aiimi

The Aiimi Insight Engine.

ARCHITECTURE OVERVIEW

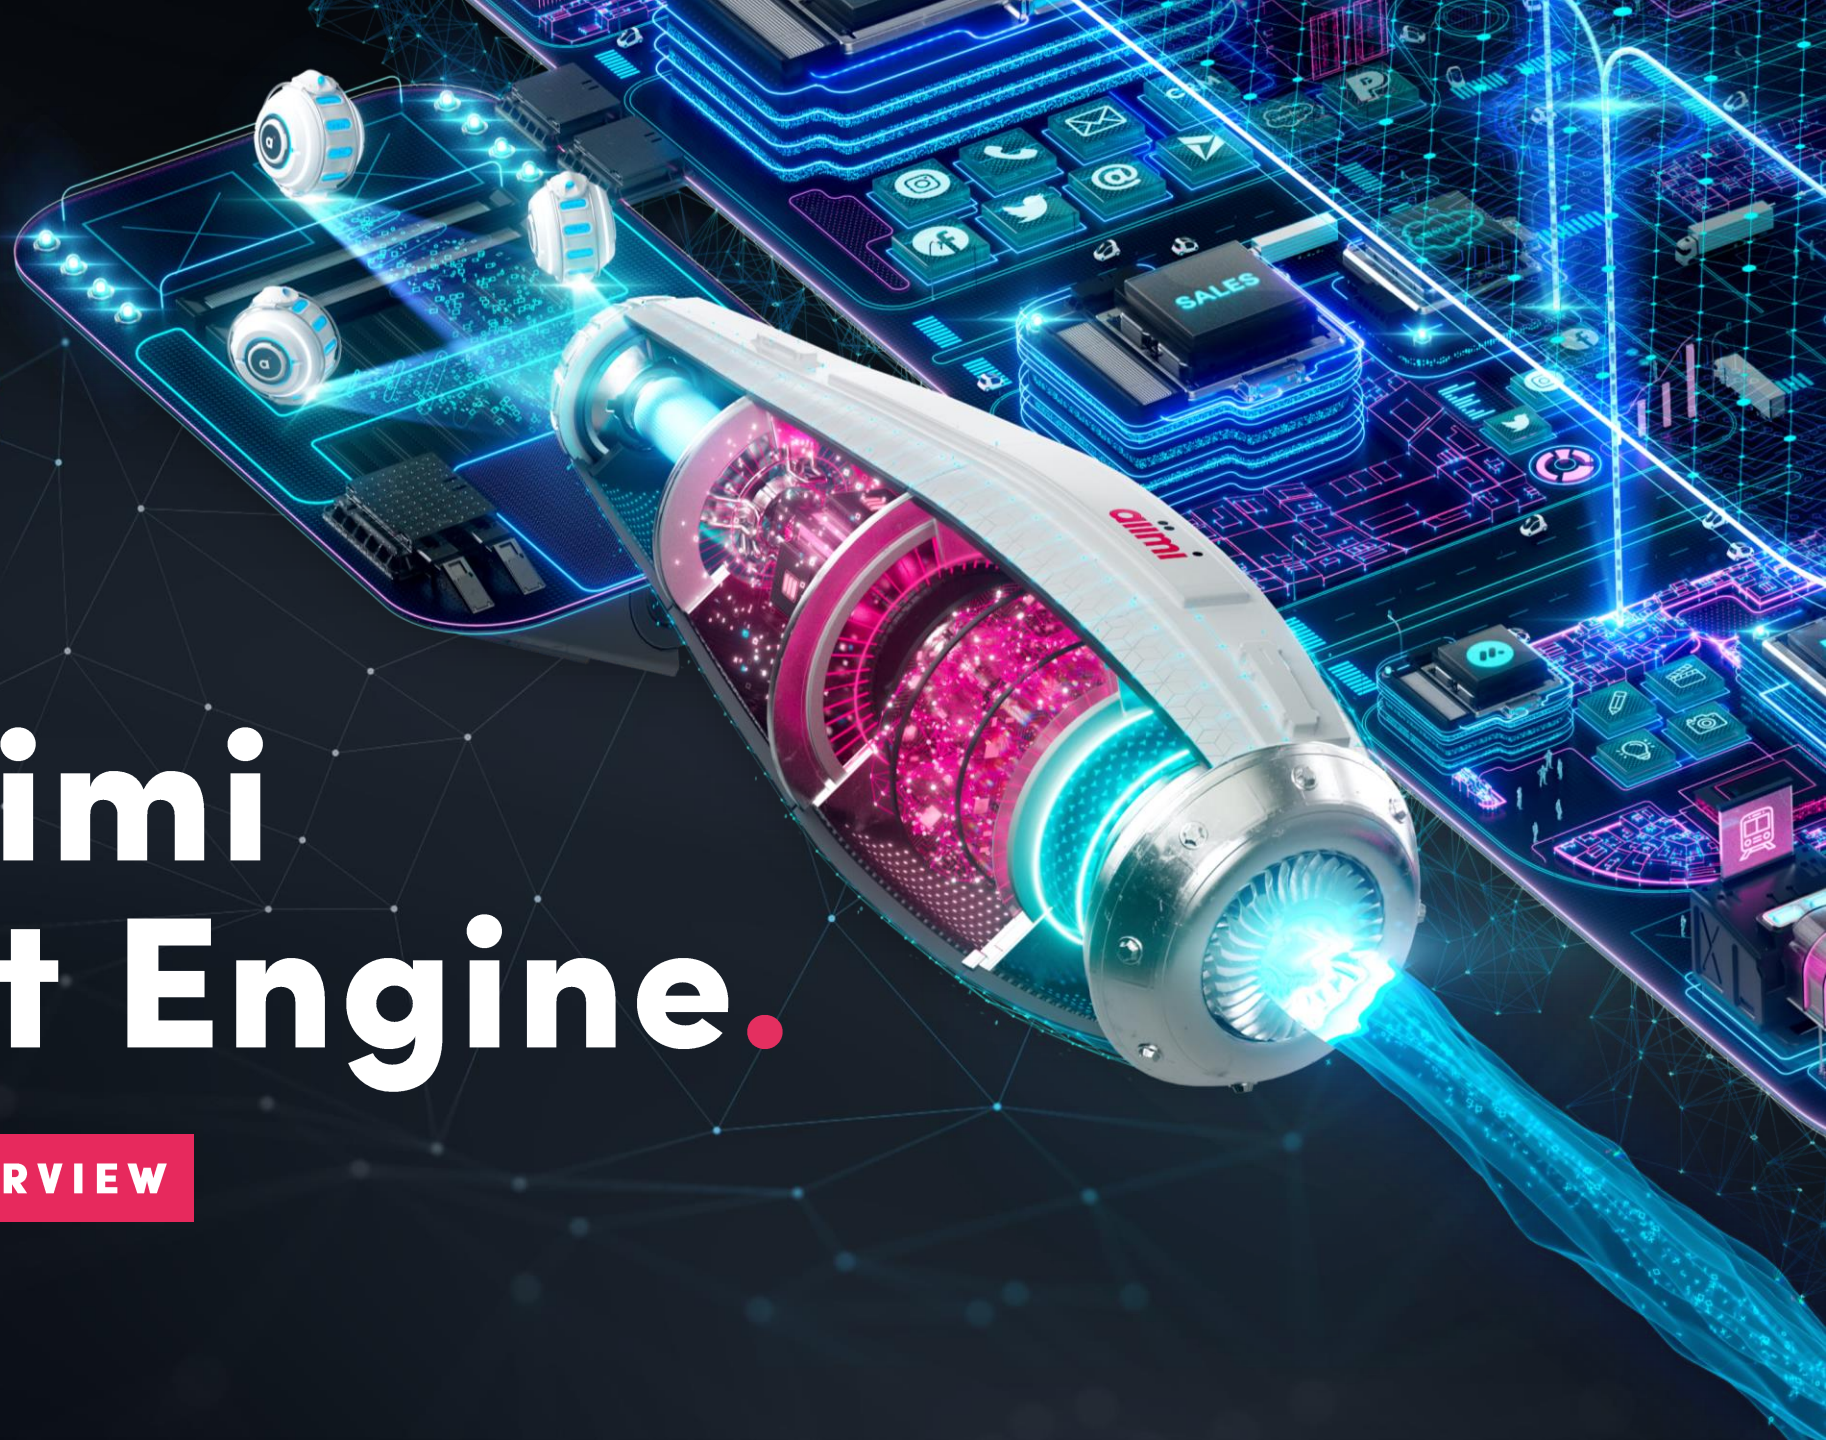

Agenda.

01 / Overview & Key Concepts.

02 / Architecture Overview.

03 / Logical Architecture.

04 / Backup & Disaster Recovery.

05 / Software Versions.
(Client/Server)

06 / Scaling the Platform.

07 / Securing the Platform.

HOW IT WORKS

The Aiimi Insight Engine.

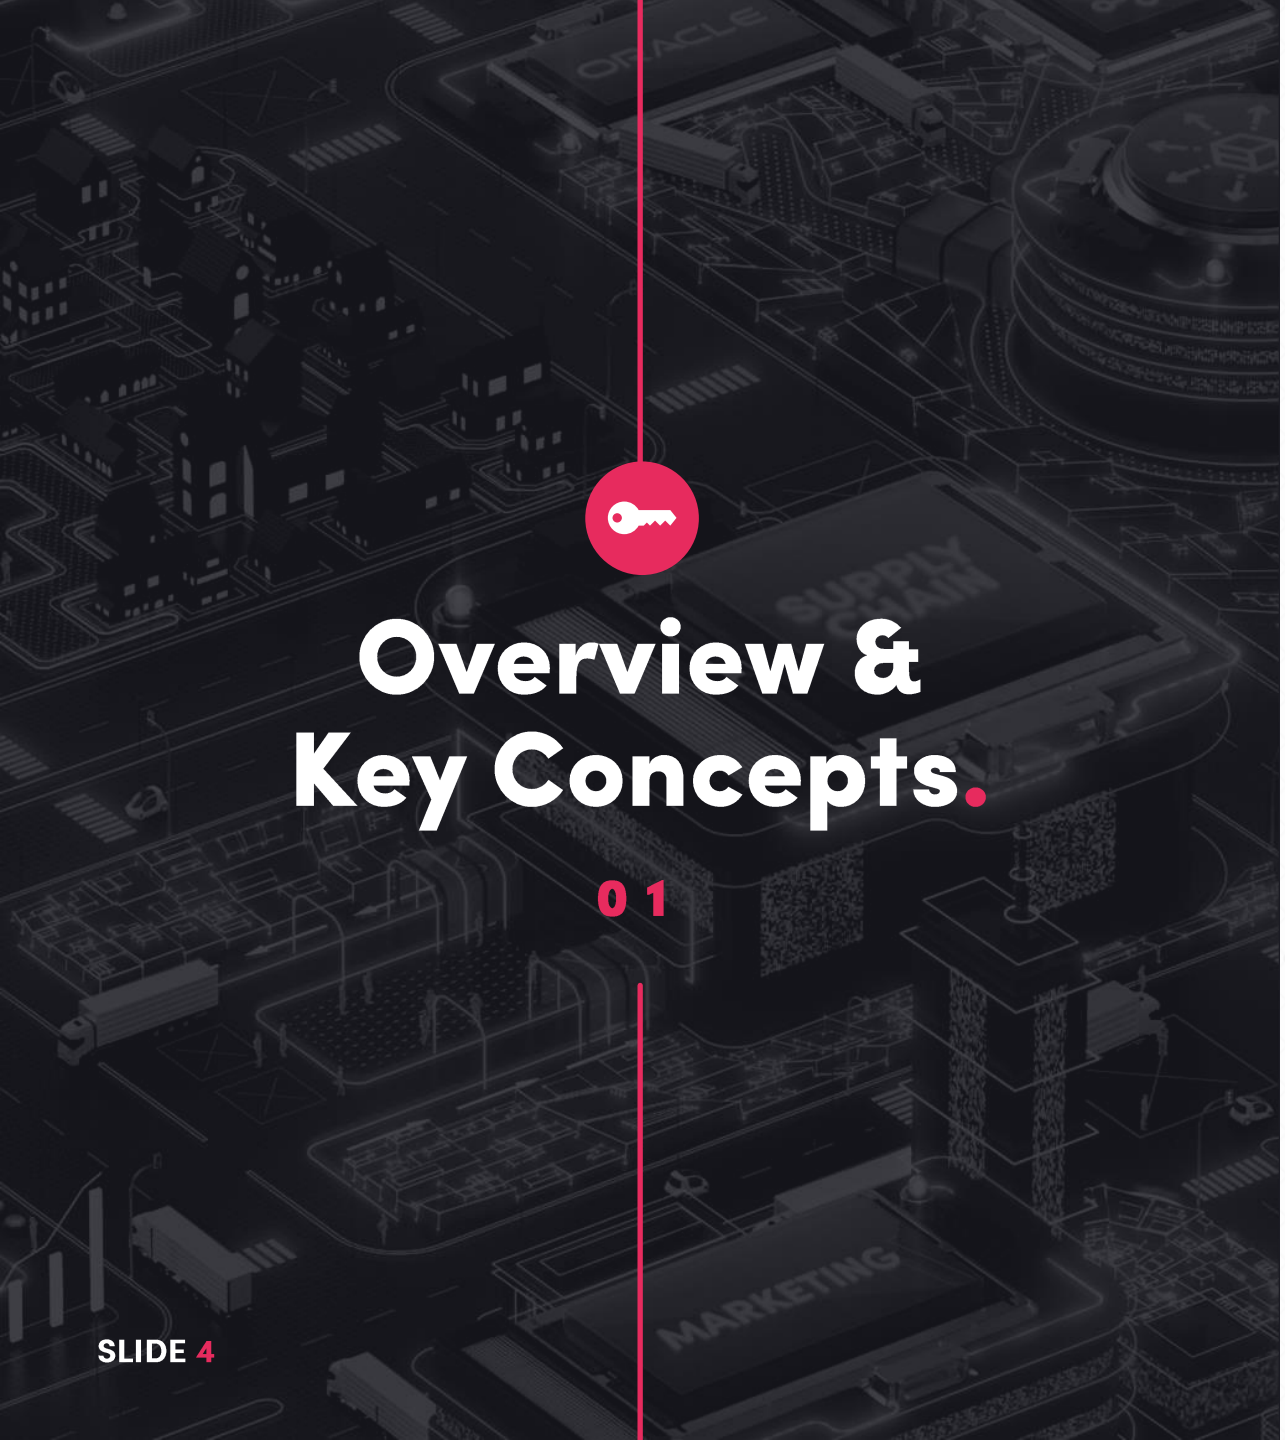

Overview & Key Concepts.

0 1

- Aiimi Insight Engine is an information management platform that brings unstructured and structured information together.
- It offers users a fundamentally different way to consume that information by leveraging the links between things.
- It makes extensive use of machine learning.
- Logically, Aiimi Insight Engine comprises of 5 components.
- Each component has a distinct responsibility.
- Each can be scaled according to the loading characteristics of the respective components.

Ingestion.

OVERVIEW & KEY CONCEPTS 01

- Comprises of a set of source connectors that run in an agent.
- Multiple source agents can run.
- Responsible for both crawling and fetching content from source systems.
- Support for both structured and unstructured data.
- They perform the initial load and then process deltas to pull in changes.
- Deltas run on a schedule or can run continuously.
- In some cases they work on a notification basis (source system dependant).

Enrichment.

OVERVIEW & KEY CONCEPTS 02

- Comprises of a set of enrichment steps that perform operations on documents and data, and then add additional metadata or entities.
- Multiple enrichment agents can run.
- Enrichment steps are configured together into a pipeline.
- Agent processes are responsible for running pipelines.
- Pipelines run on a schedule or continuous.
- Sources can share pipelines if they have the same enrichment steps.
- Aiimi provide numerous out of the box, and it's easy to create new ones using .NET or Python (or any REST/JSON supporting platform).

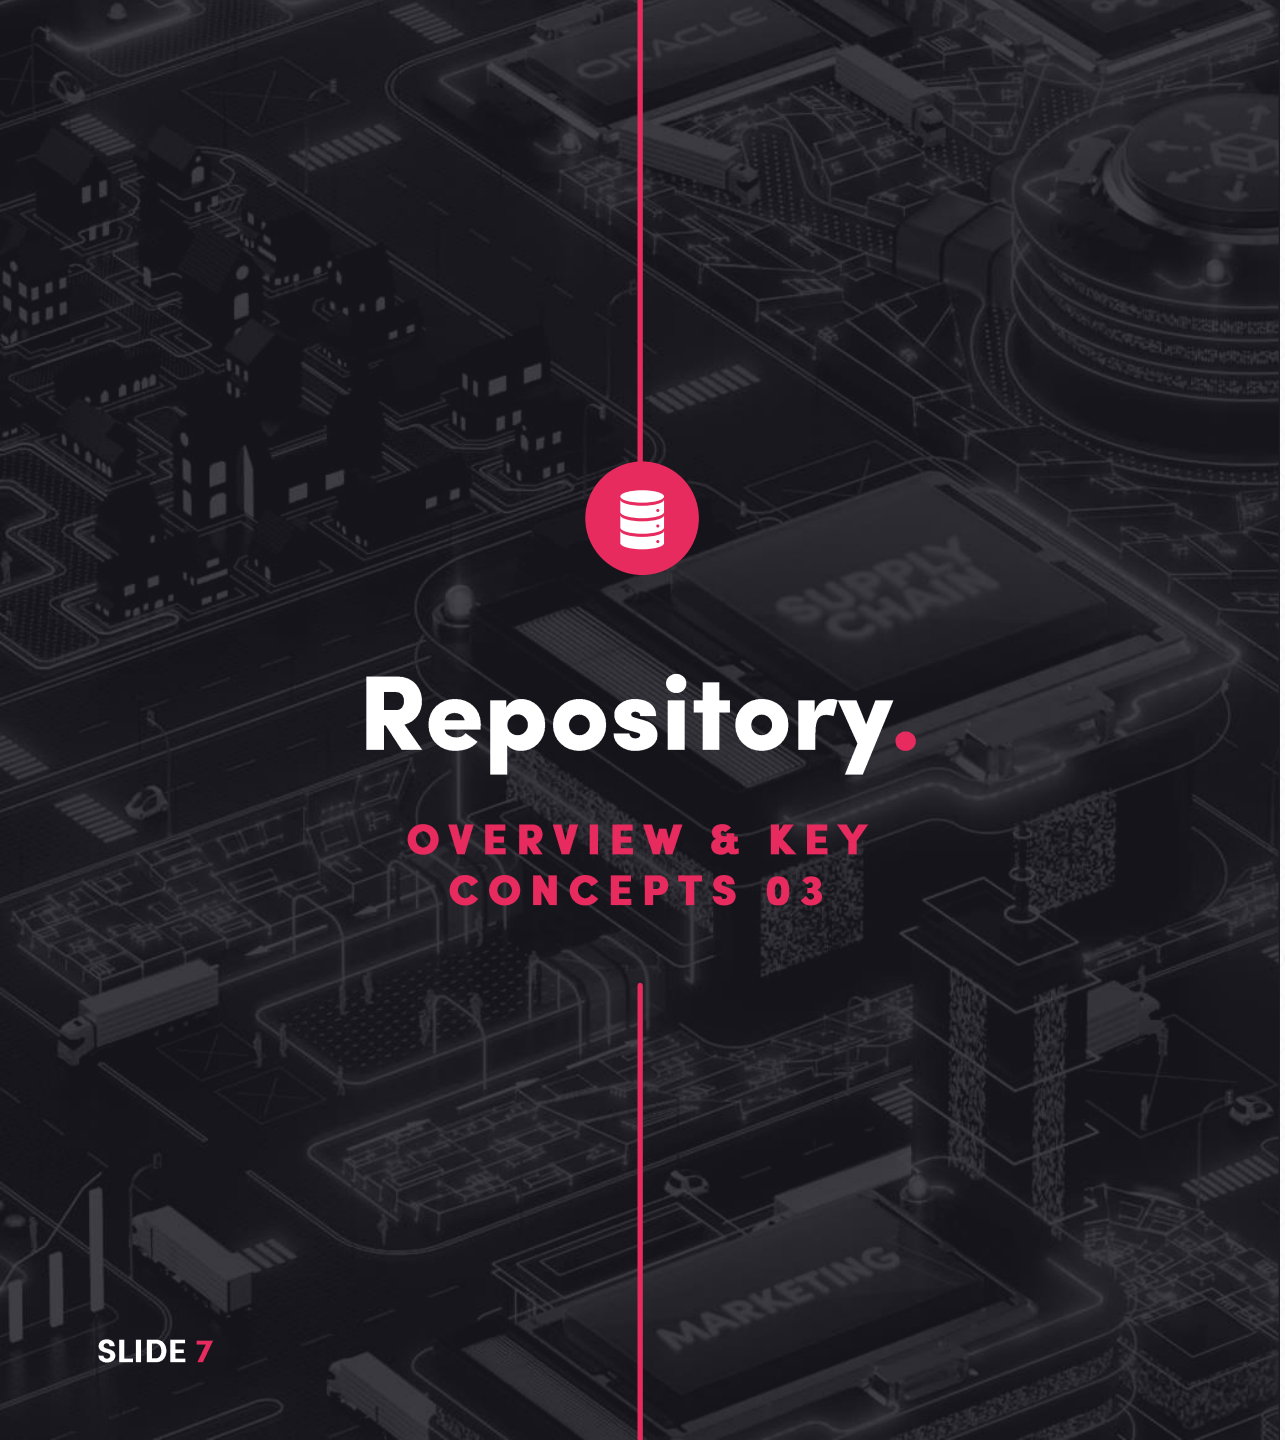

Repository.

OVERVIEW & KEY CONCEPTS 03

- Based on Elasticsearch.
- Provides search and analytical capabilities.
- Easy to scale to accommodate huge volumes of data and documents.
- Very fast response times to user queries.
- Very fast when ingesting new data and documents.
- Natively supports graph queries.
- Natively supports alerts and watchers.

Gateway.

OVERVIEW & KEY CONCEPTS 04

- Sits between the repository and the user interface and the APIs.
- Provides an authentication service for users and API requests.
- Provides an authorization service for requests to ensure people can only see the data, documents and attributes (fields) that they have permission to see.
- Provides a comprehensive auditing service, making sure all requests are fully audited.
- Can be load balanced for resilience and larger user numbers.
- Exposes an app API and a data science API.

User Interface & API.

OVERVIEW & KEY CONCEPTS 05

A series of out of the box **user interfaces** are provided for end users:

- Enterprise search and discovery
- Monitor
- PII data detection
- PCI data detection

Control Hub provides a comprehensive administration portal:

A **Python wrapper** is provided for the Data Science API:

Although any language that understands REST/JSON can use the DS API.

Easy to **create your own apps** using the API:

Typically created with HTML/JS/CSS, although this is not mandatory.

Architecture Overview.

02

- Agents run the ‘back end’ processes.
- Agents can be scaled across multiple servers.
- An Elasticsearch cluster will house the data.
- The Elasticsearch cluster will have data and proxy nodes.
- The gateway and user interface can run on a single server or across a load balanced set of servers.
- The back-end agent servers connect to the respective source systems to pull data and documents in.

Logical Architecture.

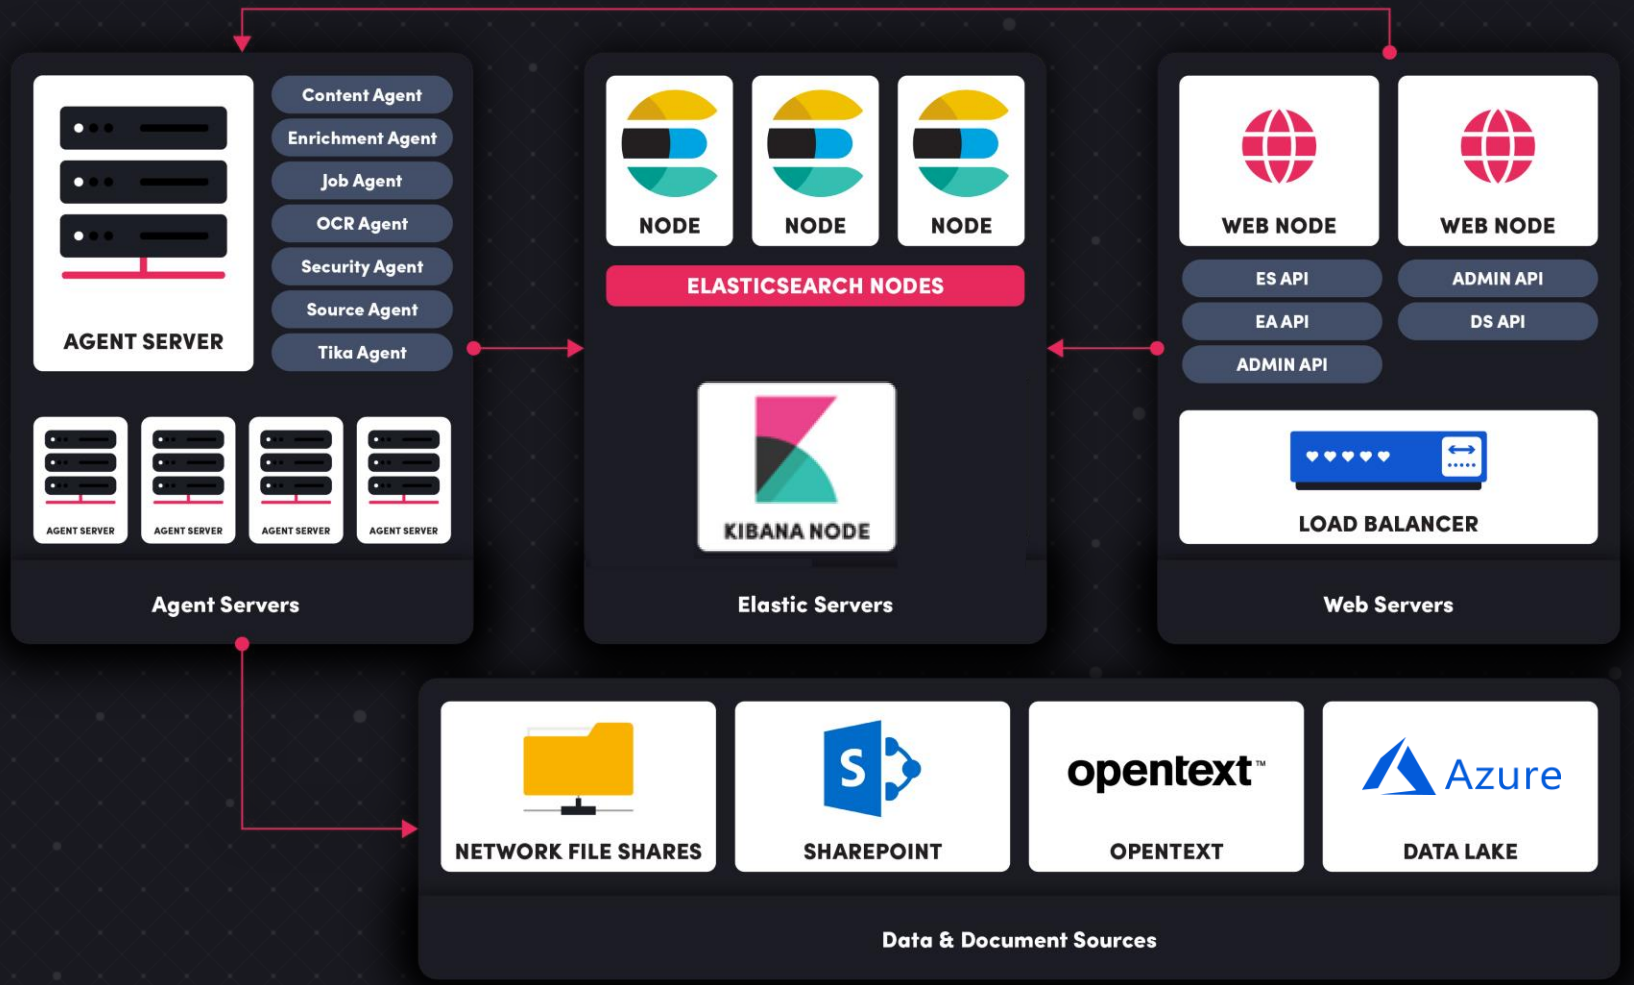

Logical Architecture.

Agents can run across servers.

The only singleton service is the Security Agent, but this can be in a cluster.

A Kibana service can be run for diagnostics.

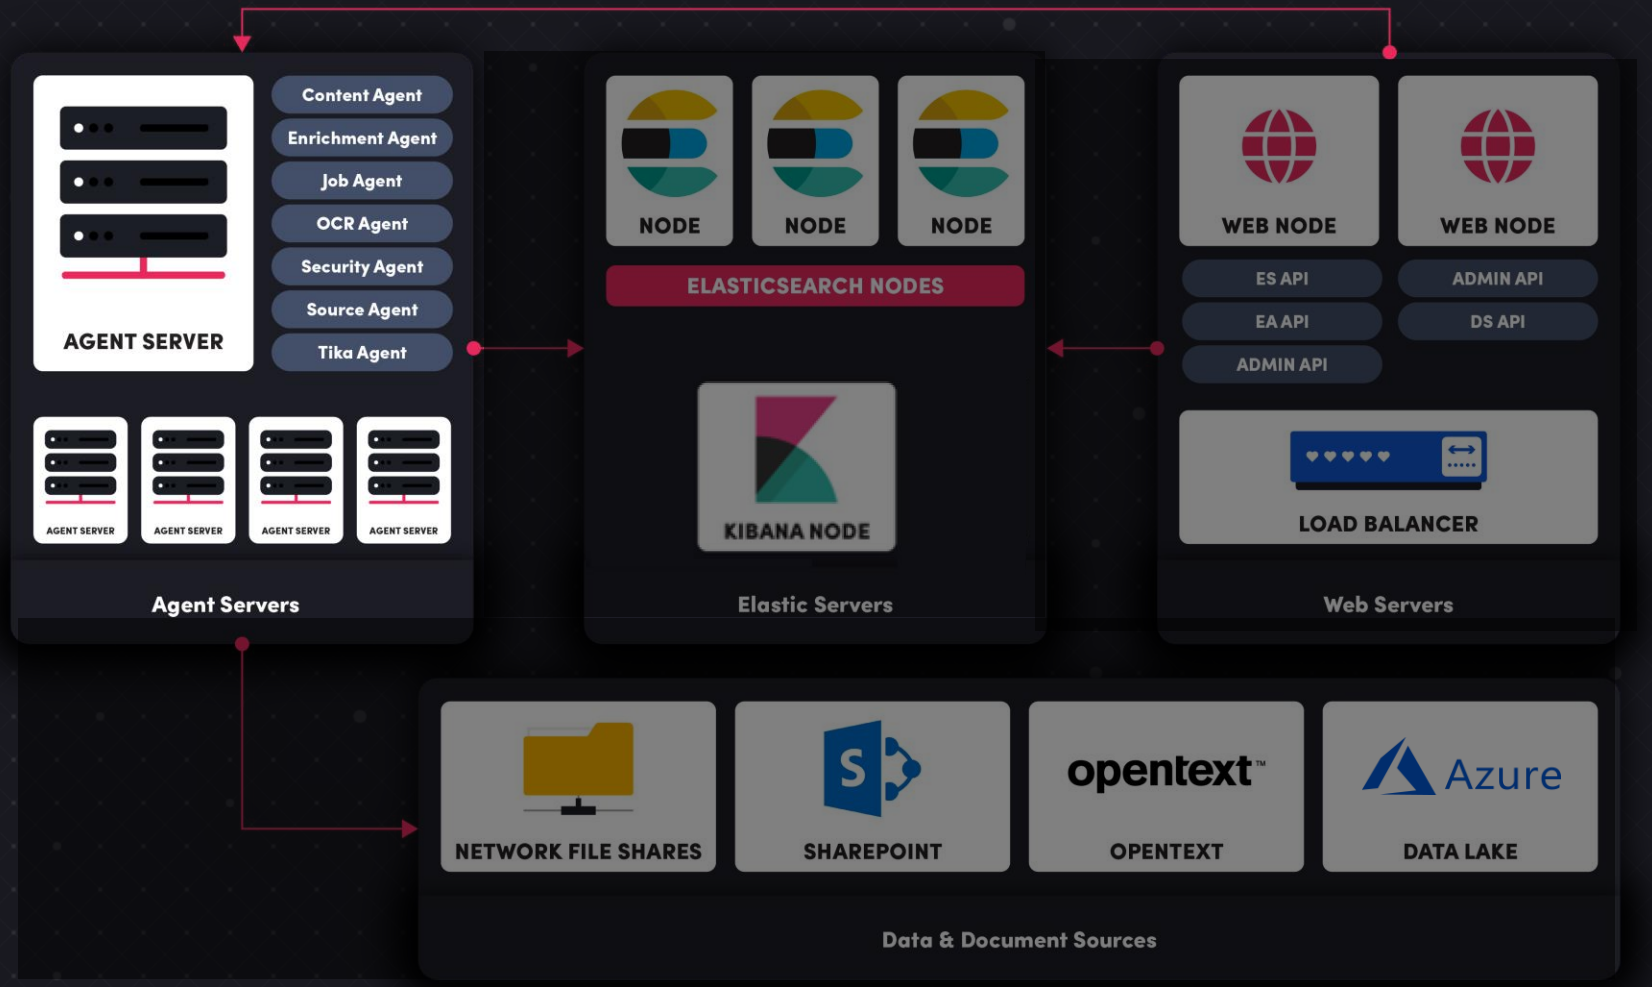

Logical Architecture.

Agents can run across servers.

The only singleton service is the Security Agent, but this can be in a cluster.

A Kibana service can be run for diagnostics.

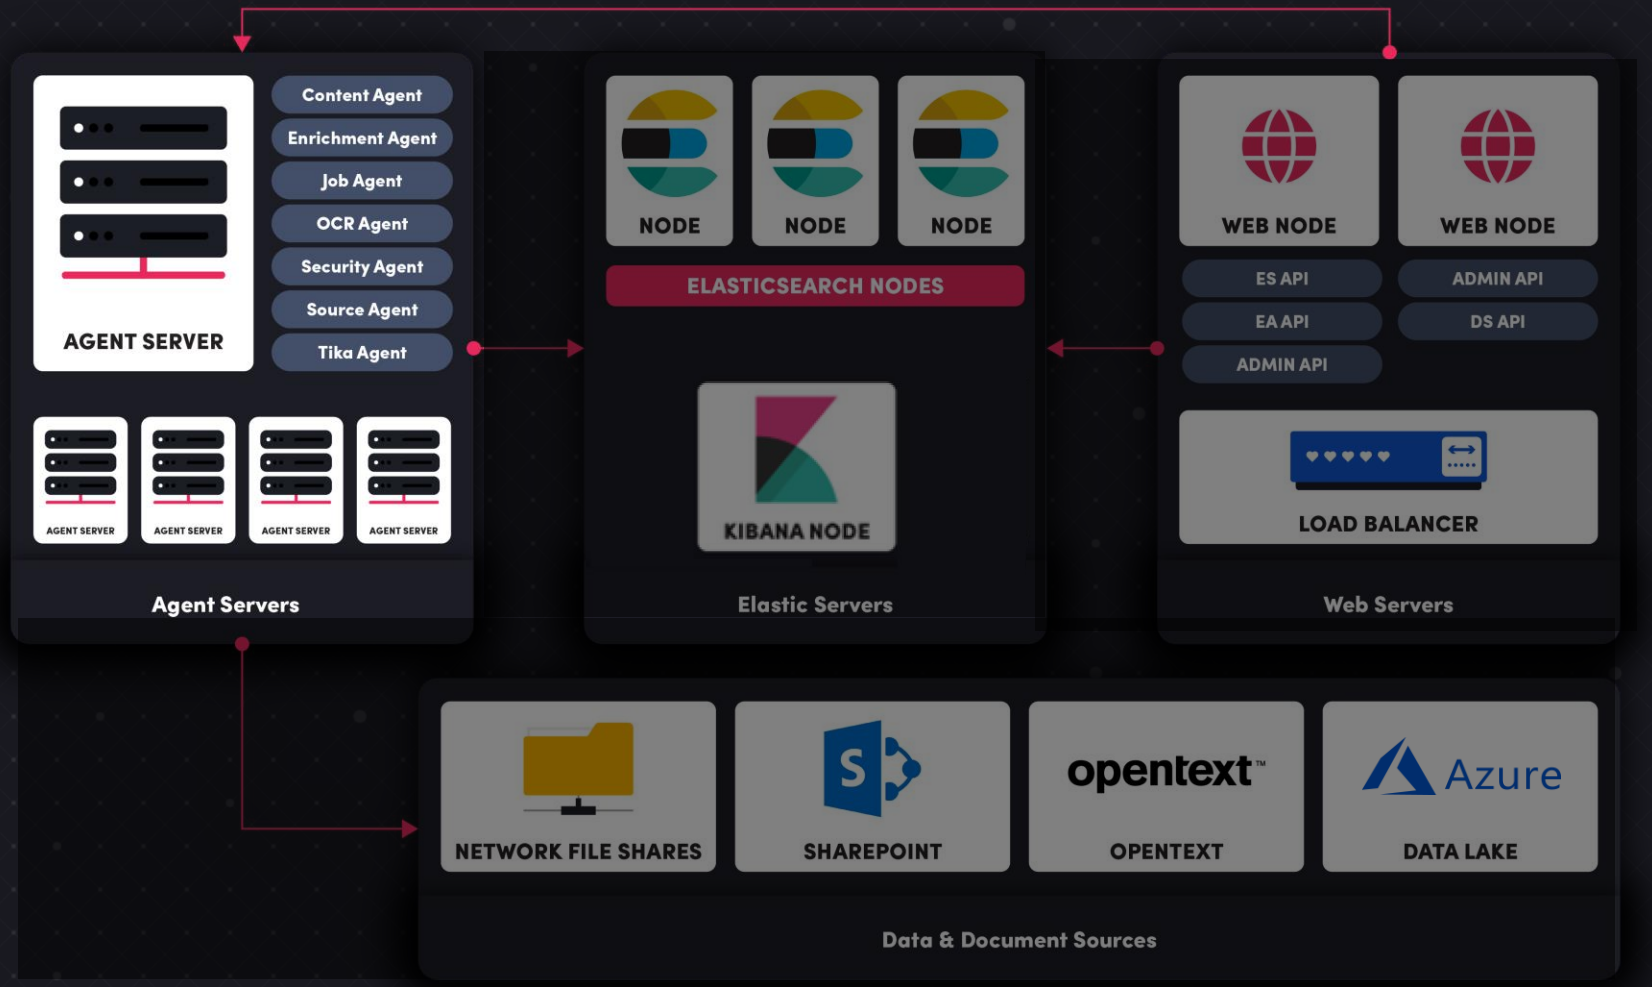

Logical Architecture.

Agents can run across servers.

The only singleton service is the Security Agent, but this can be in a cluster.

A Kibana service can be run for diagnostics.

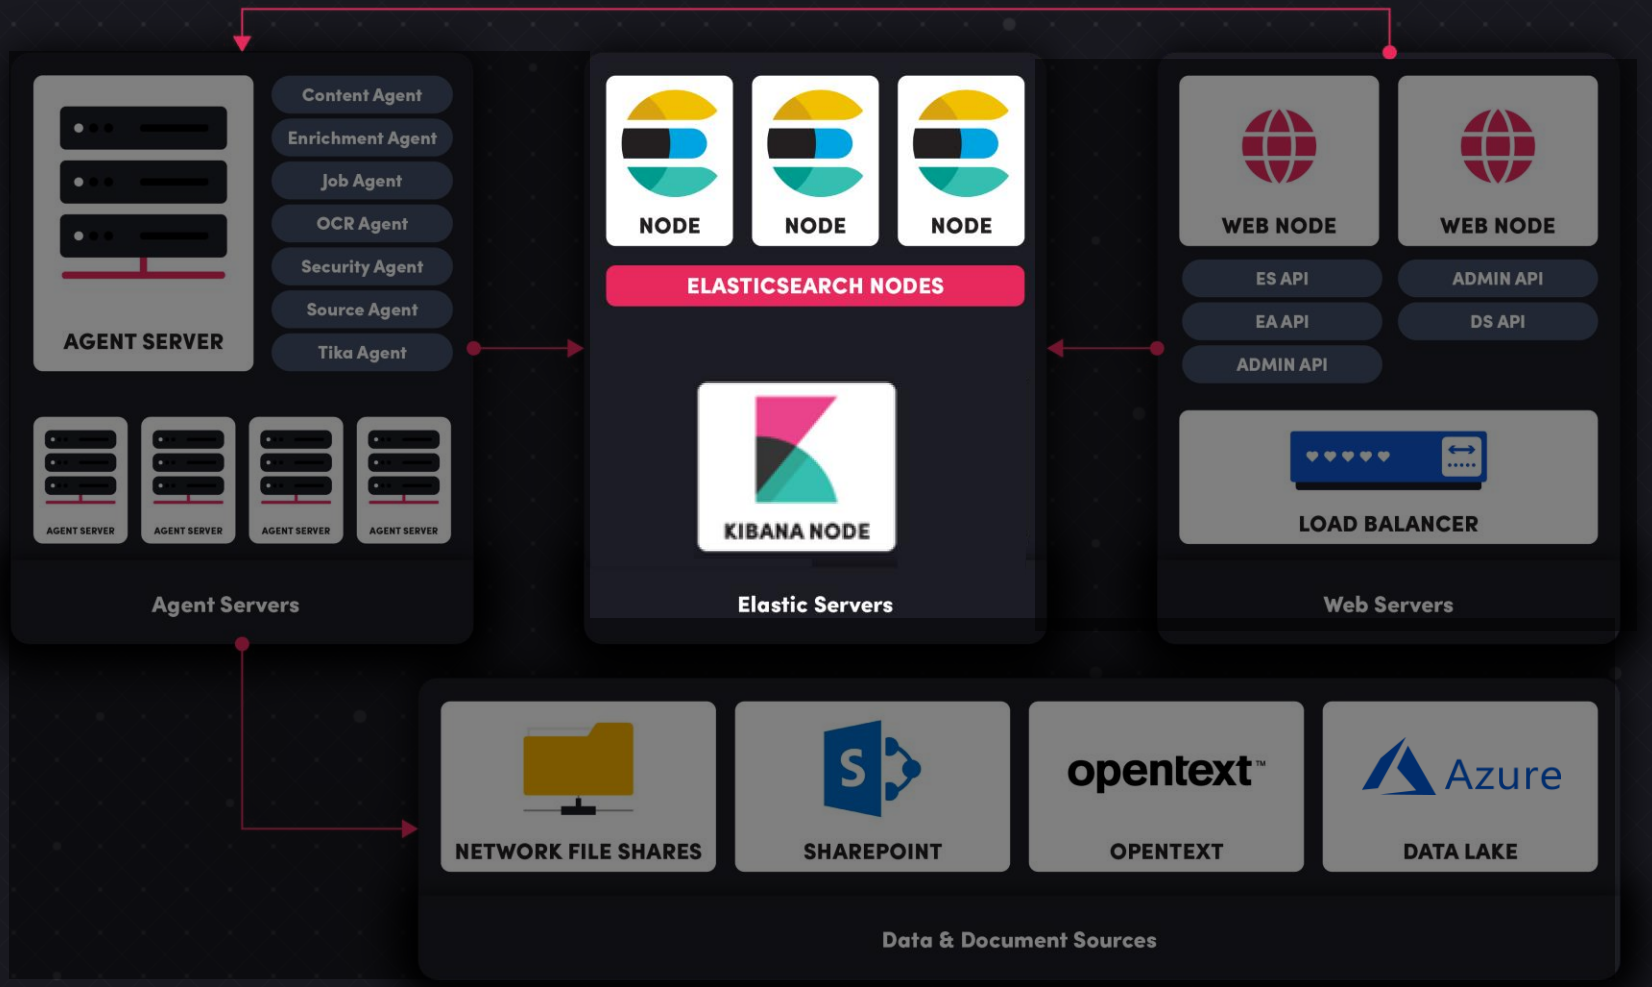

Backup & Disaster Recovery.

0 4

- Elasticsearch snapshots can be used to taken periodic backups of the indexes and configurations.
- Snapshots and restore are a built-in feature of Elasticsearch.
- We have an agent process to automate these, and they can run on a schedule.
- Disaster Recovery can be handled with cloud IaaS resilience features .
- Alternatively we can maintain a shadow DR system.
- Given the nature of the system, it can easily ‘catch up’ from a previous snapshot by running the source connectors which will perform a delta.

Software Versions.

CLIENT / SERVER
05

Server

- Red Hat Enterprise version 7 & 8
- Elasticsearch 8.3
- 6 Week Release Cycle

Client

- Chrome
- Edge Chromium

Scaling the Platform.

06

As content volumes grow:

- Tune Elasticsearch RAM ratios (heap to file system cache)
- Horizontally scale Elasticsearch
- Vertically and/or horizontally scale Aiimi Insight Engine Agent Servers

A user concurrency grows

(either more users or busier users):

- Horizontally scale web application servers

Securing the Platform.

07

SSL & TLS 1.2 / 1.3

- Security between client and web tier
- Security between nodes
- Security between agents/web tier and Elasticsearch nodes

Firewall rules to close unnecessary ports

Firewall rules to control which servers can talk to each other

Credentials for source systems encrypted

SAML2, LDAP Integration

- Third party auth platforms
- MFA

All users authenticated

All user requests undergo authorisation and permission trimming


ORACLE
SUPPLY CHAIN
SAP
MARKETING
IT
aws
citrix