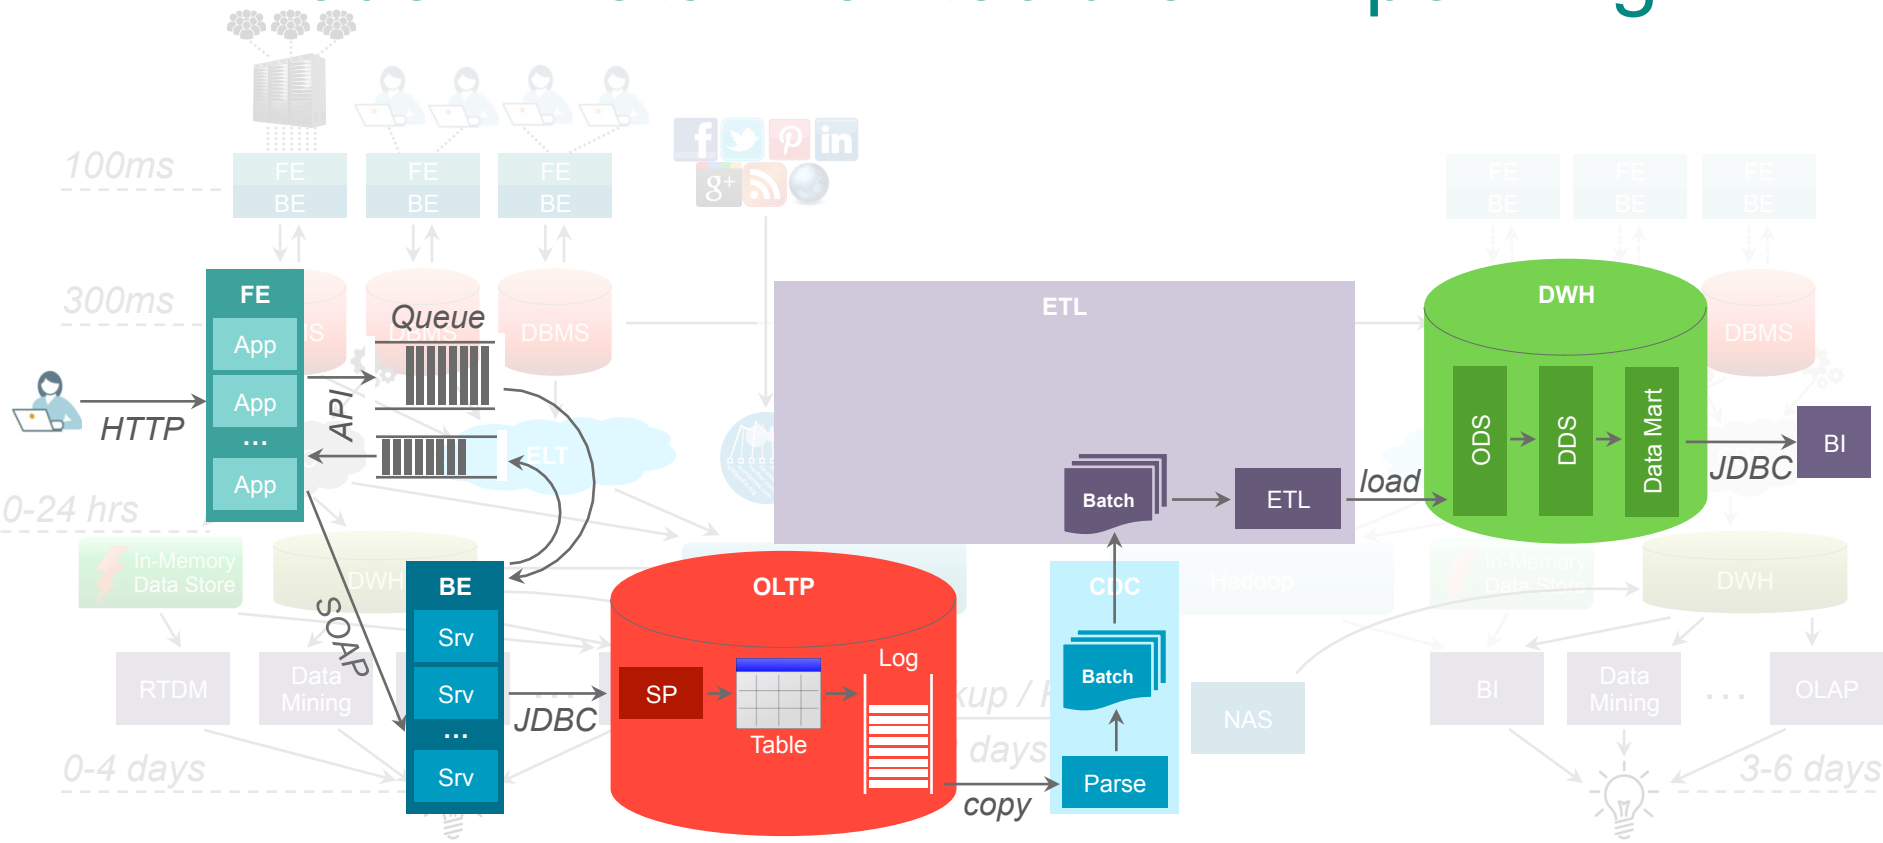
# Modern Data Architecture – Pipelining



# Modern Data Architecture – Pipelining

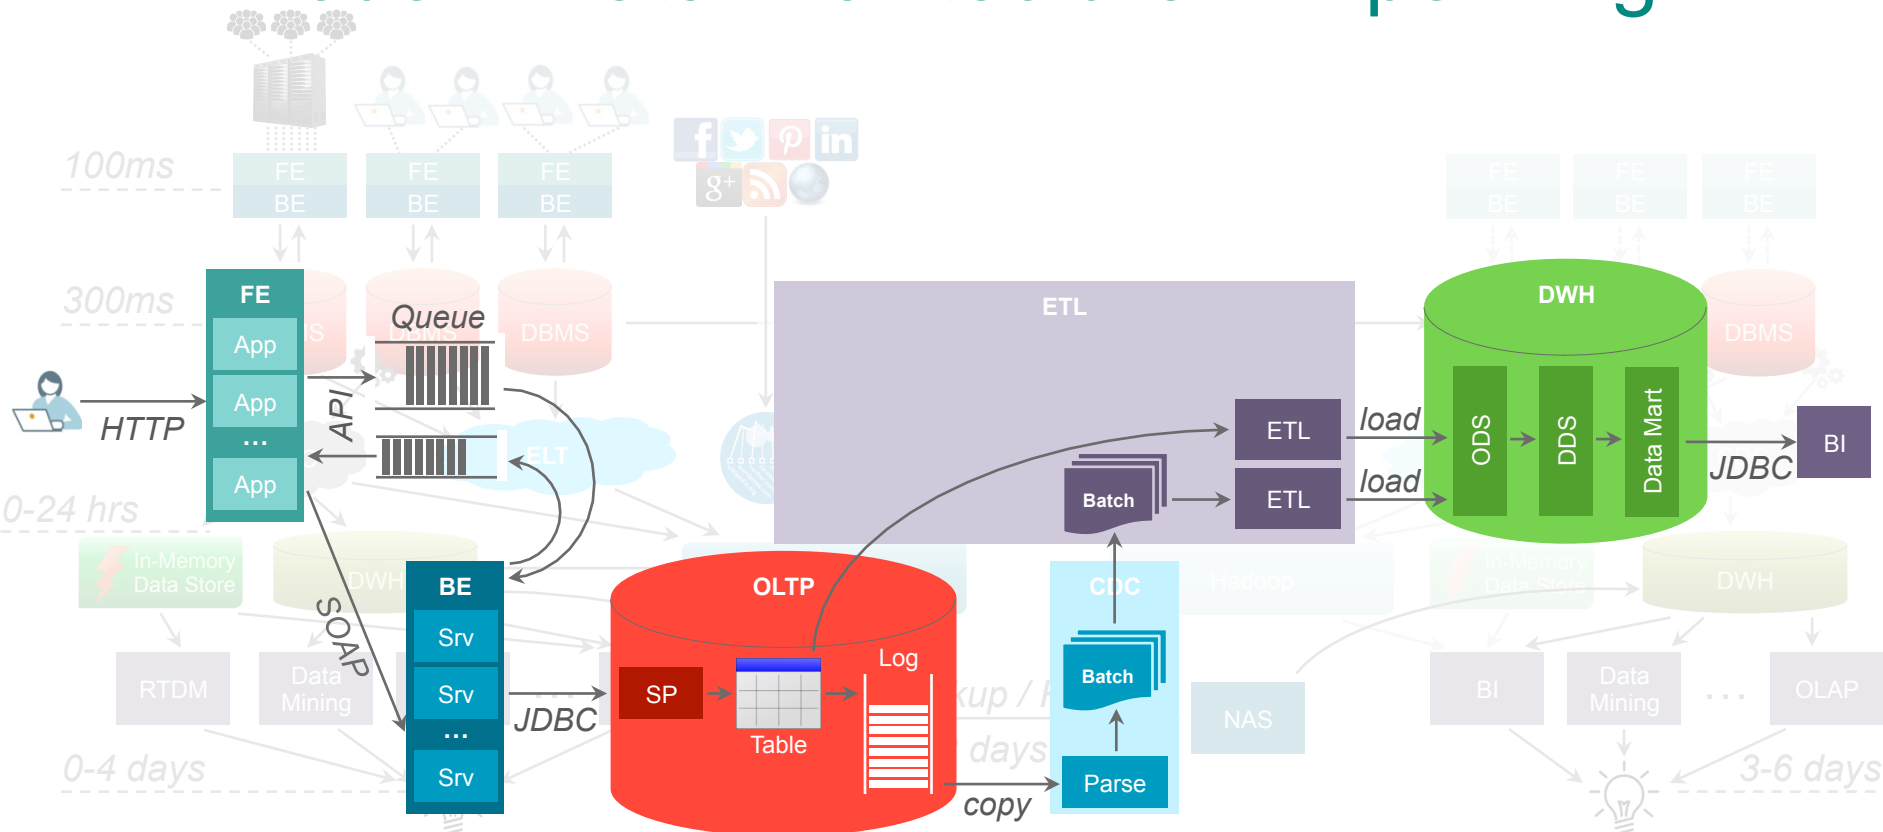


# Modern Data Architecture – Pipelining



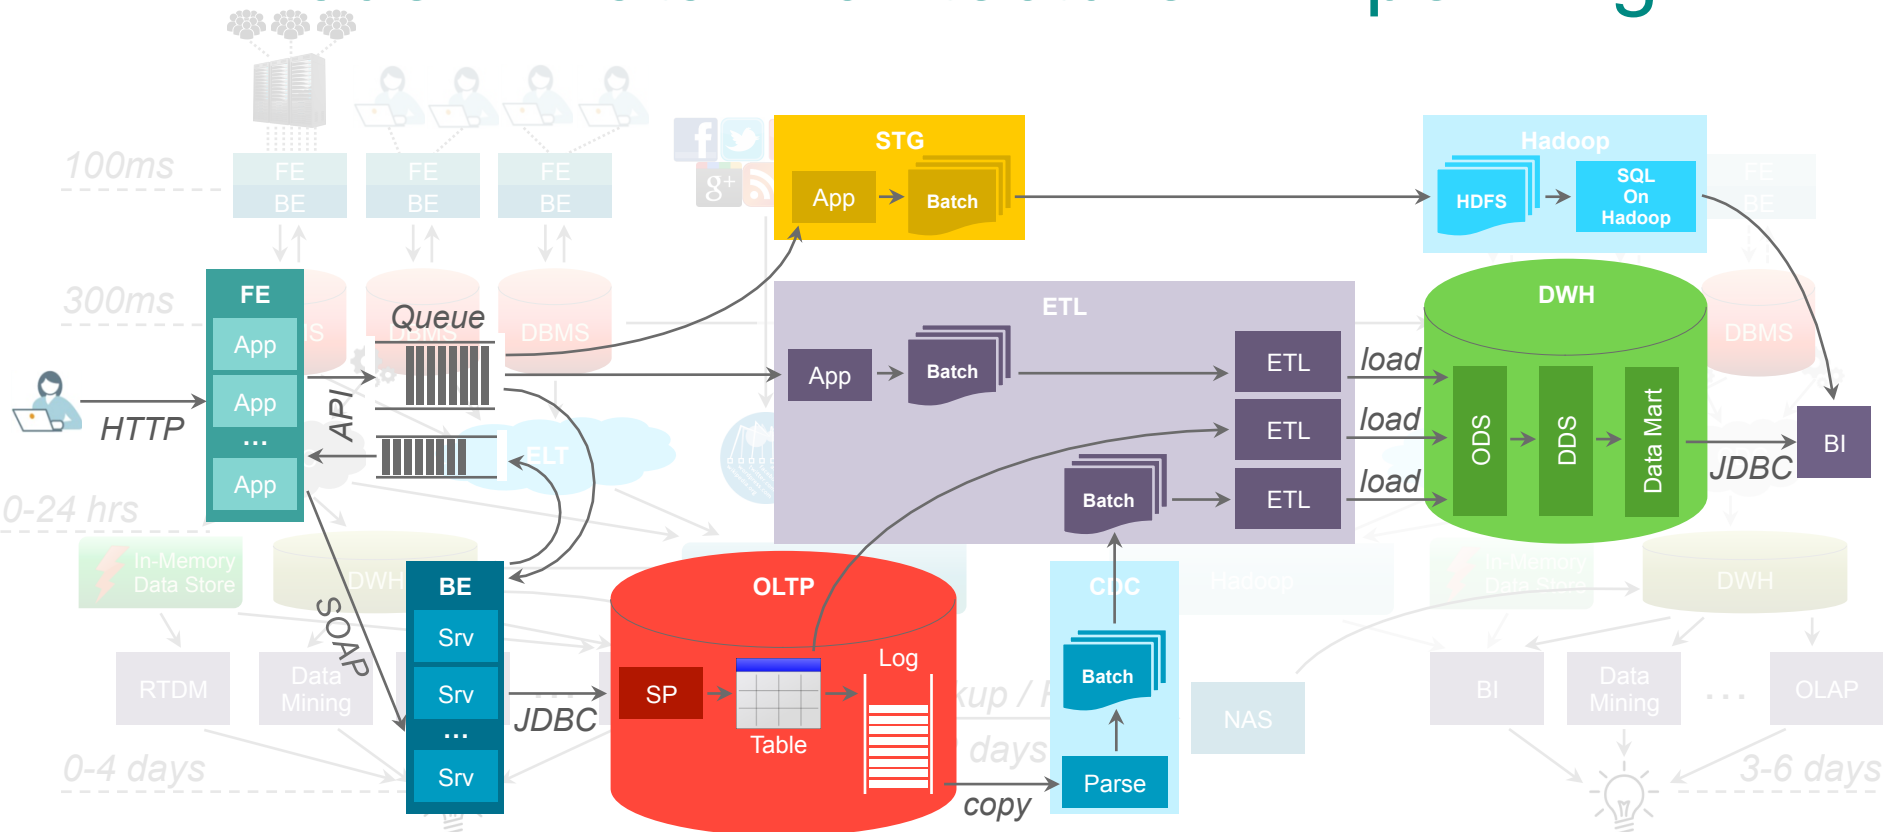


# Modern Data Architecture – Pipelining



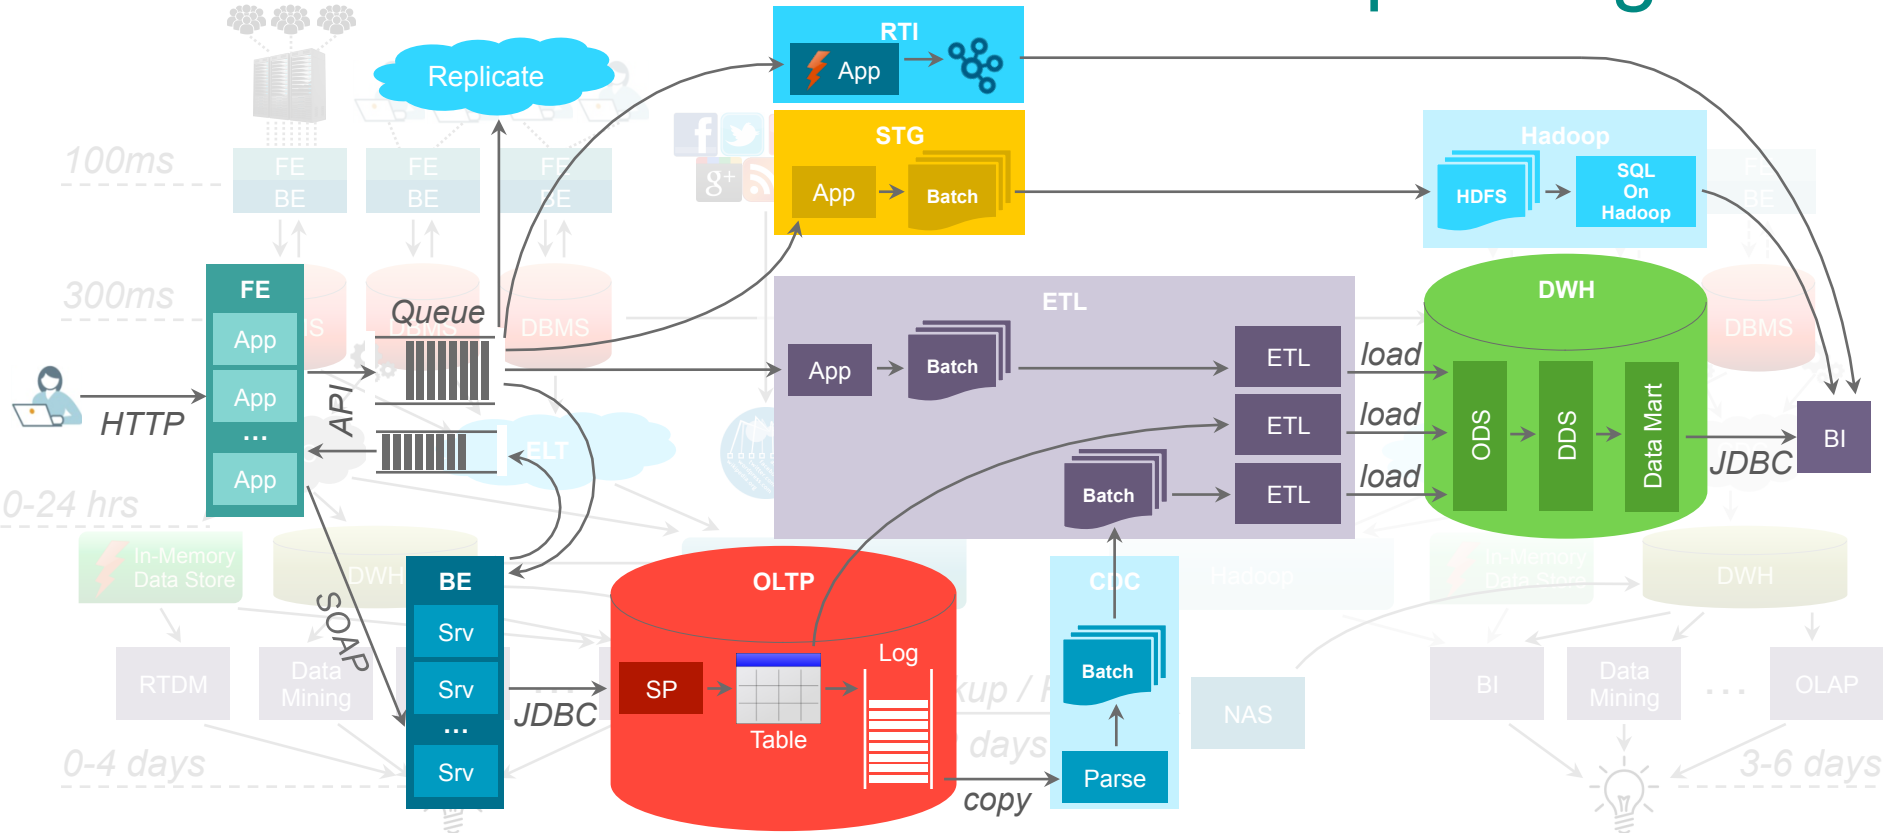


# Modern Data Architecture – Pipelining

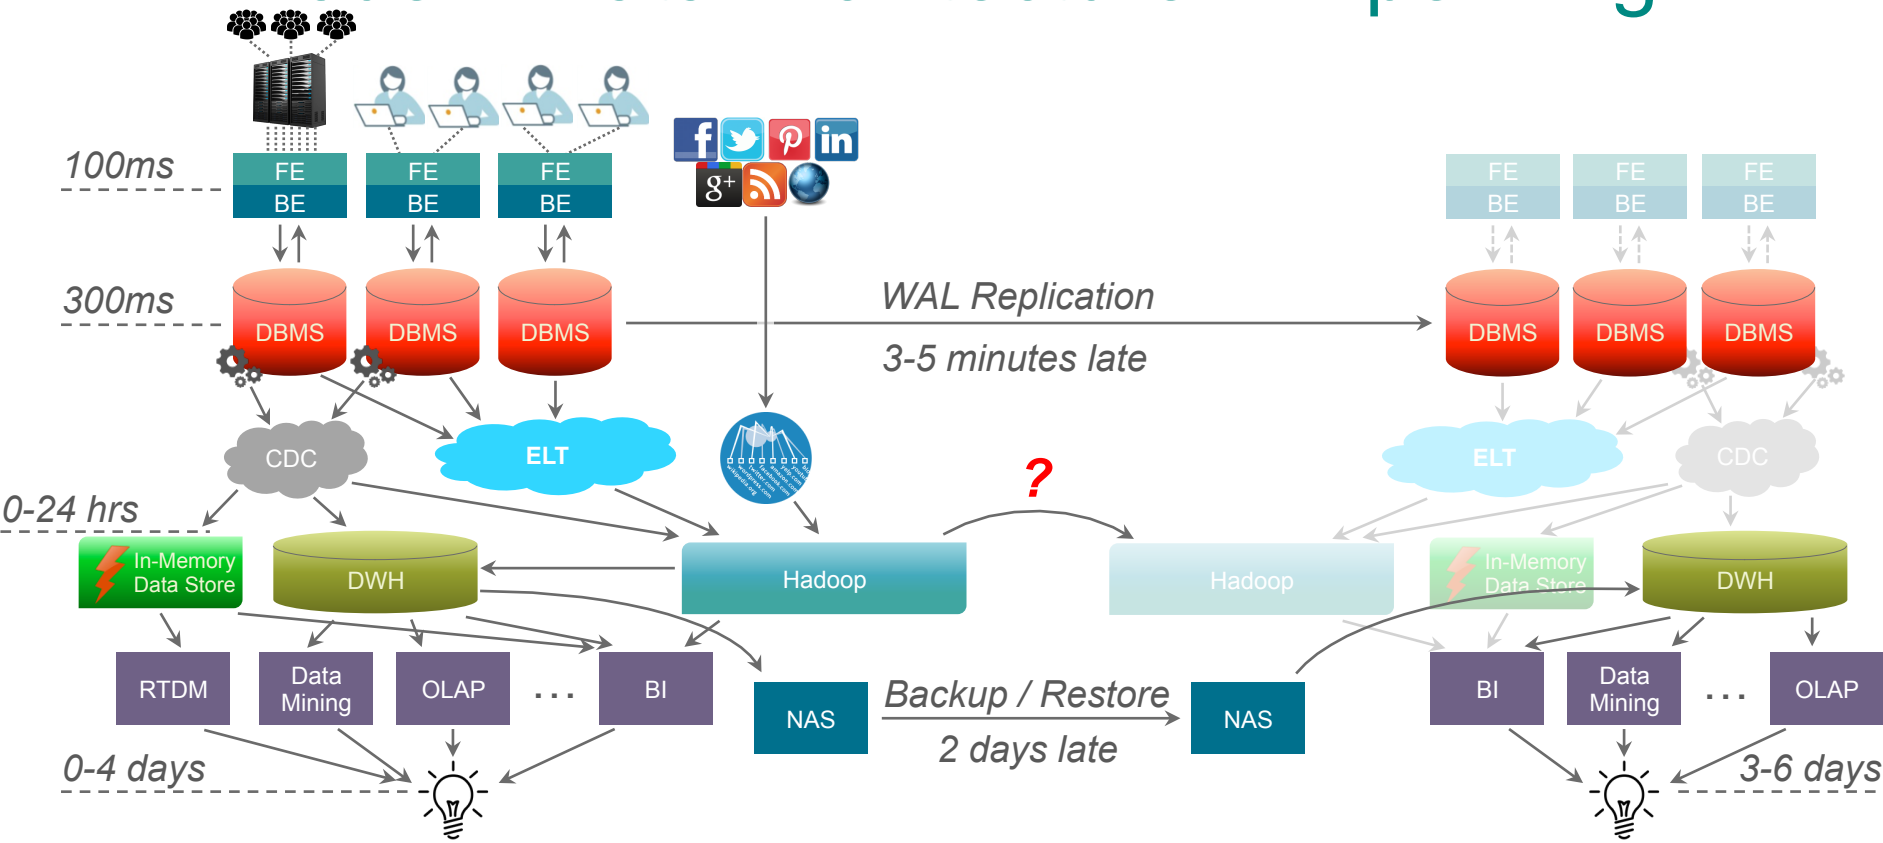




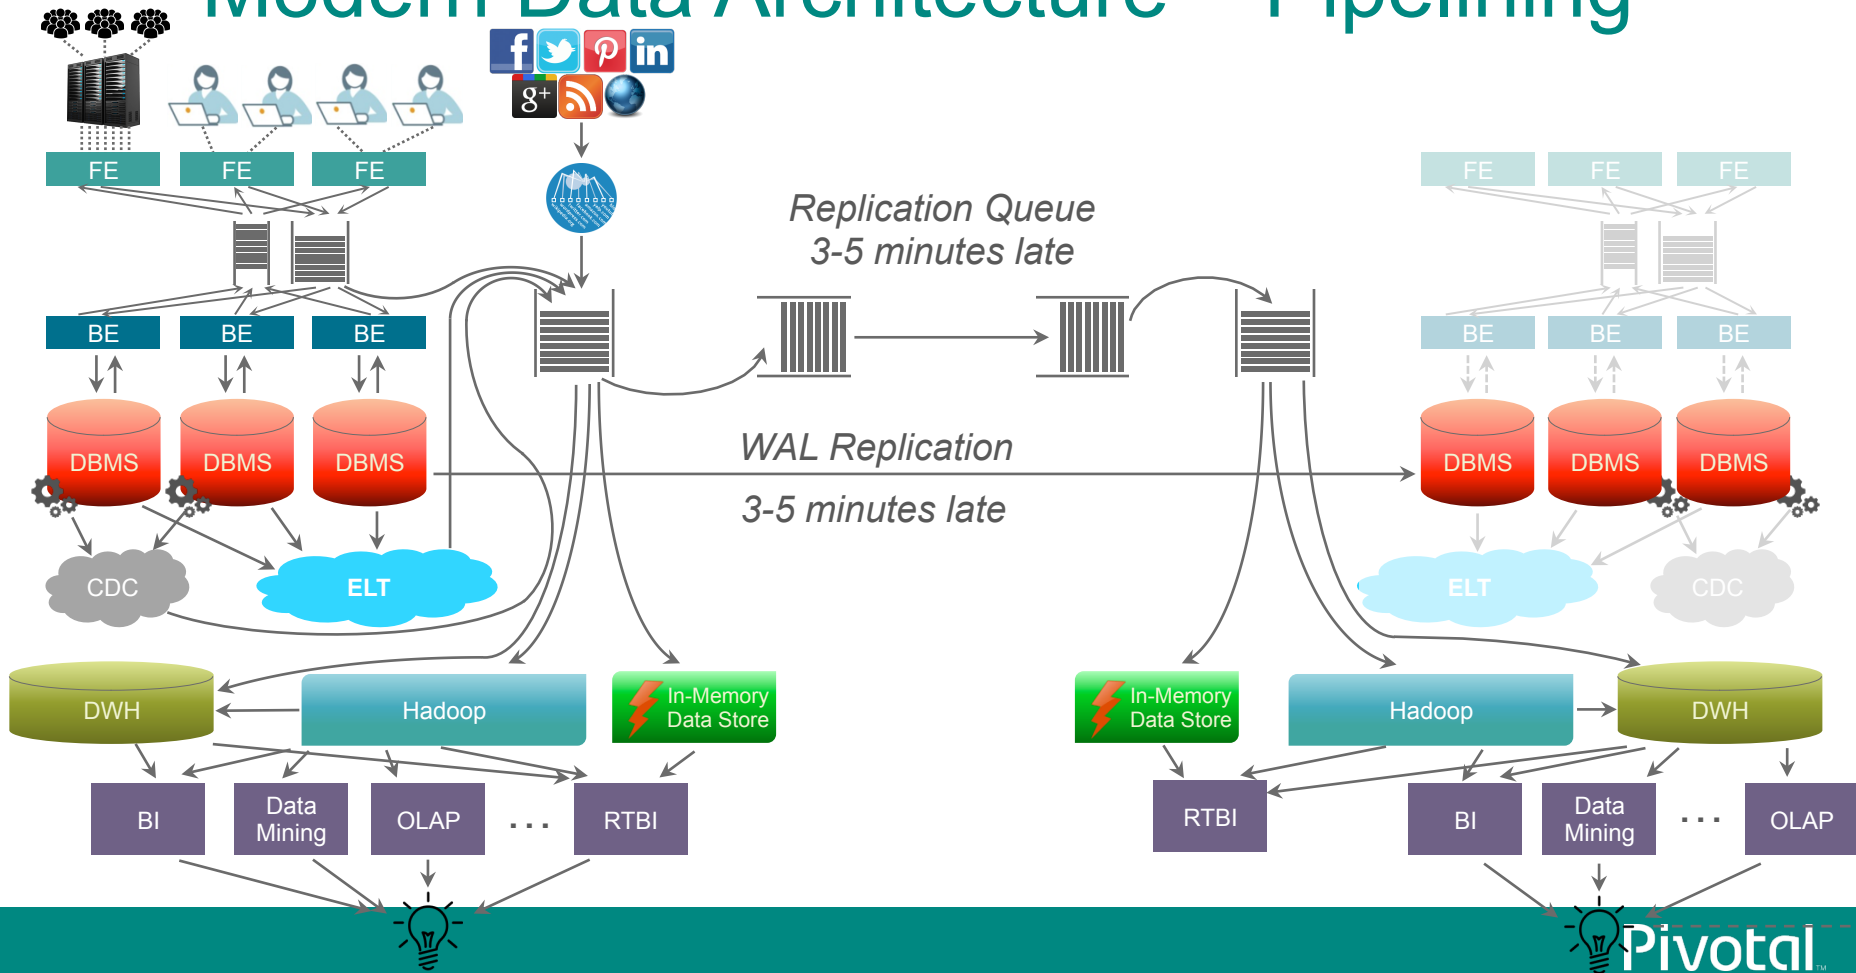
# Modern Data Architecture – Pipelining



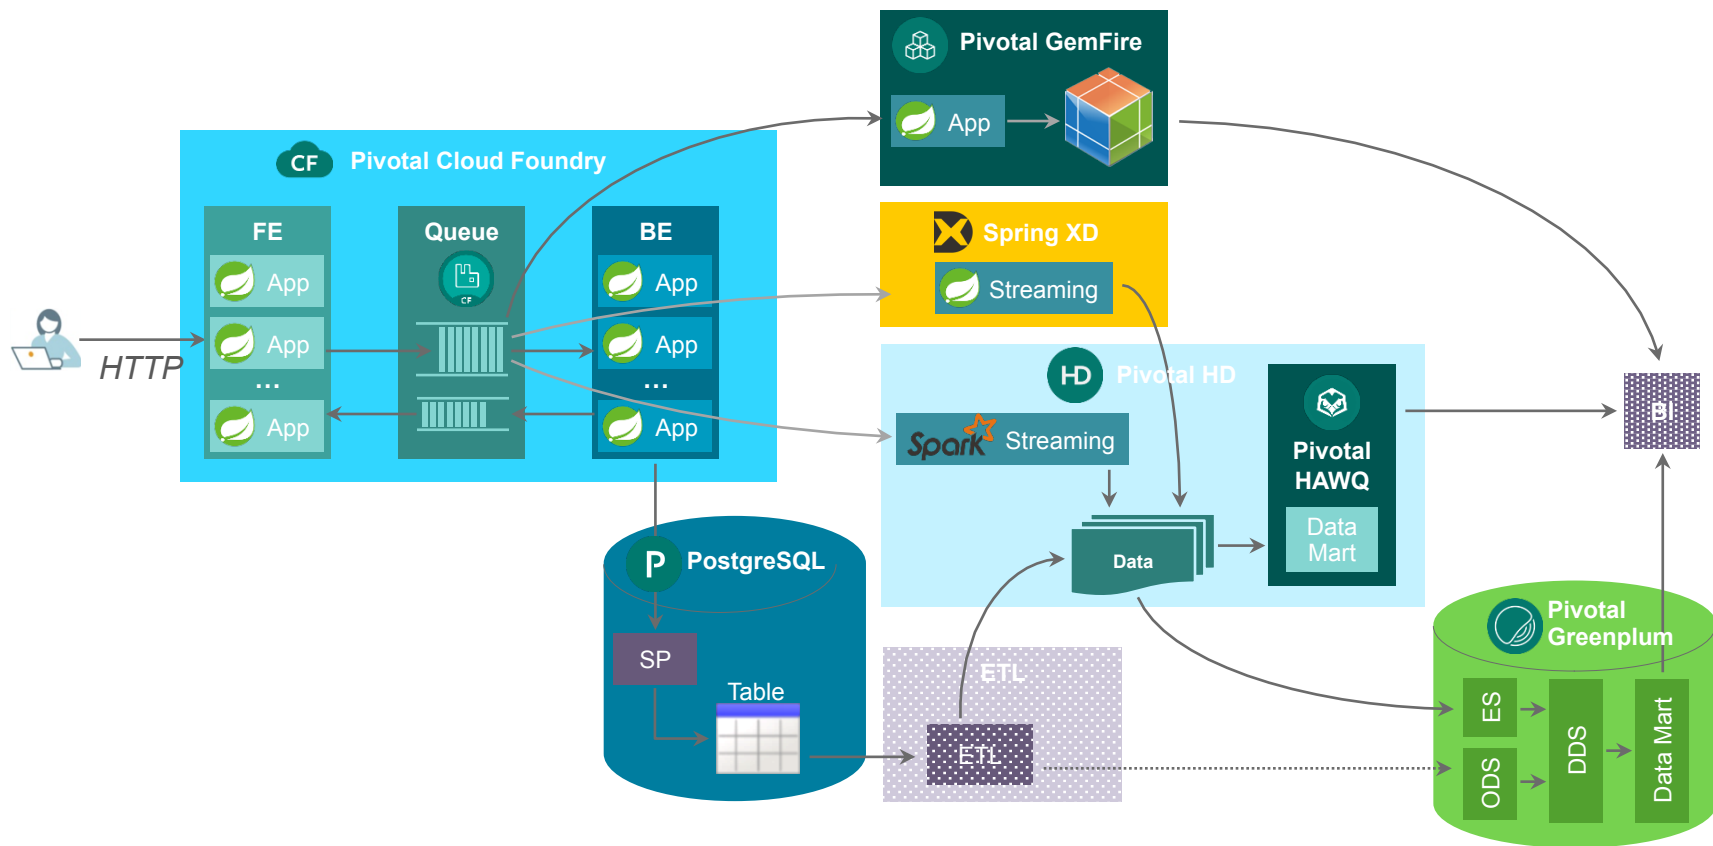
# Modern Data Architecture – Pipelining



# Modern Data Architecture – Pipelining

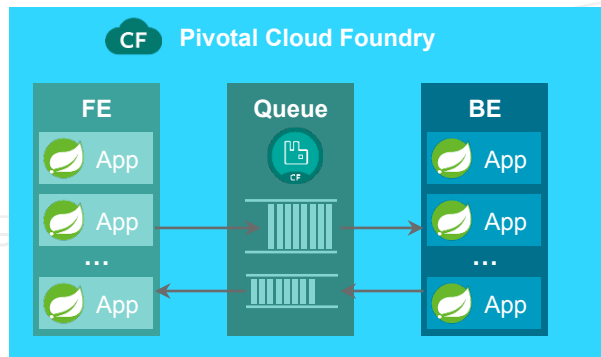


# Pivotal and Modern Data Architecture

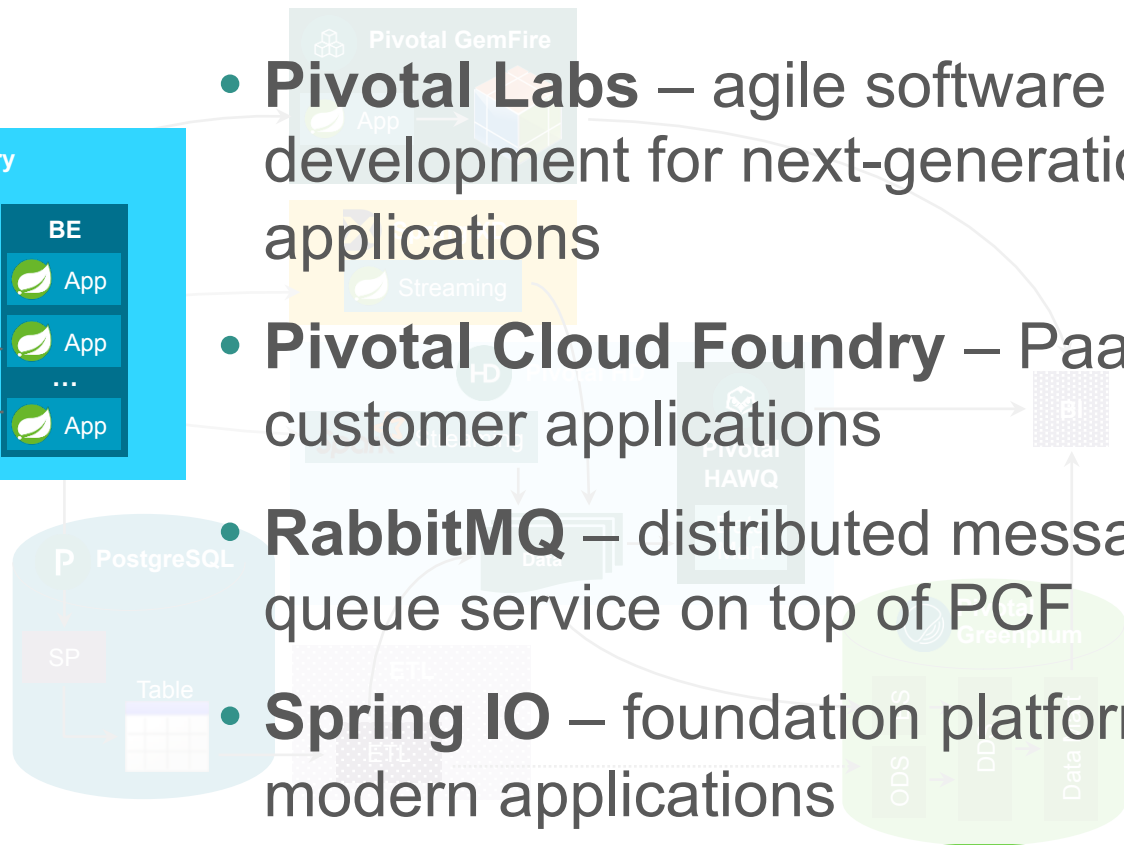




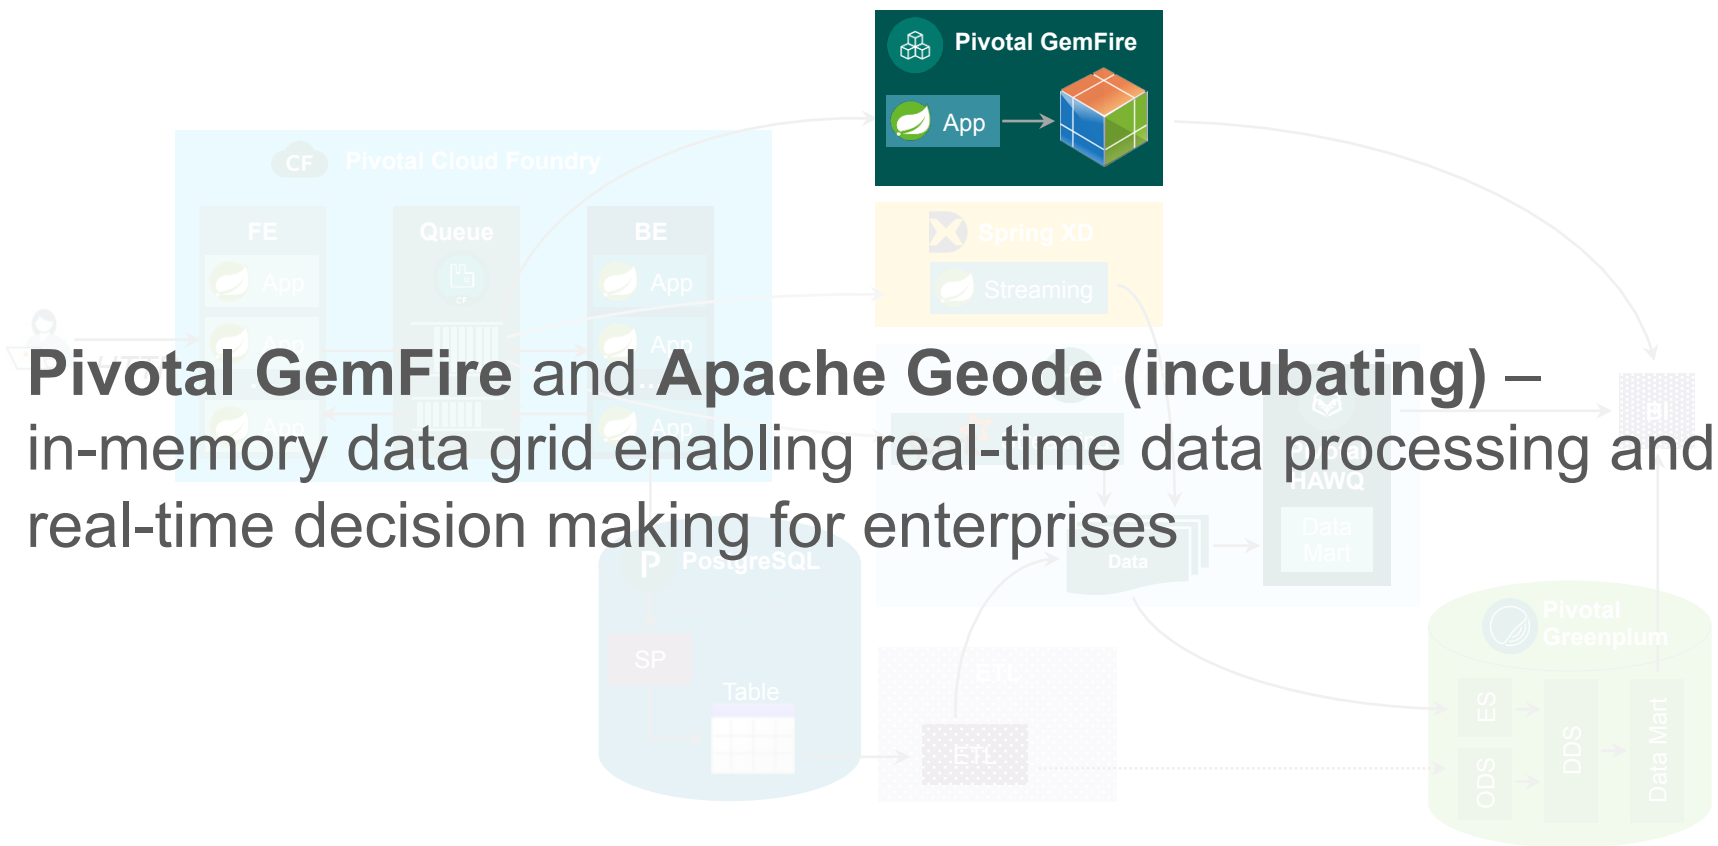
# Pivotal and Modern Data Architecture



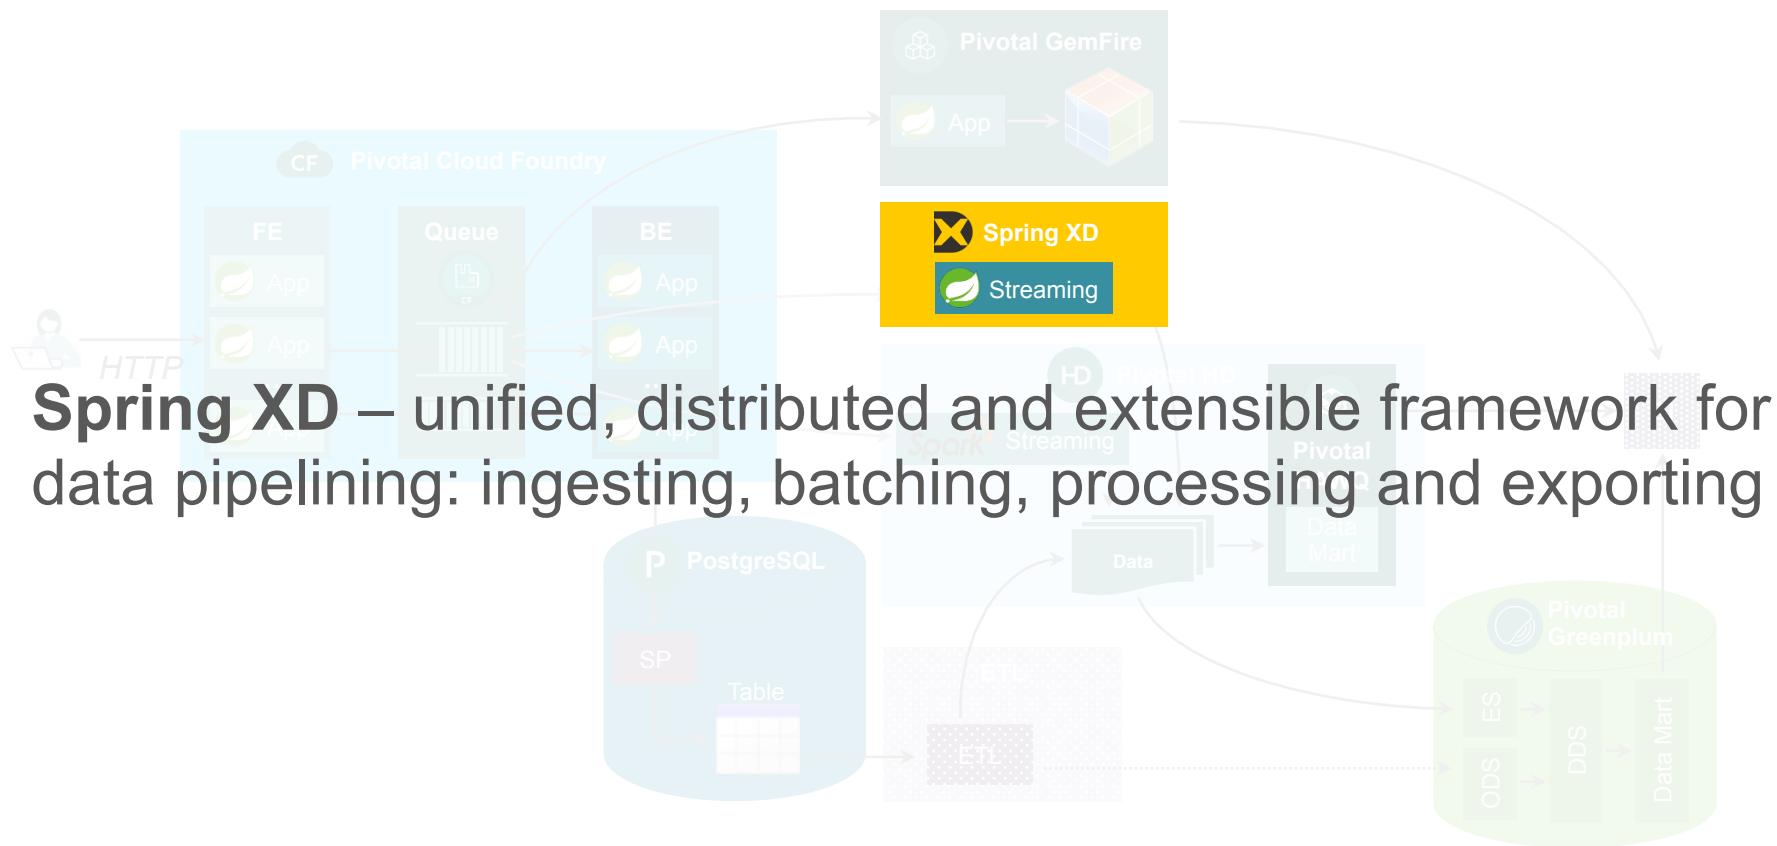
- **Pivotal Labs** – agile software development for next-generation applications
- **Pivotal Cloud Foundry** – PaaS for customer applications
- **RabbitMQ** – distributed message queue service on top of PCF
- **Spring IO** – foundation platform for modern applications



# Pivotal and Modern Data Architecture

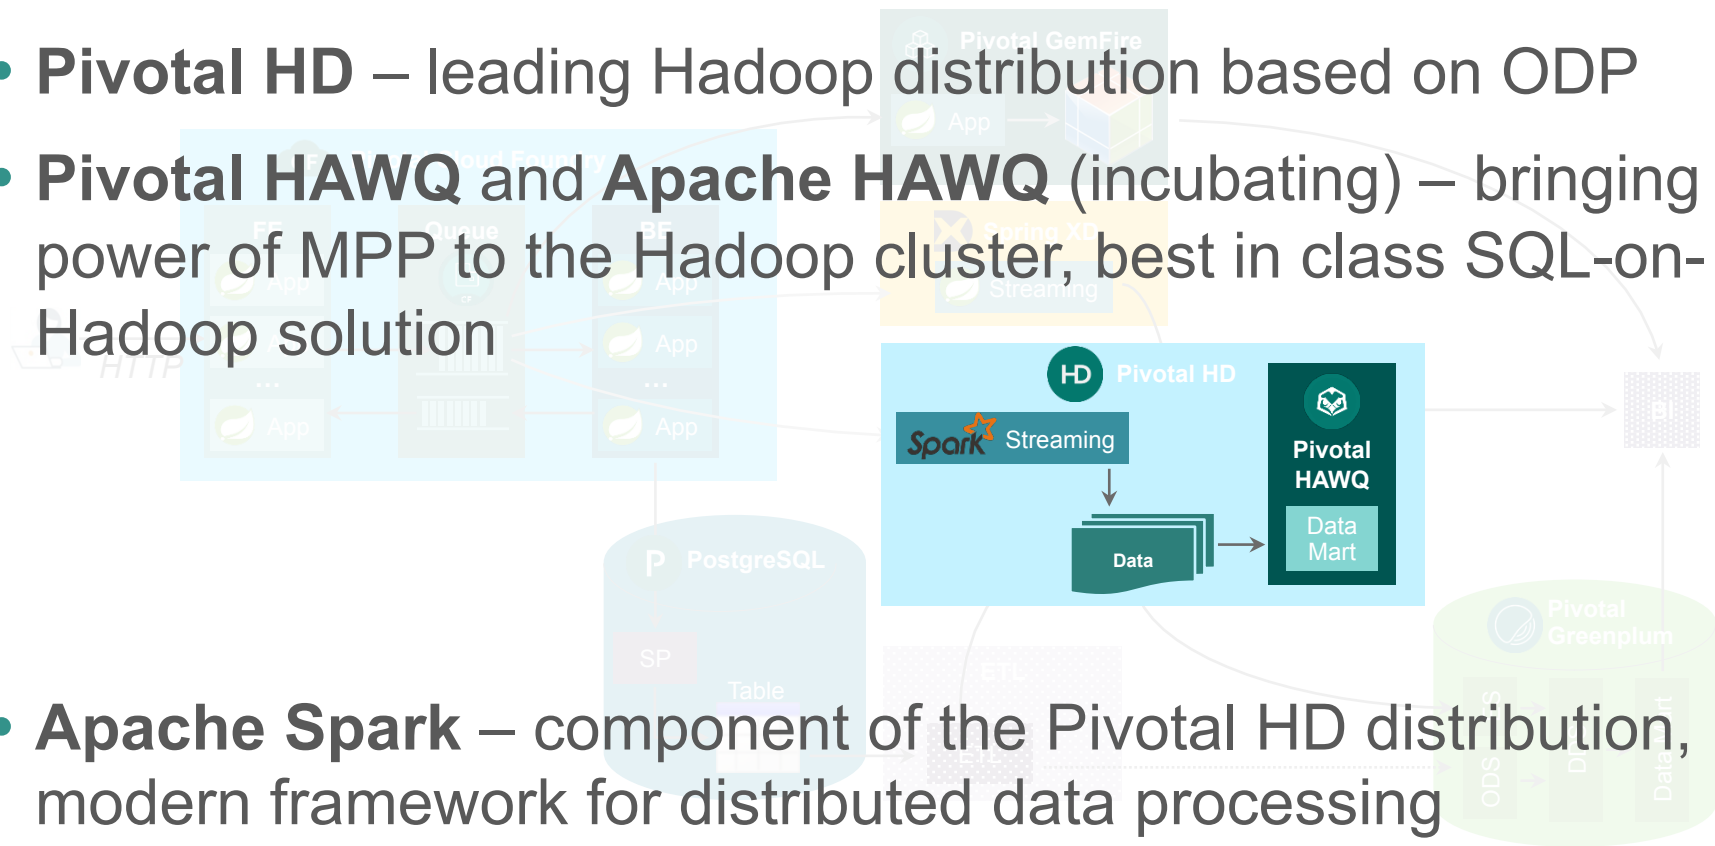


# Pivotal and Modern Data Architecture



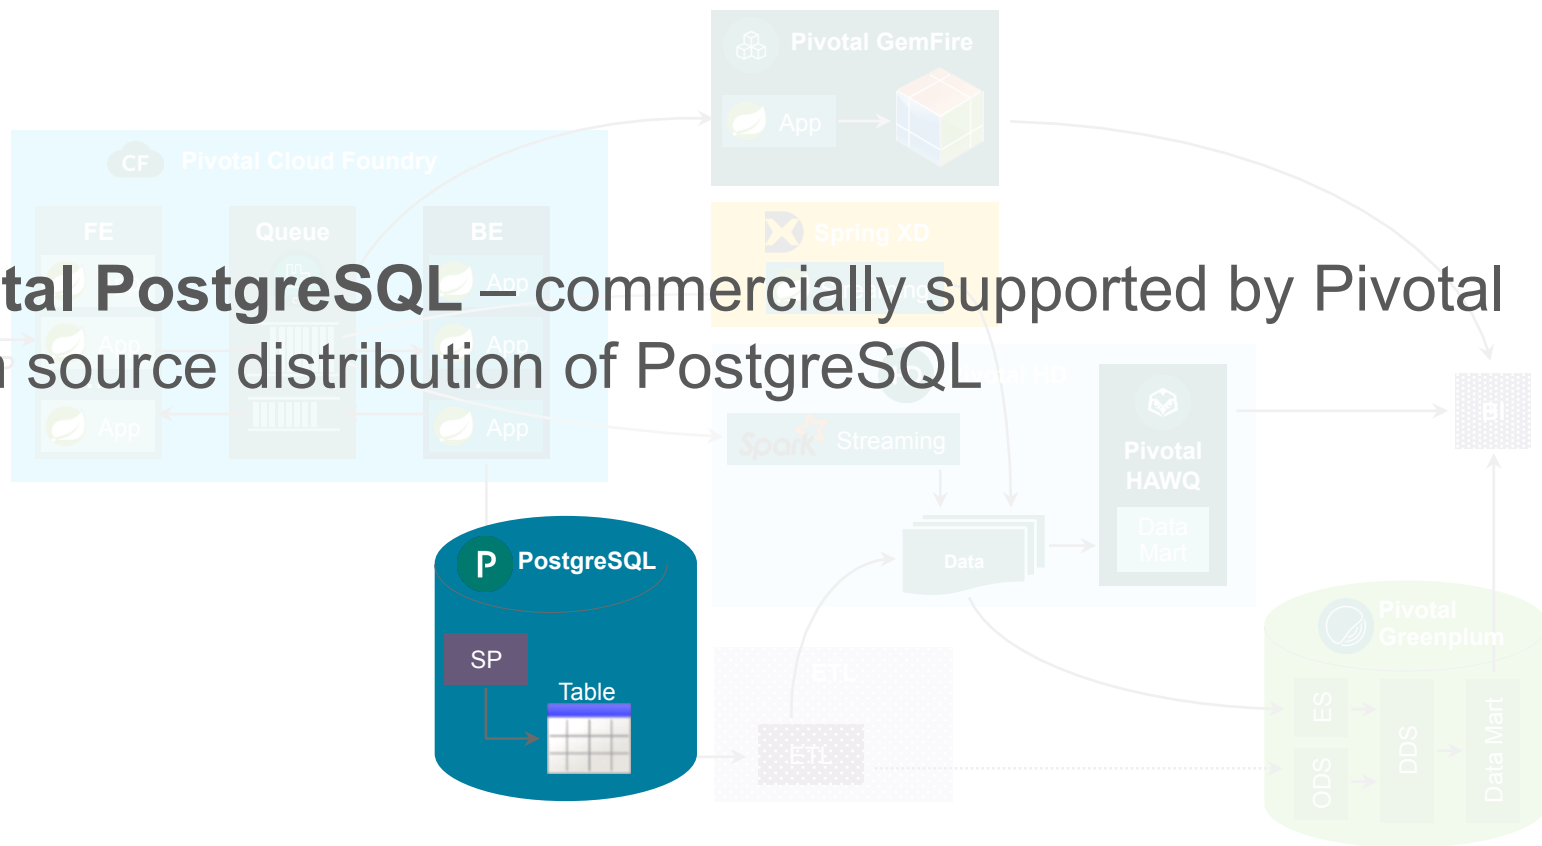
# Pivotal and Modern Data Architecture

- **Pivotal HD** – leading Hadoop distribution based on ODP
- **Pivotal HAWQ** and **Apache HAWQ** (incubating) – bringing the power of MPP to the Hadoop cluster, best in class SQL-on-Hadoop solution
- **Apache Spark** – component of the Pivotal HD distribution, modern framework for distributed data processing



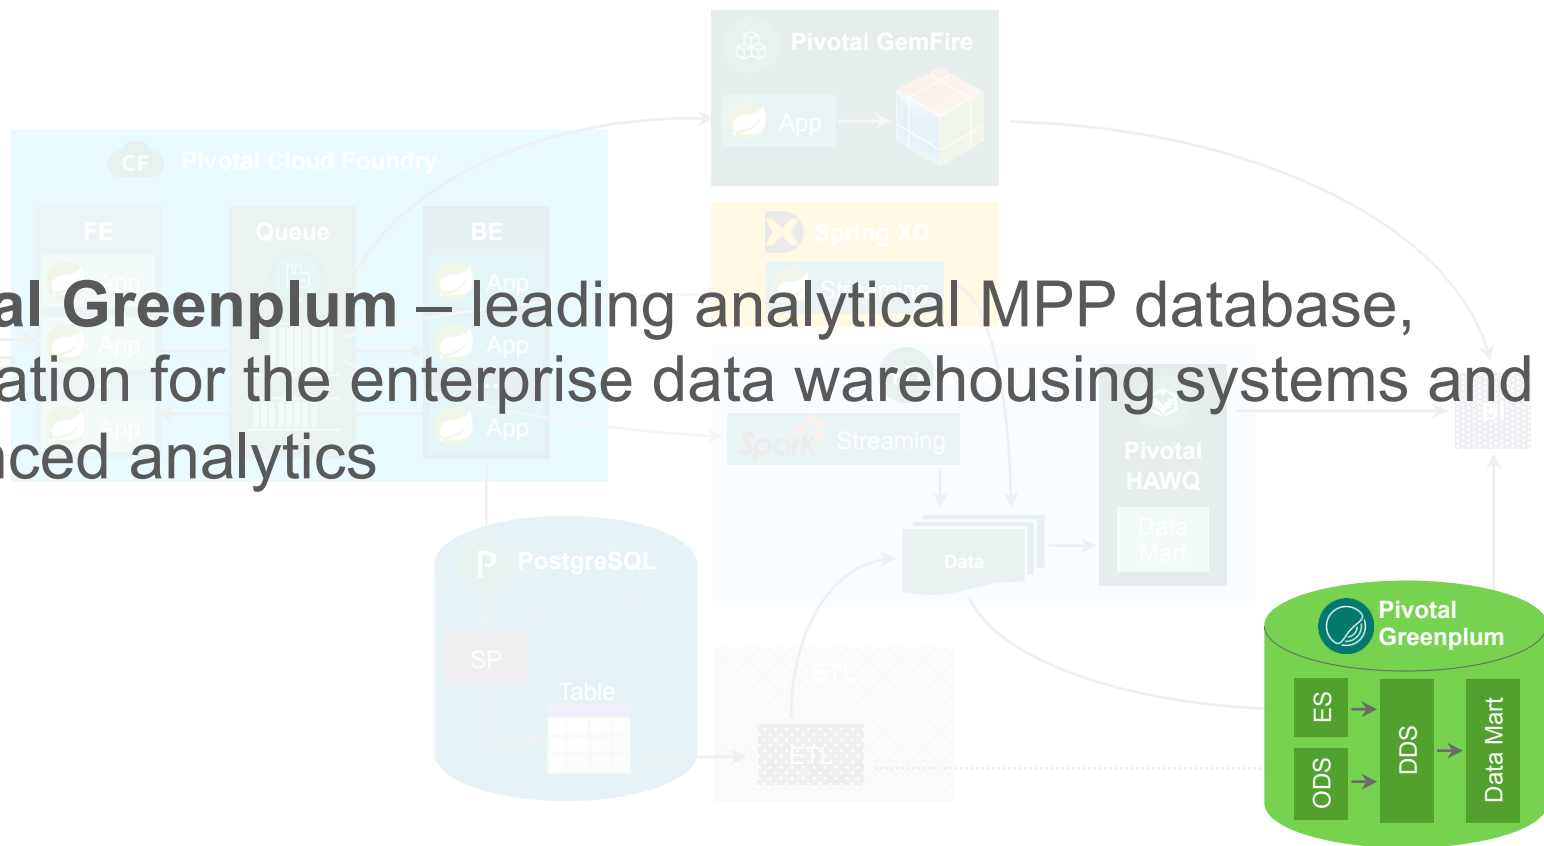
# Pivotal and Modern Data Architecture

- **Pivotal PostgreSQL** – commercially supported by Pivotal open source distribution of PostgreSQL

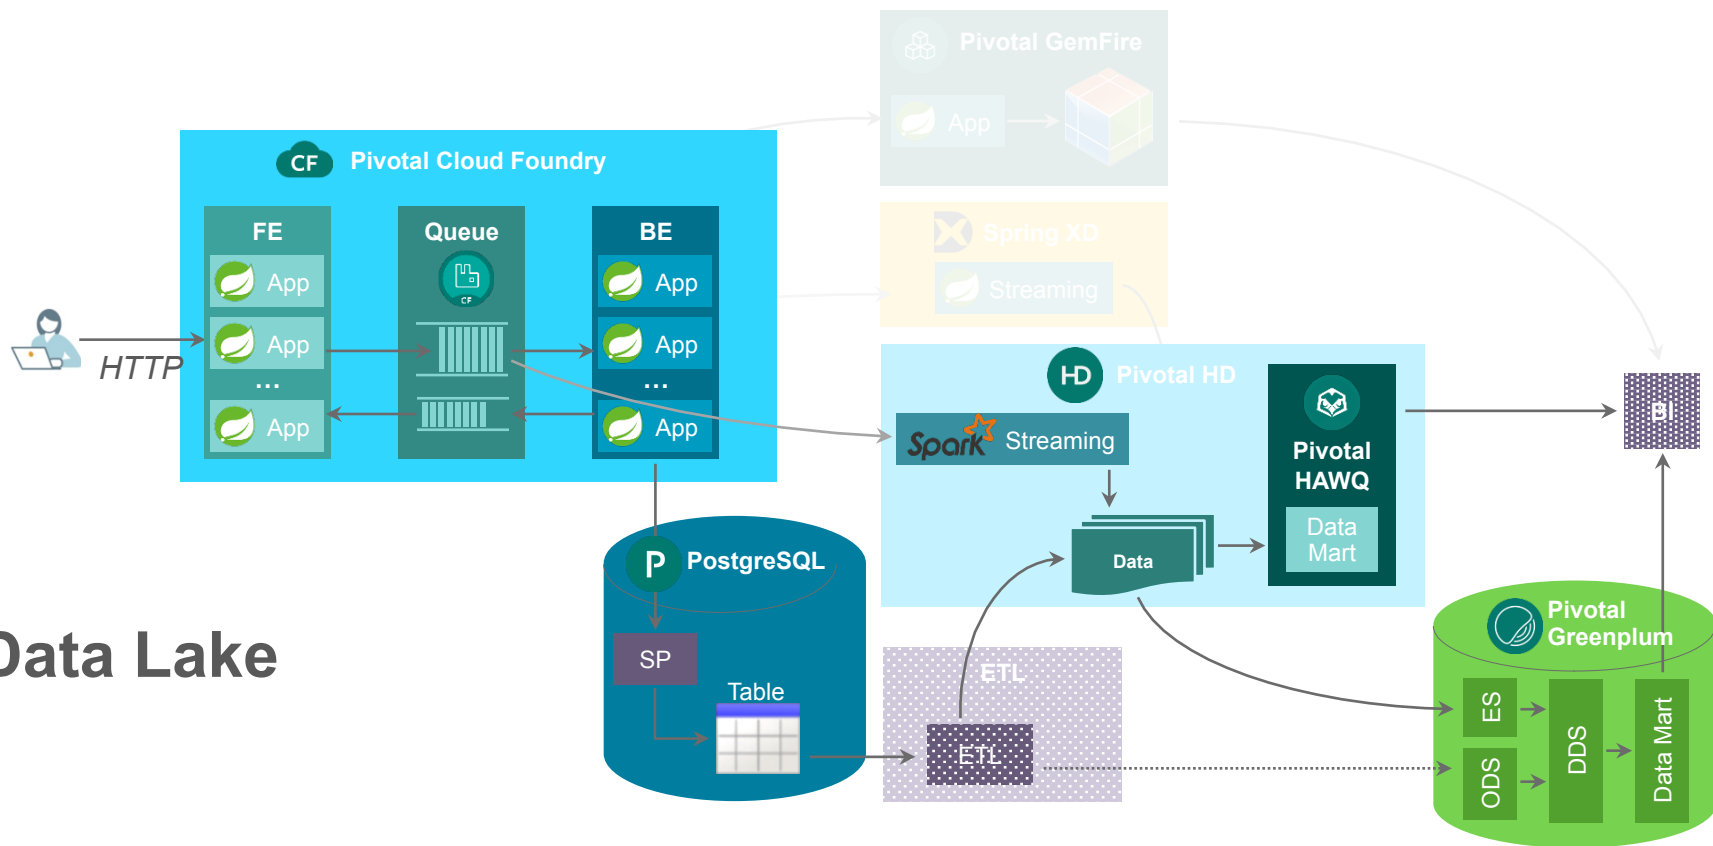


# Pivotal and Modern Data Architecture

**Pivotal Greenplum** – leading analytical MPP database, foundation for the enterprise data warehousing systems and advanced analytics

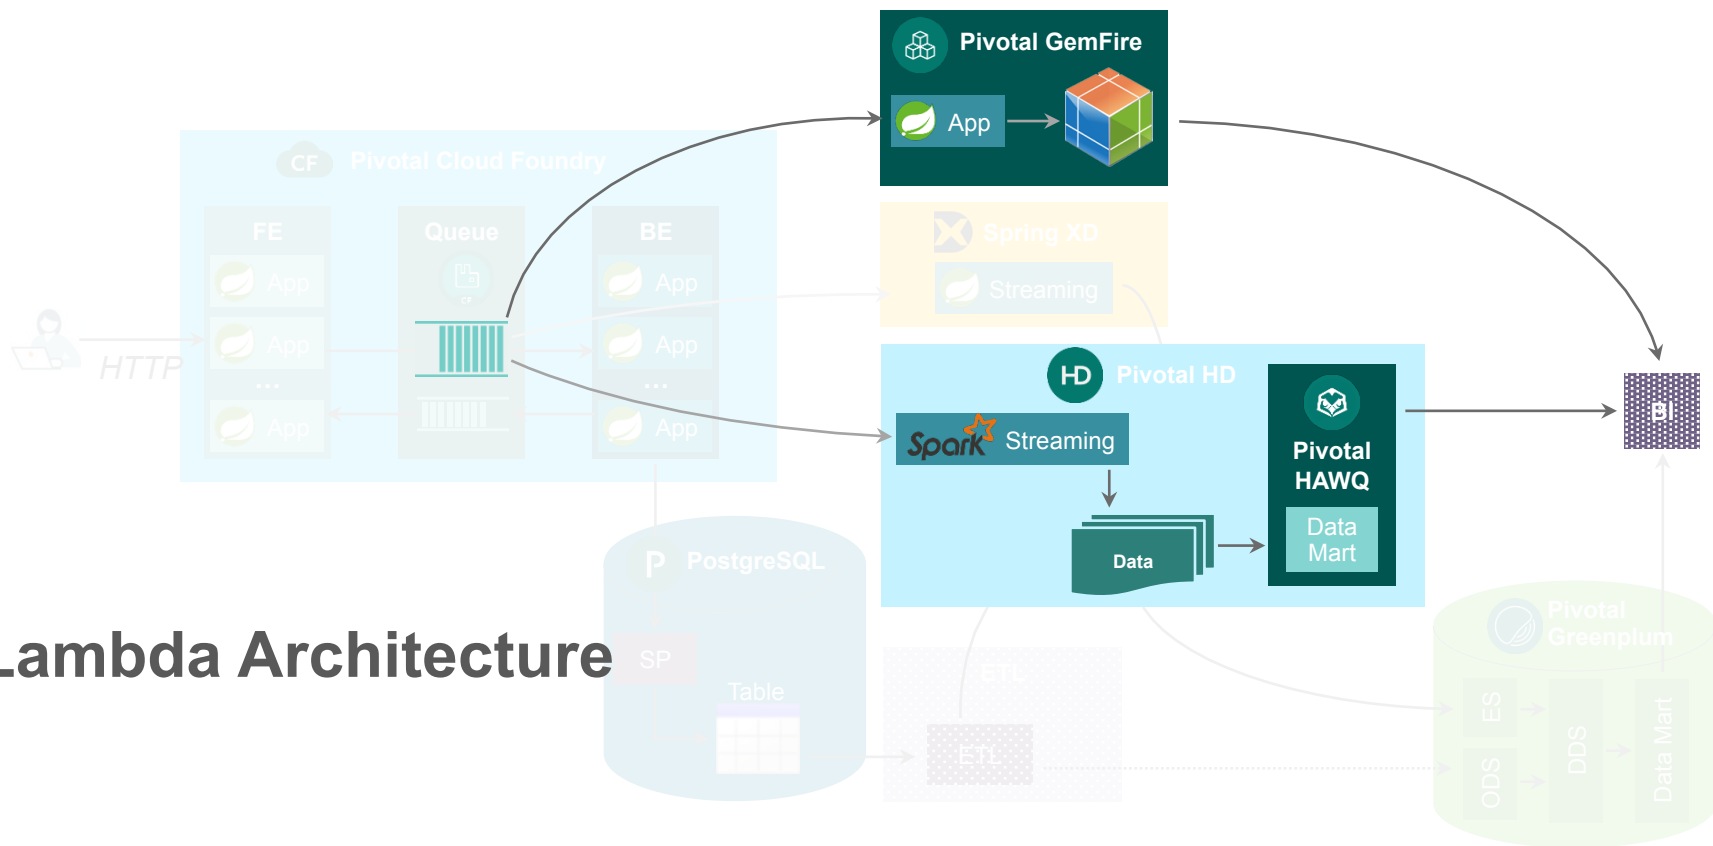


# Pivotal and Modern Data Architecture



Data Lake

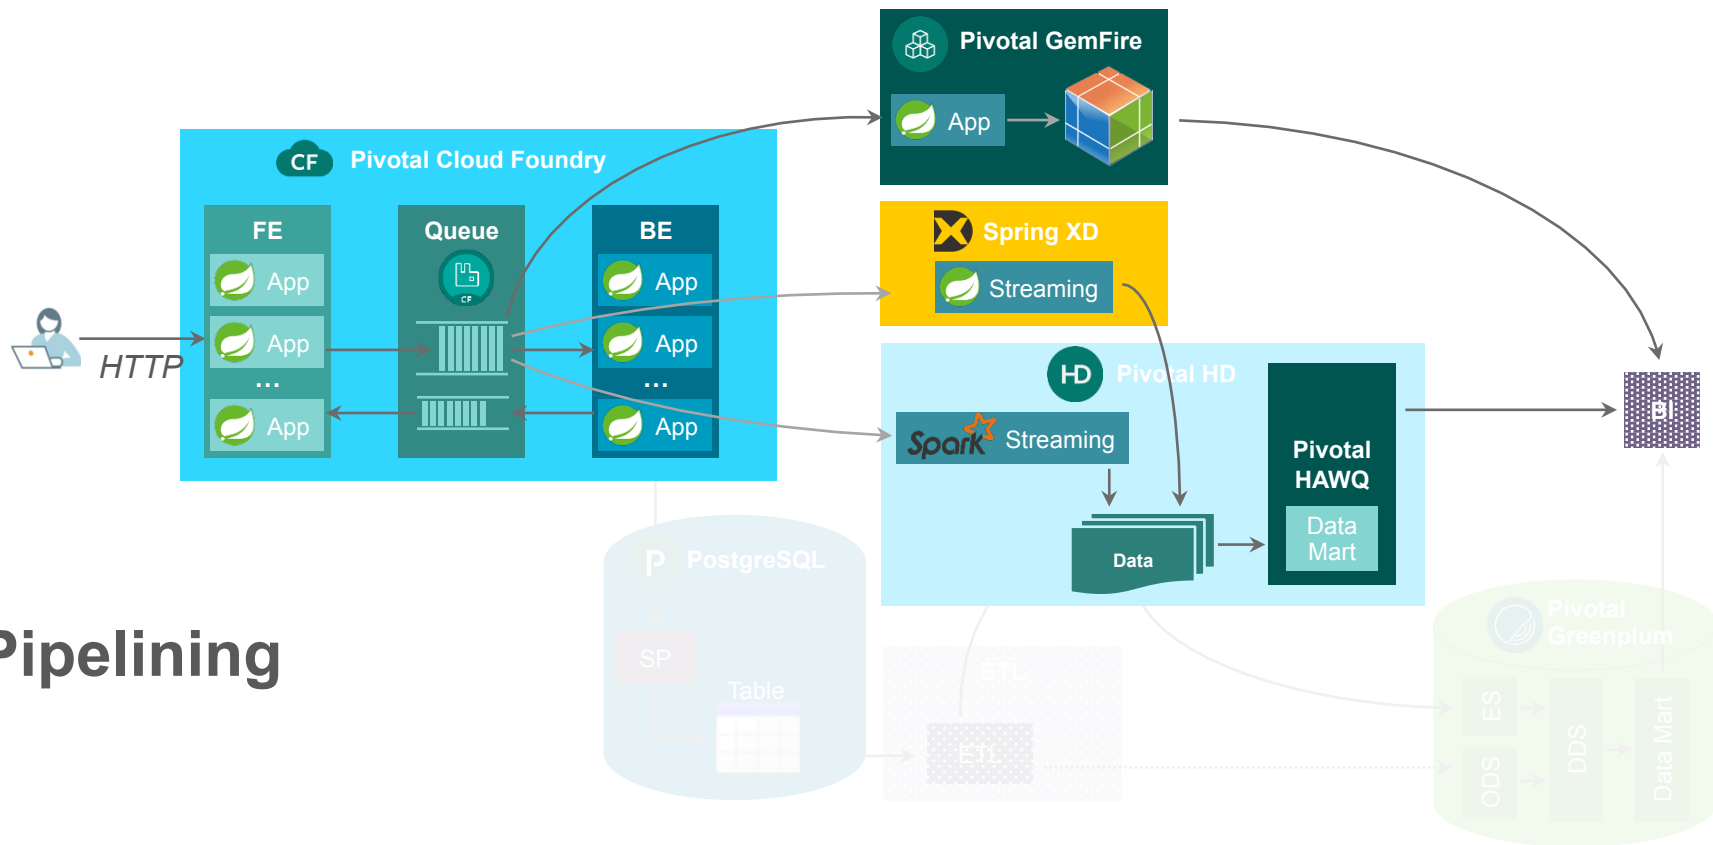
# Pivotal and Modern Data Architecture



**Lambda Architecture**

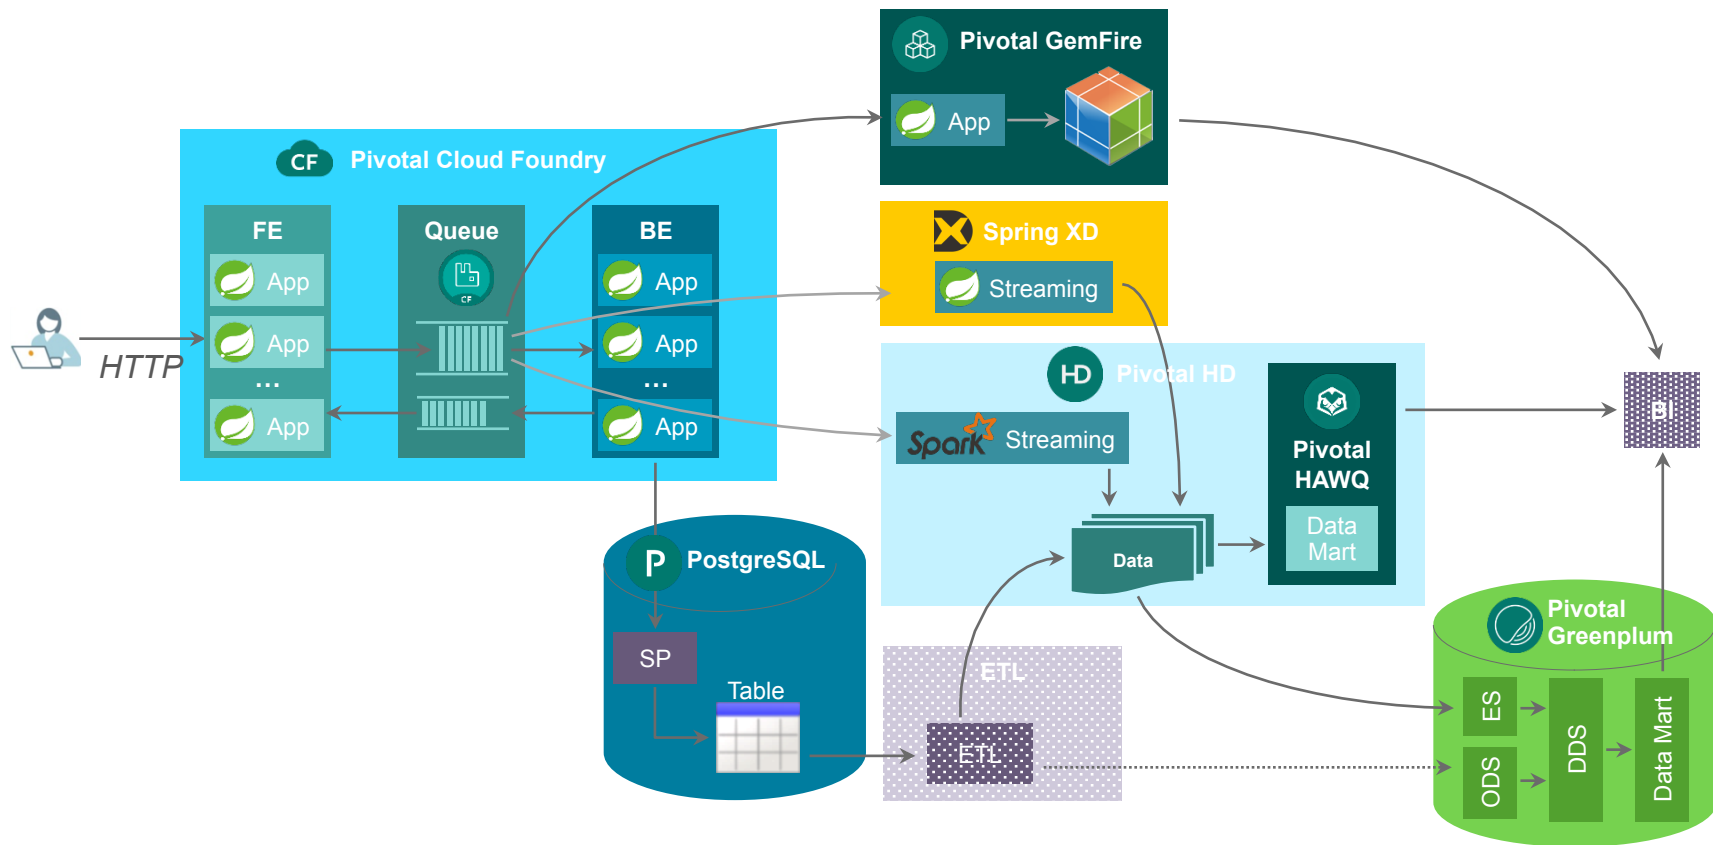


# Pivotal and Modern Data Architecture



Pipelining

# Pivotal and Modern Data Architecture



Questions?  
Questions?

# Pivotal

BUILT FOR THE SPEED OF BUSINESS