



IBM ECM UK UserGroup 2008  
**Connecting the  
ECM Community**

Technical Breakout



## Speakers

- **Phil Lyons**      **ECM Consultant**

- [phil.lyons@uk.ibm.com](mailto:phil.lyons@uk.ibm.com)



- **Giles Metcalf**      **ECM System Architect**

- [giles.metcalf@uk.ibm.com](mailto:giles.metcalf@uk.ibm.com)



- **Alan Bluck**      **ECM Systems Architect**

- [alan.bluck@uk.ibm.com](mailto:alan.bluck@uk.ibm.com)



- **Stephen Orrell**      **ECM Systems Architect**

- [stephen.orrell@uk.ibm.com](mailto:stephen.orrell@uk.ibm.com)



## Topics

- **Global Company, Global Architecture**
- **High Availability, Supporting Mission Critical Applications**
- **Security, Keep it Secret, Keep it Safe**
- **Open Discussion**

## Global Company, Global Architecture

- **Phil Lyons, ECM Services Consultant**
- **Overview of Challenges and Focus Areas**
- **Decentralised Implementation**
  - Using new distributed architecture features of P8 4.0
- **Real Life Example**
  - Global Architecture Implemented by Large Financial Conglomerate
- **Discussion**

## Global Company, Global Architecture

- **Key Challenges**
  - Infrastructure
  - Data Centres
  - WAN Capability
  - Satellite Offices
  - Local Data Privacy Laws, Policies, Banking Regulations
  - Distributed User Base
  - Duplicated Requirements
  - Duplicated and Distributed Implementation Streams
  - Duplicated Costs
  - IT Controversy
  - Business Process Controversy
  - Avoiding Silos



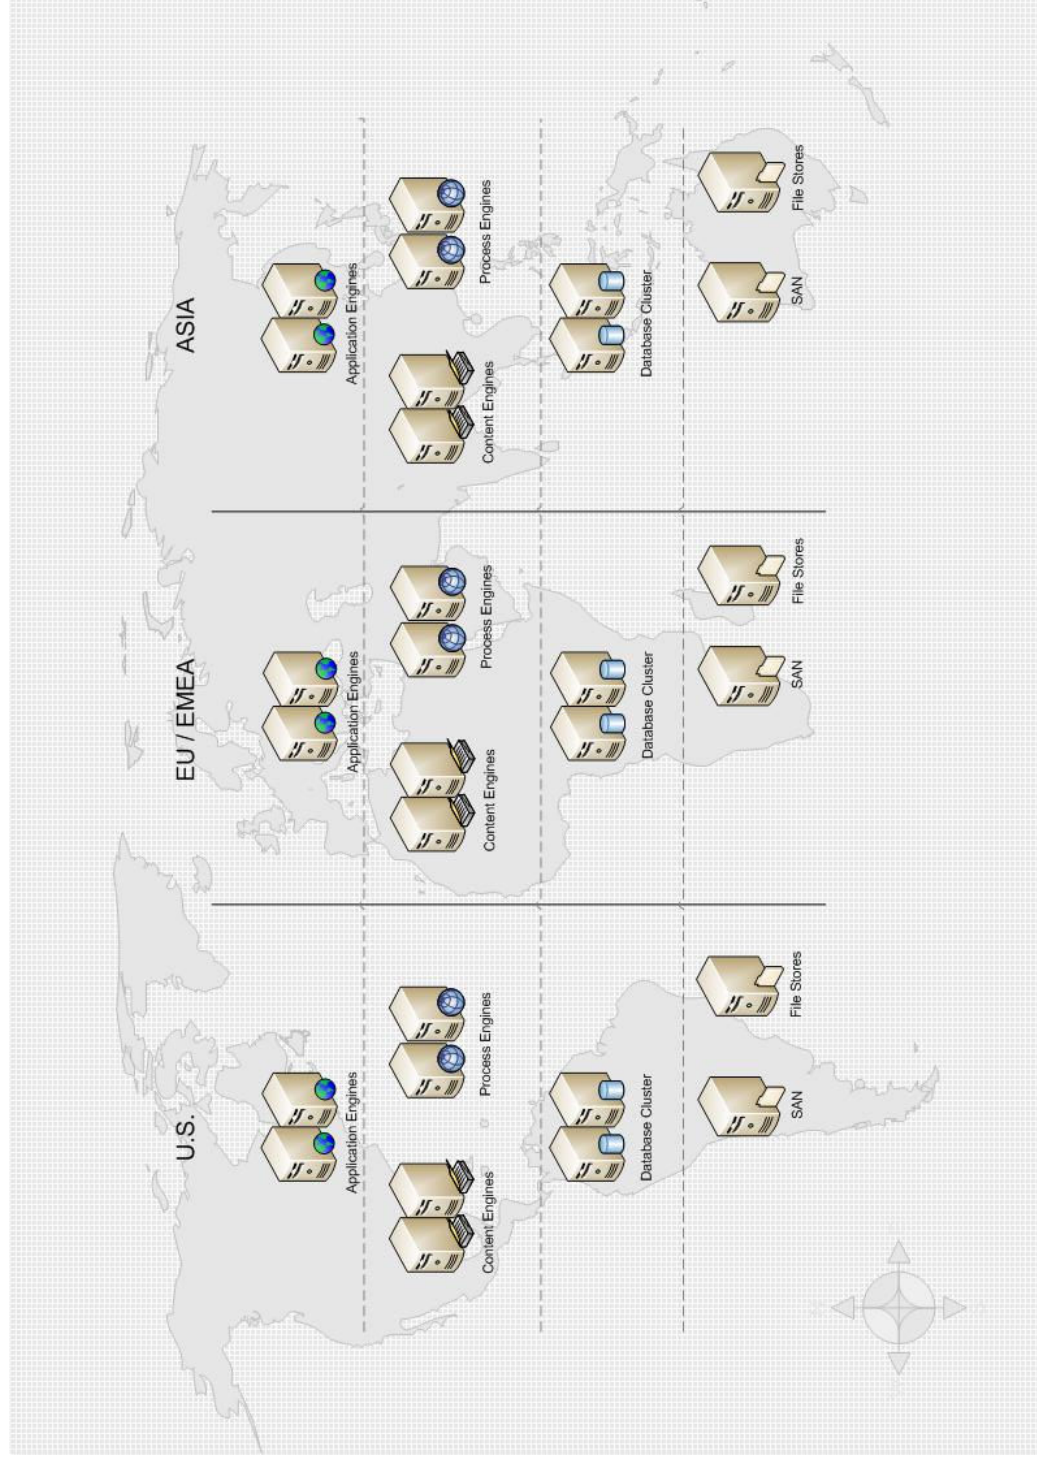
## Global Company, Global Architecture

- **Areas of Focus**
  - Geographical Regions / Poles, US, EU & Asia
  - Infrastructure
    - Decentralised Implementation
    - Data Centre Requirements
    - WAN Capability
  - Data Privacy Requirements
    - Distributed Storage
  - Visibility as a Single Domain
    - Single Point of Access
      - User
      - Application Integration
      - Enterprise Search






# Decentralised Global Implementation



## Single Domain?


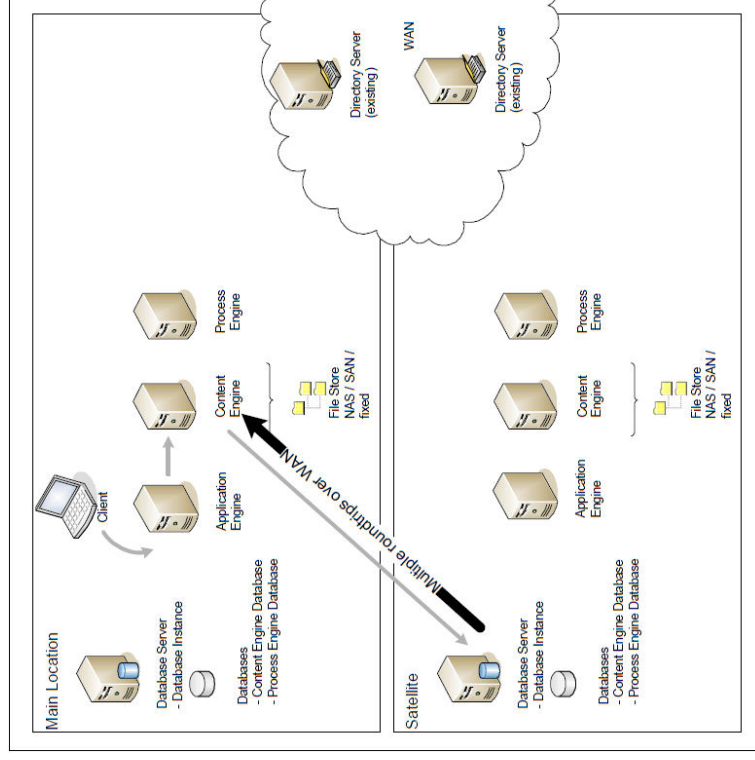
- Use Upgrade Tool to Merge into a single P8 Domain
- Bi-Directional GCD Database Replication





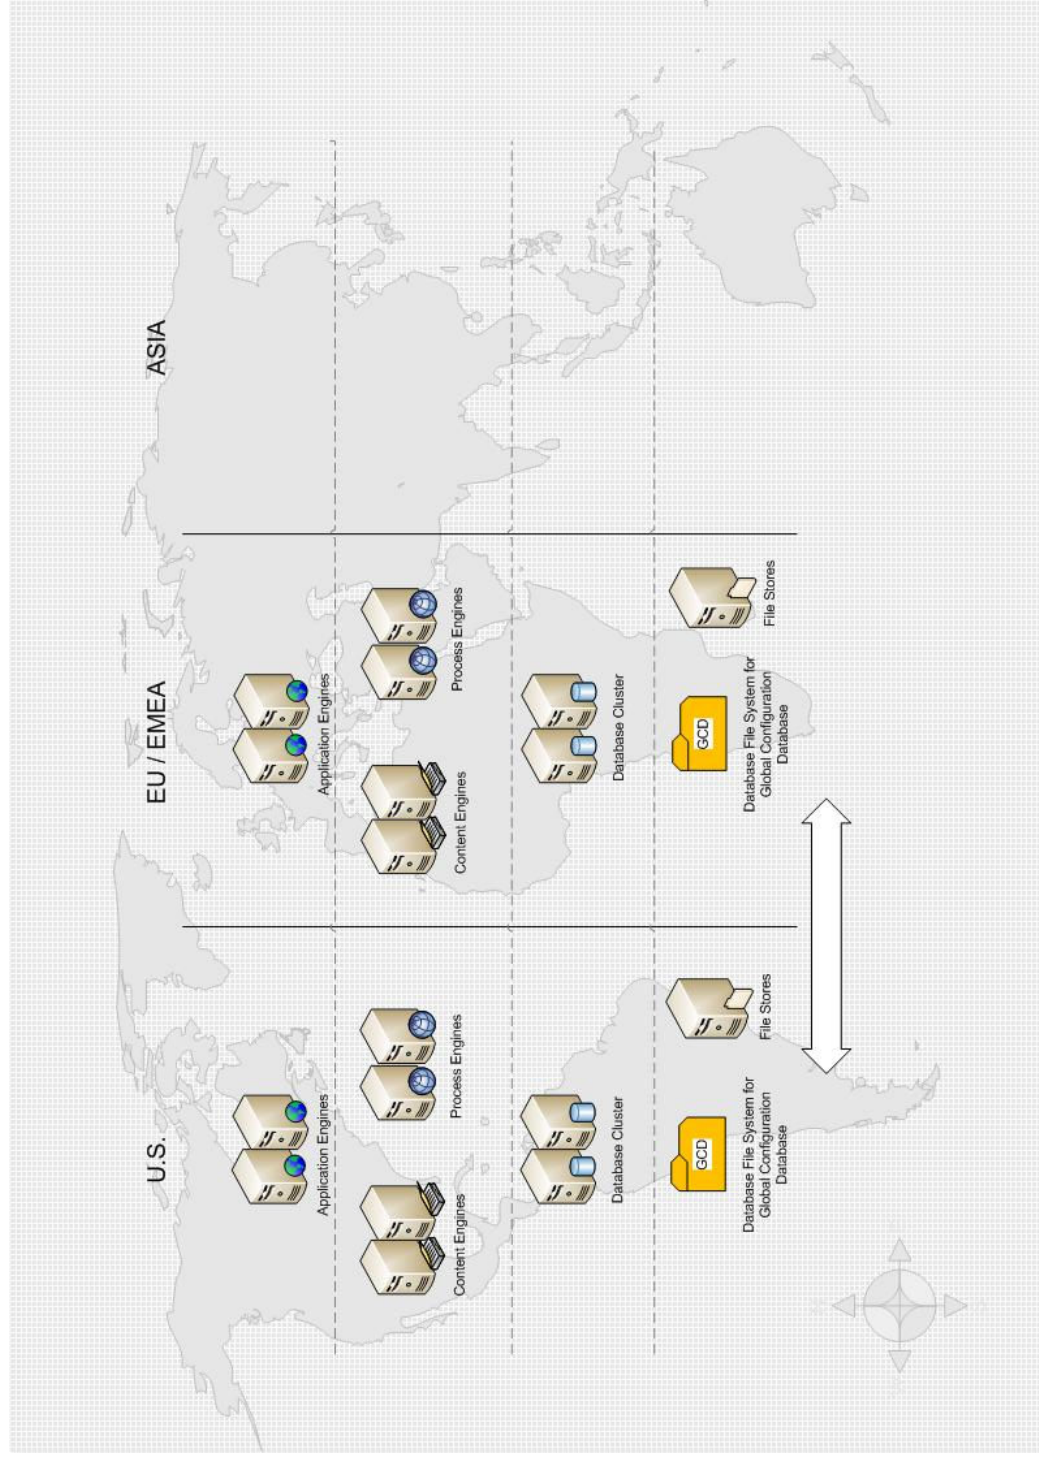
# Request Forwarding

- When talking about distributed systems, the efficient use of the network bandwidth between the locations is essential. In this area, IBM FileNet P8 4.0 includes major benefits.





# Real Life Example – Large Financial Conglomerate



## Real Life Example - Objectives

- **Single Repository for all Documentation**
- **Support Enterprise Search**
  - Single Domain
- **Managed by Central CoE in the U.S.**
  - IT Contollership
  - Best Practice
  - Reusability
  - Simplification
  - Application Roadmap and Strategy
- **Central Point for Application Integration using WebSphere**
  - ECM Everywhere, enable Front Office and Back Office applications to provide Document Management Functionality at transactional level.

## Discussion

- **Distributed Storage**
- **Centre Of Excellence**
- **Monitoring**
  - FileNet System Monitor
  - User Experience Monitoring Tools



## High Availability, Supporting Mission Critical Applications

- **Giles Metcalf**    **ECM Systems Architect**
- **Comparison of High Availability Approaches**
  - Farm versus Cluster in P8 Published Supported HA Practices
- **Some Real World Examples**
  - P8 ECM High Availability best practices put into effect
- **Discussion**

## *High Availability and Disaster Recovery are not equivalent!*

- **High Availability** mitigates against minor system failures
- **Disaster Recovery** mitigates against extended loss of major systems
- **High Availability** systems are generally not geographically separated
- **Disaster Recovery** systems are *always* geographically separated
- **High Availability** failovers are instantaneous and transparent to users
- **Disaster Recovery** failovers may involve a time lag before service is restored


***A High Availability architecture does not remove the need for a good backup and DR policy!***



## Some options for High Availability architectures


### Server Farms

- Multiple equivalent systems and load balancers
- Best approach for application servers with minimal shared data
- Not optimal for data-centric servers (such as database servers)
- May show performance drop on failover unless adequately sized
- Requires load-balancing hardware or software
- Scalable



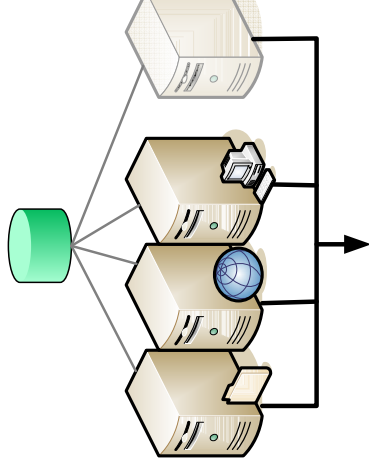
### Asymmetric (Active-Passive) Server Clusters

- Equivalent systems and shared storage
- Best approach for data-centric servers with frequent updates
- Most expensive option for idle hardware
- No performance drop on failover
- Requires clustering software
- Less easily scalable



### Asymmetric (N+1, N+M) Server Clusters

- Best for data-centric servers with frequent updates
- Multiple options for additional standby systems
- No performance drop on failover if N+1 or N+M architecture used
- Requires clustering software
- Can be scaled
- Multipurpose systems and shared storage




## High-Availability Best Practice

- **Cluster Data Tier servers for high availability active - passive via symmetric, asymmetric or N+1 clusters type**
  - Exception: Oracle RAC can be deployed active - active
- **Farm P8 engine servers for high availability and scalability**
  - CE behind a hardware or software (e.g. WebSphere) load balancer
  - PE *must* be behind a hardware load balancer
  - AE behind a hardware or software (e.g. WebSphere) load balancer, using sticky sessions where P8 Workplace is utilised
  - Use shared disk between farmed AE instances to centralise Workplace configuration

## High-Availability Best Practice – Other Factors


- **Design Factors**
  - Applications running on the system must be designed with great care to avoid crashes or other failures.
  - Exception handling, both by administrators and application programs, must be well designed.
- **Deployment factors**
  - Comprehensive testing and staging of the system is paramount to avoiding production downtime.
  - Deployment of a new application into production must be planned and tested carefully to minimise complications.
- **Education factors**
  - System administrators need to be well-trained and dedicated full-time to their systems.
  - Thorough user training will help keep the system performing well, abuse can affect overall system performance.

# Utility Company




User tier

# Central Government Department



User tier

# Real-life Examples – High Street Bank



User tier



## Discussion

### Some topics that we may wish to consider:

- **Resilience versus cost**
  - Need to balance cost of high levels of resilience against business cost of outage
  - True Business cost of an outage
- **Farming versus clustering**
  - IBM recommends farming for P8 servers
  - Scalability



## Introduction – Securing P8 Enterprise Content Management systems

**Stephen Orrell**    ECM Systems Architect

- **Summary of published best practices**
  - Review prior to implementing any P8 ECM system
- **Introduction to the Document Repository Service project**
  - Large UK Government Department strategic ECM system
- **Detailed look at user authentication within the DRS**
  - Real world example of P8 ECM security best practices put into effect
- **Discussion**

## Securing P8 Enterprise Content Management systems - Published Best Practices

- **Access, integrity, privacy and verification**
  - Understand which aspects of each are within the realm of P8 and which lie outside
  - Fully understand the ACL and ACE model and how it is utilised in P8 ECM prior to doing any design work
- **Provide a single source repository of record for authentication information**
  - Security service and security service federation
- **Develop an Enterprise wide standard for access security**
  - Approved JAAS login modules for all J2EE applications

## Securing P8 Enterprise Content Management systems - Published Best Practices (2)

- **Limit physical access to hardware hosting P8 ECM**
  - Authorised sysadmins only
- **Always encrypt network traffic across untrusted networks**
  - Consider requirements for encryption on internal networks
- **Isolate the security domain into its own silo during development and testing**
  - Utilise development to test principals and authentication mechanisms that are separate from the production systems

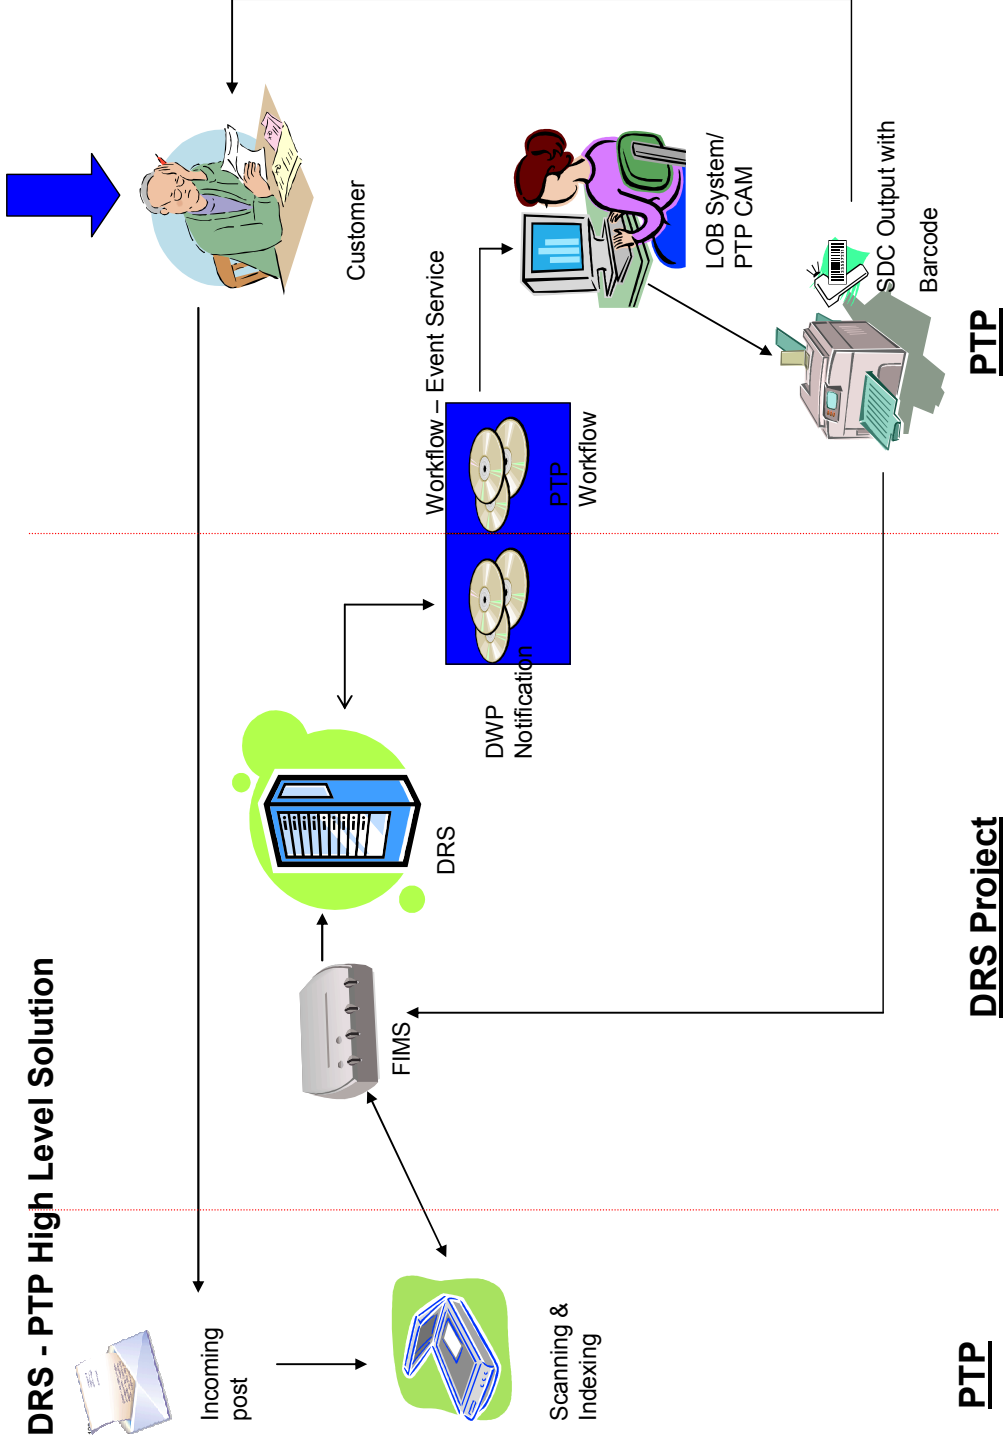
## Securing P8 Enterprise Content Management systems - In Practice... Large UK Government department

### DRS project – The Challenges

- **DRS – Document Repository service**
  - Strategic initiative for Content Management
  - Content originates from multiple sources at high volumes
  - Accessed by multiple Business Units and Projects requiring content management services
- **First Business Unit**
  - 3,000 Siebel CRM users
  - Repository required for customer correspondence
  - ECM system must be ‘invisible’ to CRM users
  - DRS is strategic for the whole department



# Securing P8 Enterprise Content Management systems - In Practice... Large UK Government department (2)

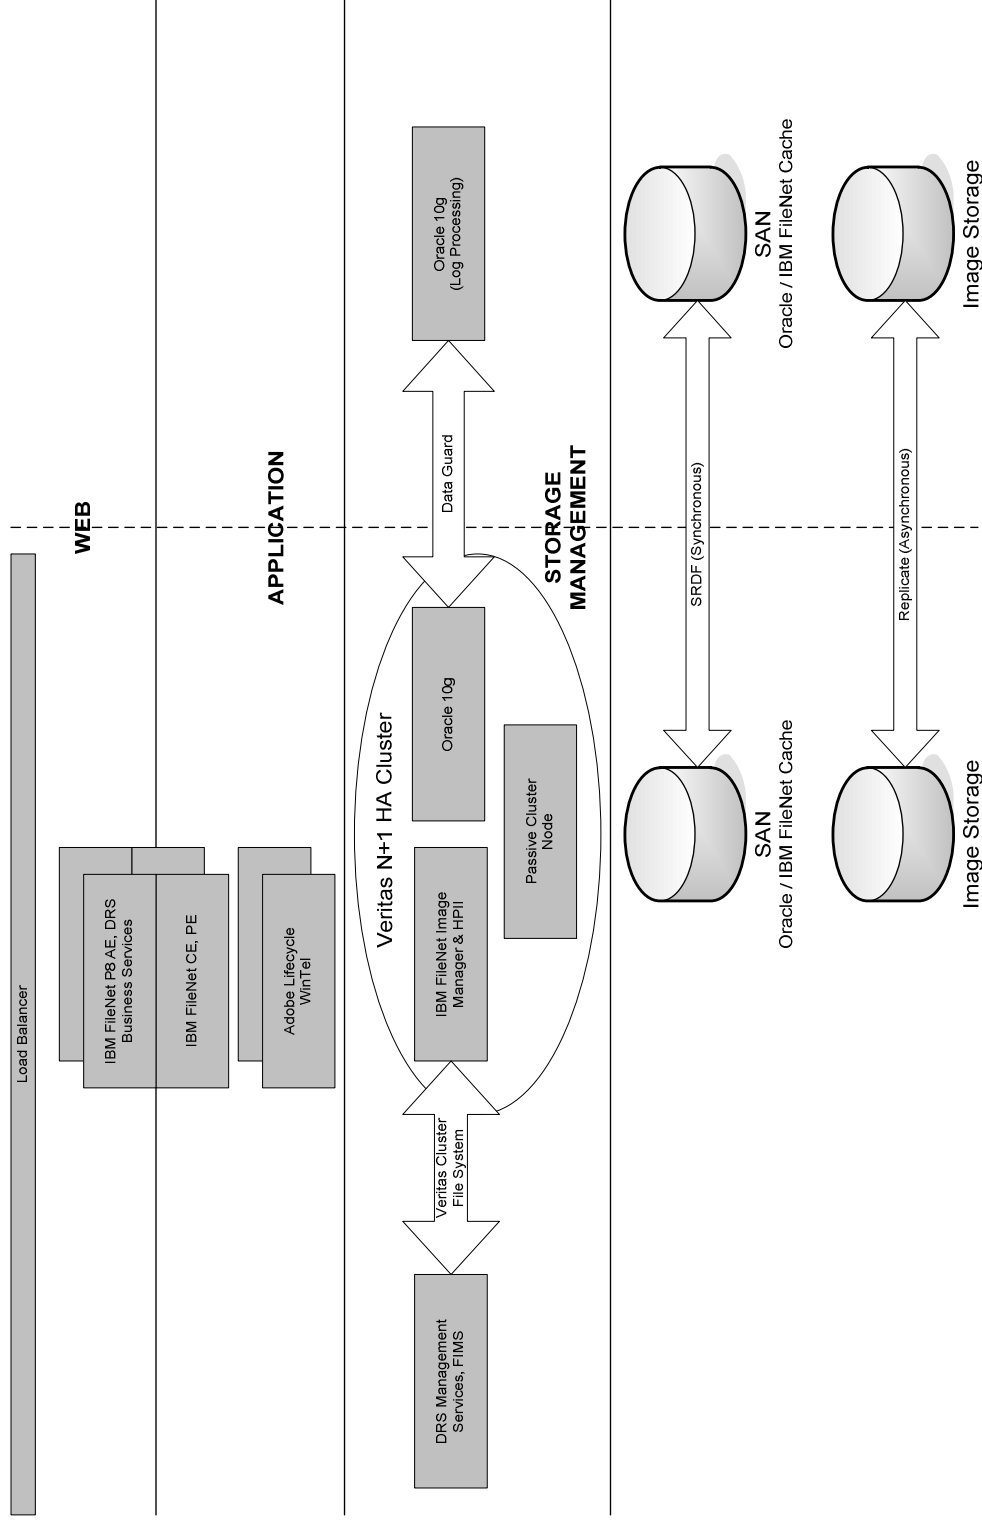






# Securing P8 Enterprise Content Management systems - In Practice... Large UK Government department (3)

DRS Pilot Architecture... Active – Passive Production – DR, 24 hour recovery




## Securing P8 Enterprise Content Management systems - In Practice... Large UK Government department (4)

### DRS Security Model

- **Single sign on**
  - Users log on via their line of business application
  - Desktop smart cards and PINs
- **Service Oriented Architecture**
  - Access to DRS content management only available via a published set of business web services
- **Design principles**
  - IBM FileNet out of the box components will be used wherever possible
- **Authorisation**
  - Document sensitivity
  - Separate business units accessing shared service
  - Takes place against corporate directory

# Securing P8 Enterprise Content Management systems - In Practice... Large UK Government department (5)



# Discussion

## Topics we may wish to consider...

- **Single sign on for P8 based applications**
  - What is supported and what isn't
- **How to enforce security in an ECM shared service**
  - Different security models
- **Security in a Service Oriented Architecture**
  - Securing web service interfaces
- **Auditing system access**
  - Out of the box or custom audit
  - What additional auditing may be required



## Open Discussion

- **The Upgrade, Moving to P8 4.0**
- **Application Integration**
- **AOB?**





# Appendix



# Topics

- P8 4.0 Architecture Changes
- Recommended Deployment Scenarios
- Upgrade Scenarios
- 3rd party Product Support

# Changes in P8 4.0 Affecting HA/DR

- **Content Engine Changes**
- **Process Engine Changes**
- **Platform support changes**
  - **Unix/Linux**
  - **File Stores → Unix file systems and NFS**

## Content Engine Changes: J2EE and Full Farming

- **CE App Server single J2EE deployment**
  - Simplifies installation and deployment
- **Farming CE Service**
  - Object Store Service, File Store Service, Cache Service become single CE Service
  - Includes CFS-IS
  - No requirement for File Store Service cluster
  - Simplifies HA configuration
- **Global Configuration Data (GCD) moved to database**
  - Simplifies DR configuration

## Platform Changes: What About Unix and CE?

- **Extended platform combinations**
- **NFS**
- **File Store architecture stays the same**
  - **Do not confuse “File Store”, File Store Service”, “File Server”, “File systems”**



## Process Engine Changes: Farming

- **Farm nodes run in the native operating system**
- **Single Shared Database Instance**
- **No dependency between PE farm members**
  - Exception is during workflow transfer
- **Homogenous PE farm**
  - Same platform and physical architecture
- **External load balancer required**

# Comparing 3.5 and 4.0 HA/DR Capabilities

Component	P8 3.5	P8 4.0
Application Engine	Farmed	Farmed
Object Store Service	Farmed/clustered	CE service: farmed
File Store Service	Clustered	CE service: farmed
File Server	Clustered or NAS	Clustered or NAS
Cache	Not shared; not HA	CE service: farmed; see file server
CFS-IS	Clustered (if separate FSS)	CE service: farmed
Database Service	Clustered; Oracle RAC: farmed	Clustered; Oracle RAC: farmed
Process Engine	Clustered	Farmed/clustered
GCD	File system: clustered	Database
Image Services	Clustered	Clustered

Clustered: active/passive      Farmed: active/active...active

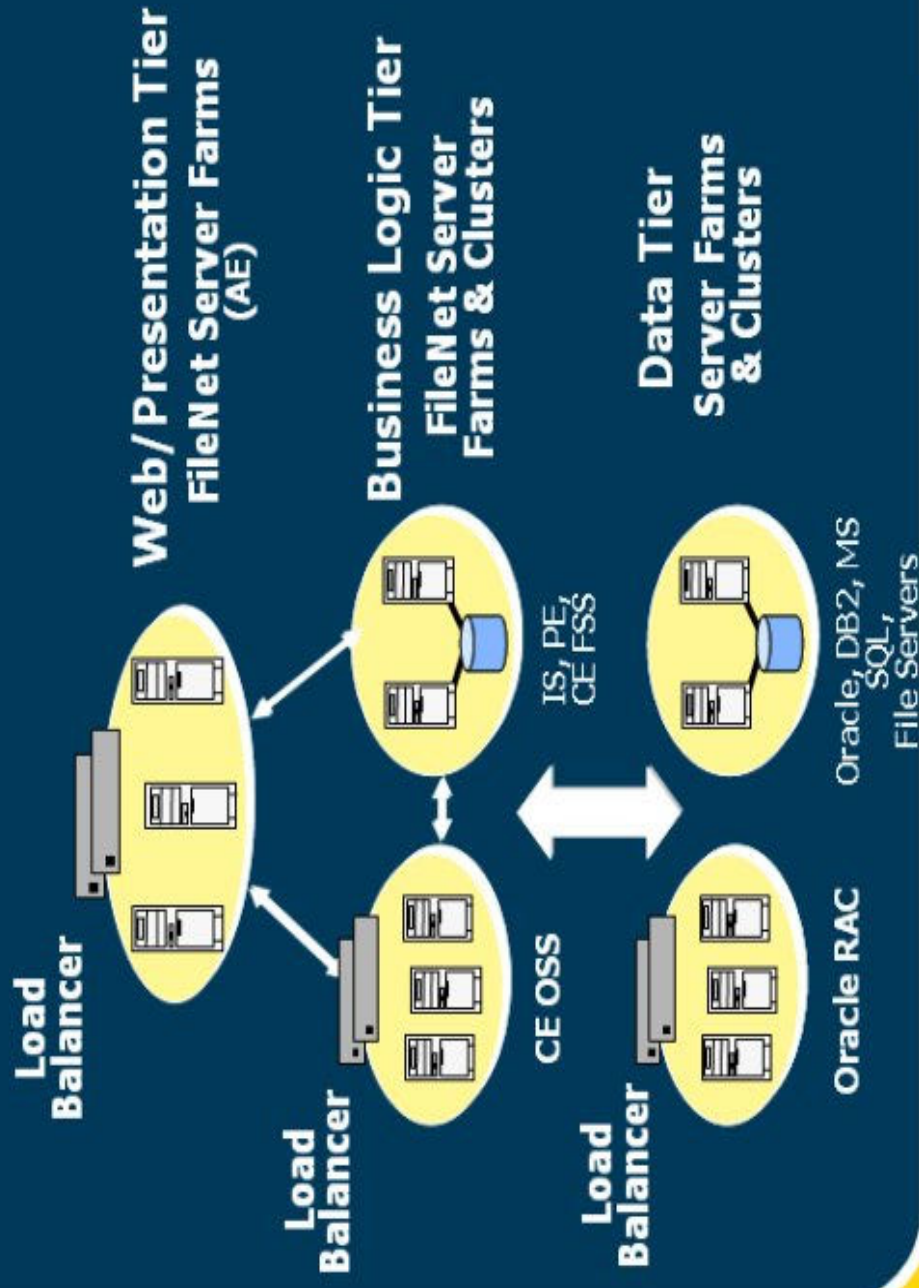
# Farming and Load Balancing

Load balancing technology fronting...					
	AE	CE	PE	CE/PE	RAC
Stateless	No	Yes	Yes	Yes	Built-in
HW	Yes	Yes	Yes	Yes	No
SW	Yes*	Yes	Yes	Yes	Built-in
* Reconnection required in case of failover					





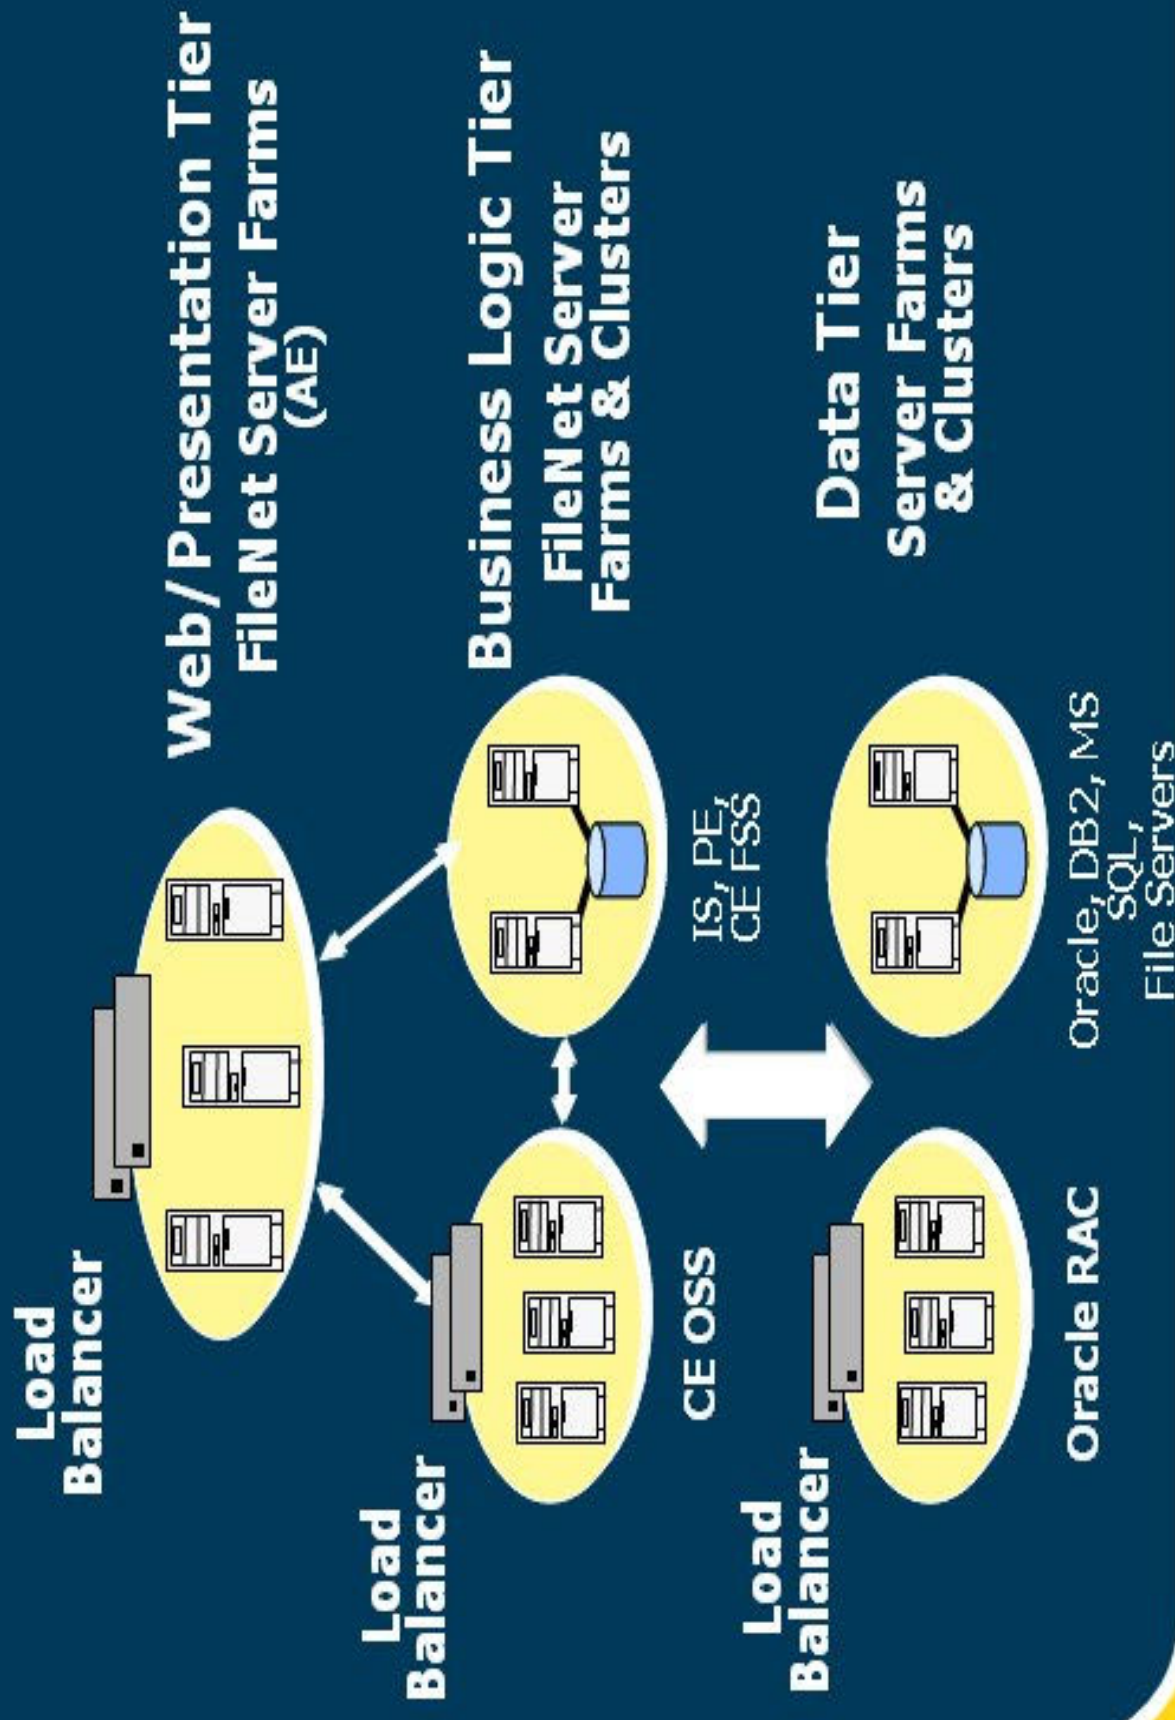
# HA/DR Best Practices for FileNet P8 3.5



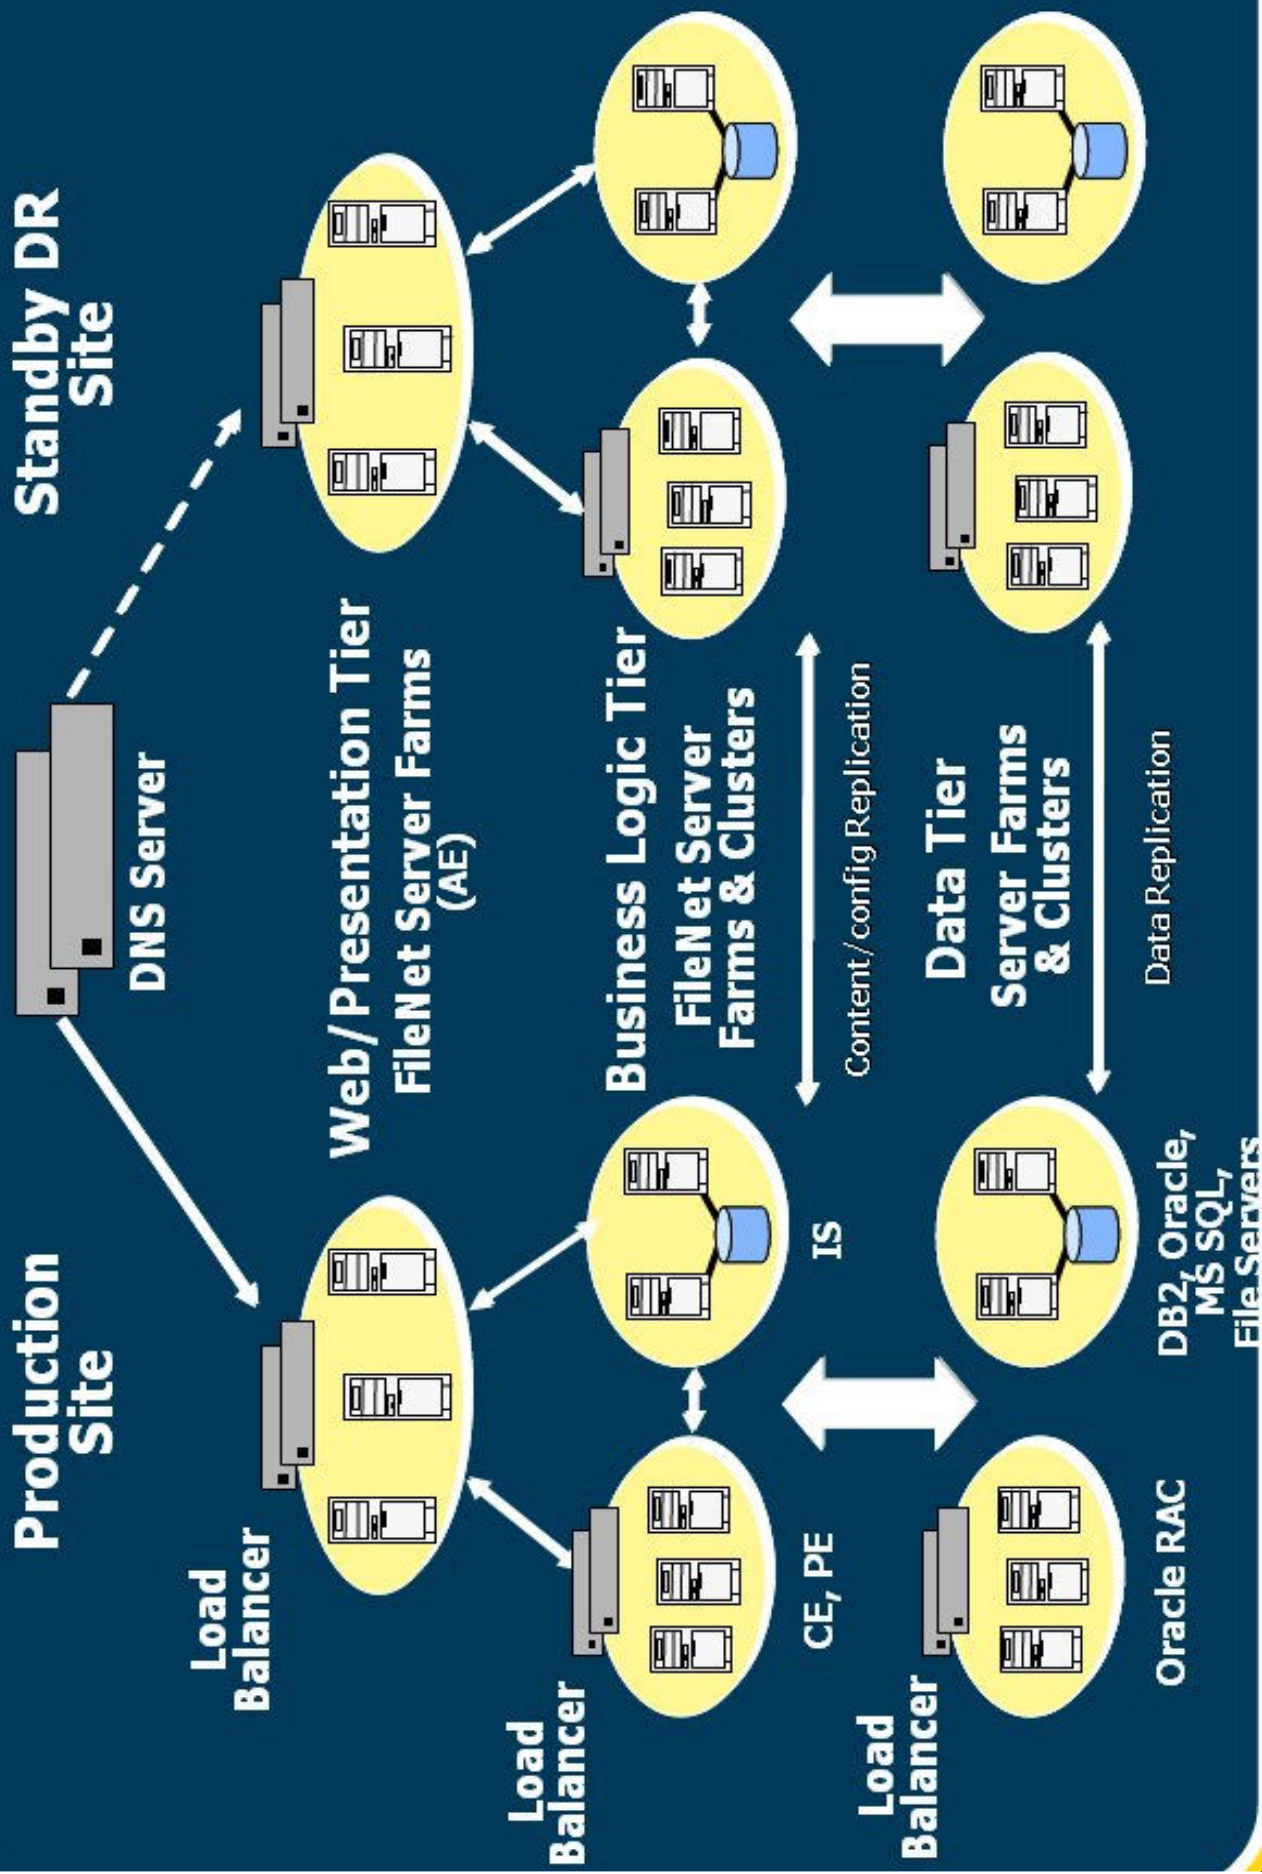
Component	P8 3.5	P8 4.0
Application Engine	Farmed	Farmed
Object Store Service	Farmed/clustered	CE service: farmed
File Store Service	Clustered	CE service: farmed
File Server	Clustered or NAS	Clustered or NAS
Cache	Not shared; not HA	CE service: farmed; see file server
CFS-IS	Clustered (if separate FSS)	CE service: farmed
Database Service	Clustered; Oracle RAC: farmed	Clustered; Oracle RAC: farmed
Process Engine	Clustered	Farmed/clustered
GCD	File system: clustered	Database
Image Services	Clustered	Clustered

Clustered: active/passive      Farmed: active/active...active

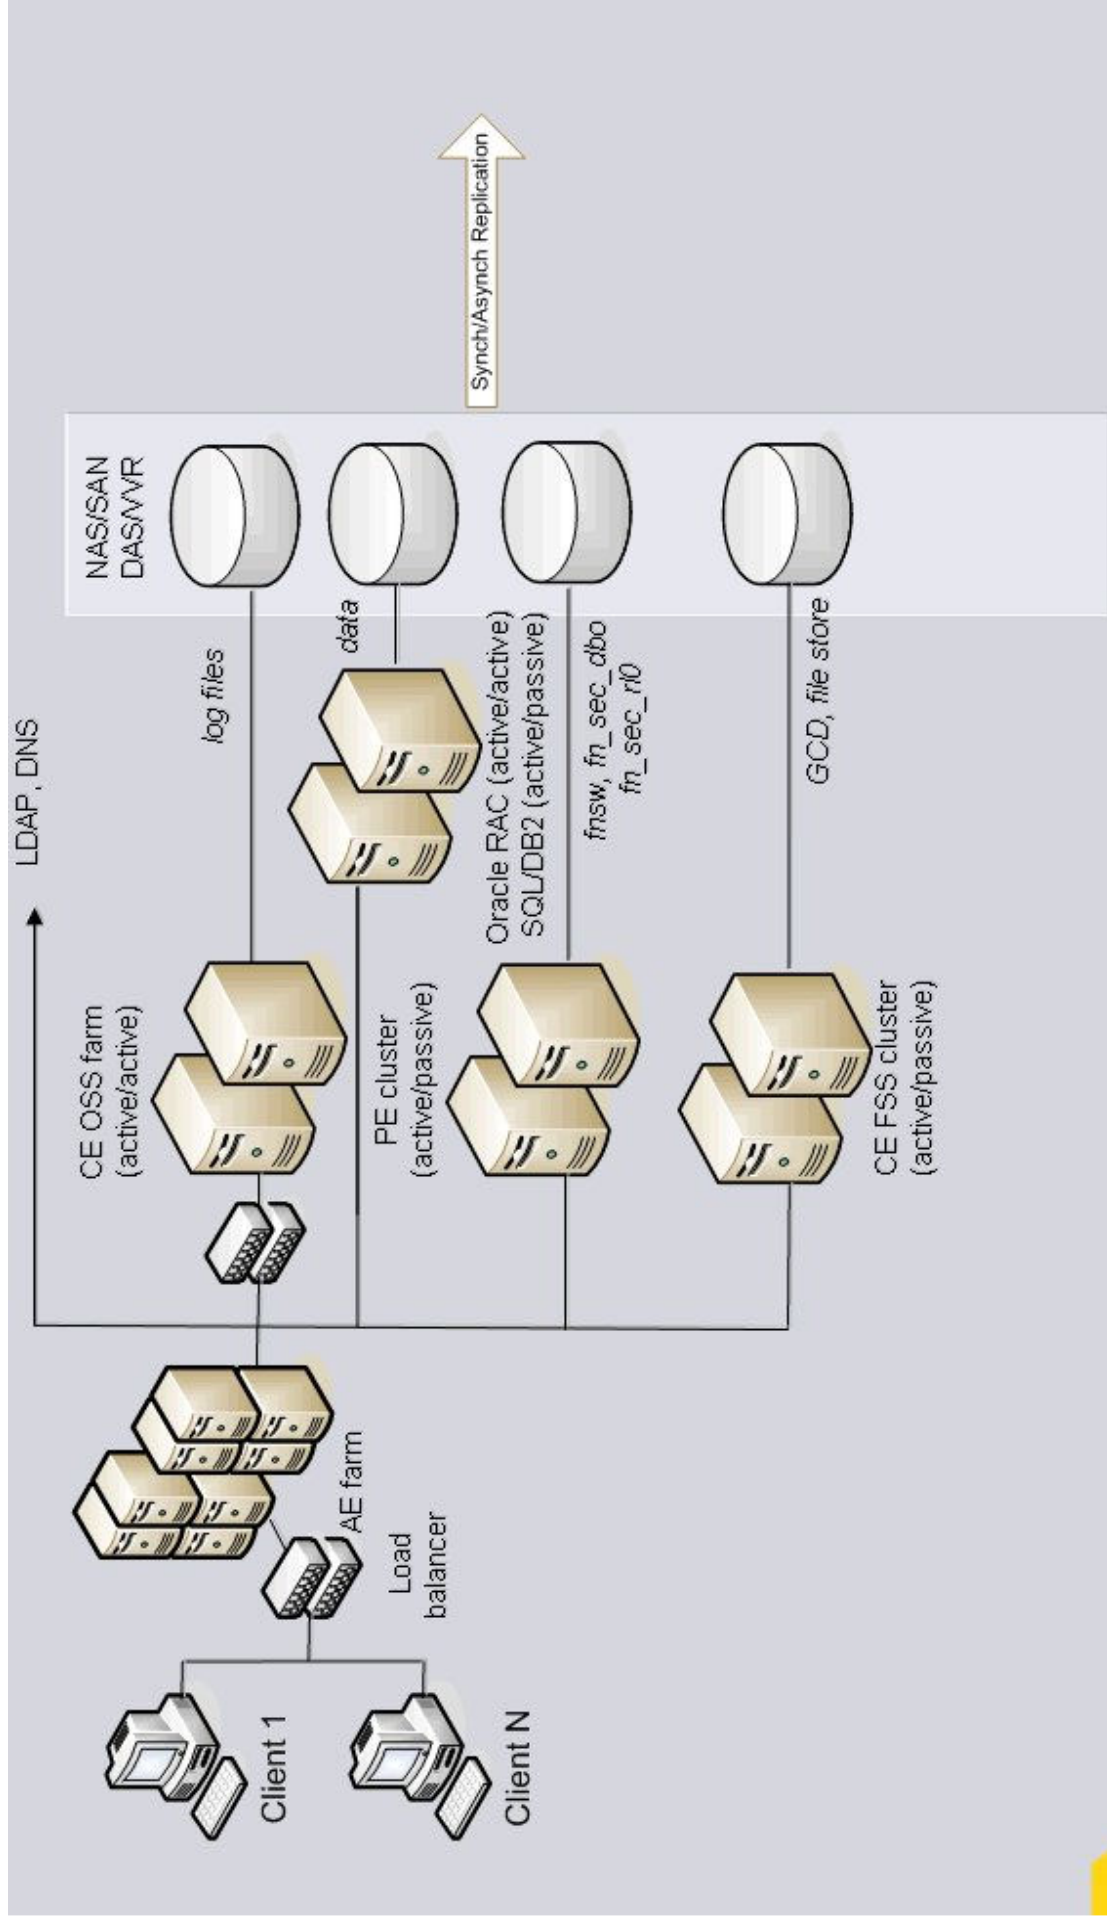
- **Why and where do we still need active/passive clustering?**
  - **File Servers (if not NAS)**
  - **Image Services**
  - **Database Services**
  - **Process Engine (optional)**
  - **Process Analyzer (candidate for post qualified / SP1)**



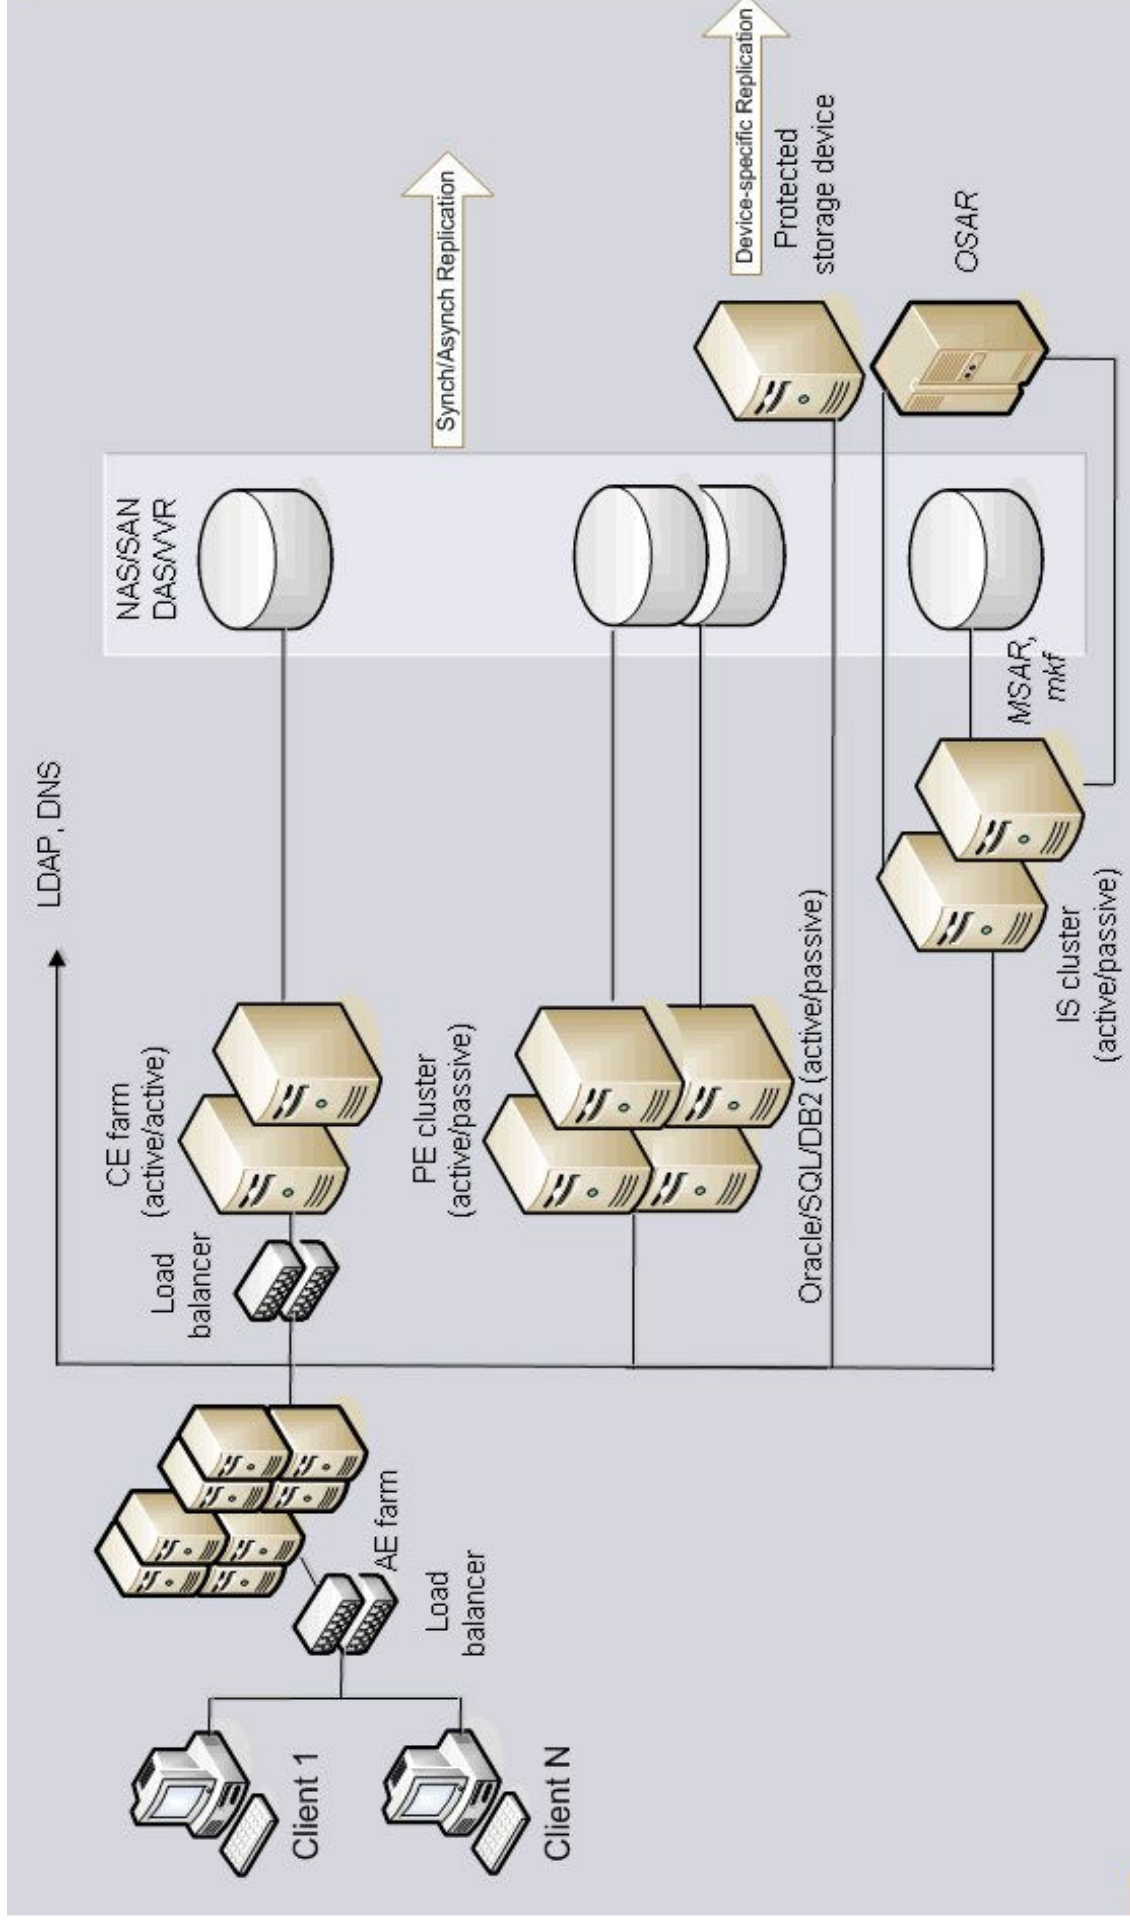




# P8 3.5 HA System

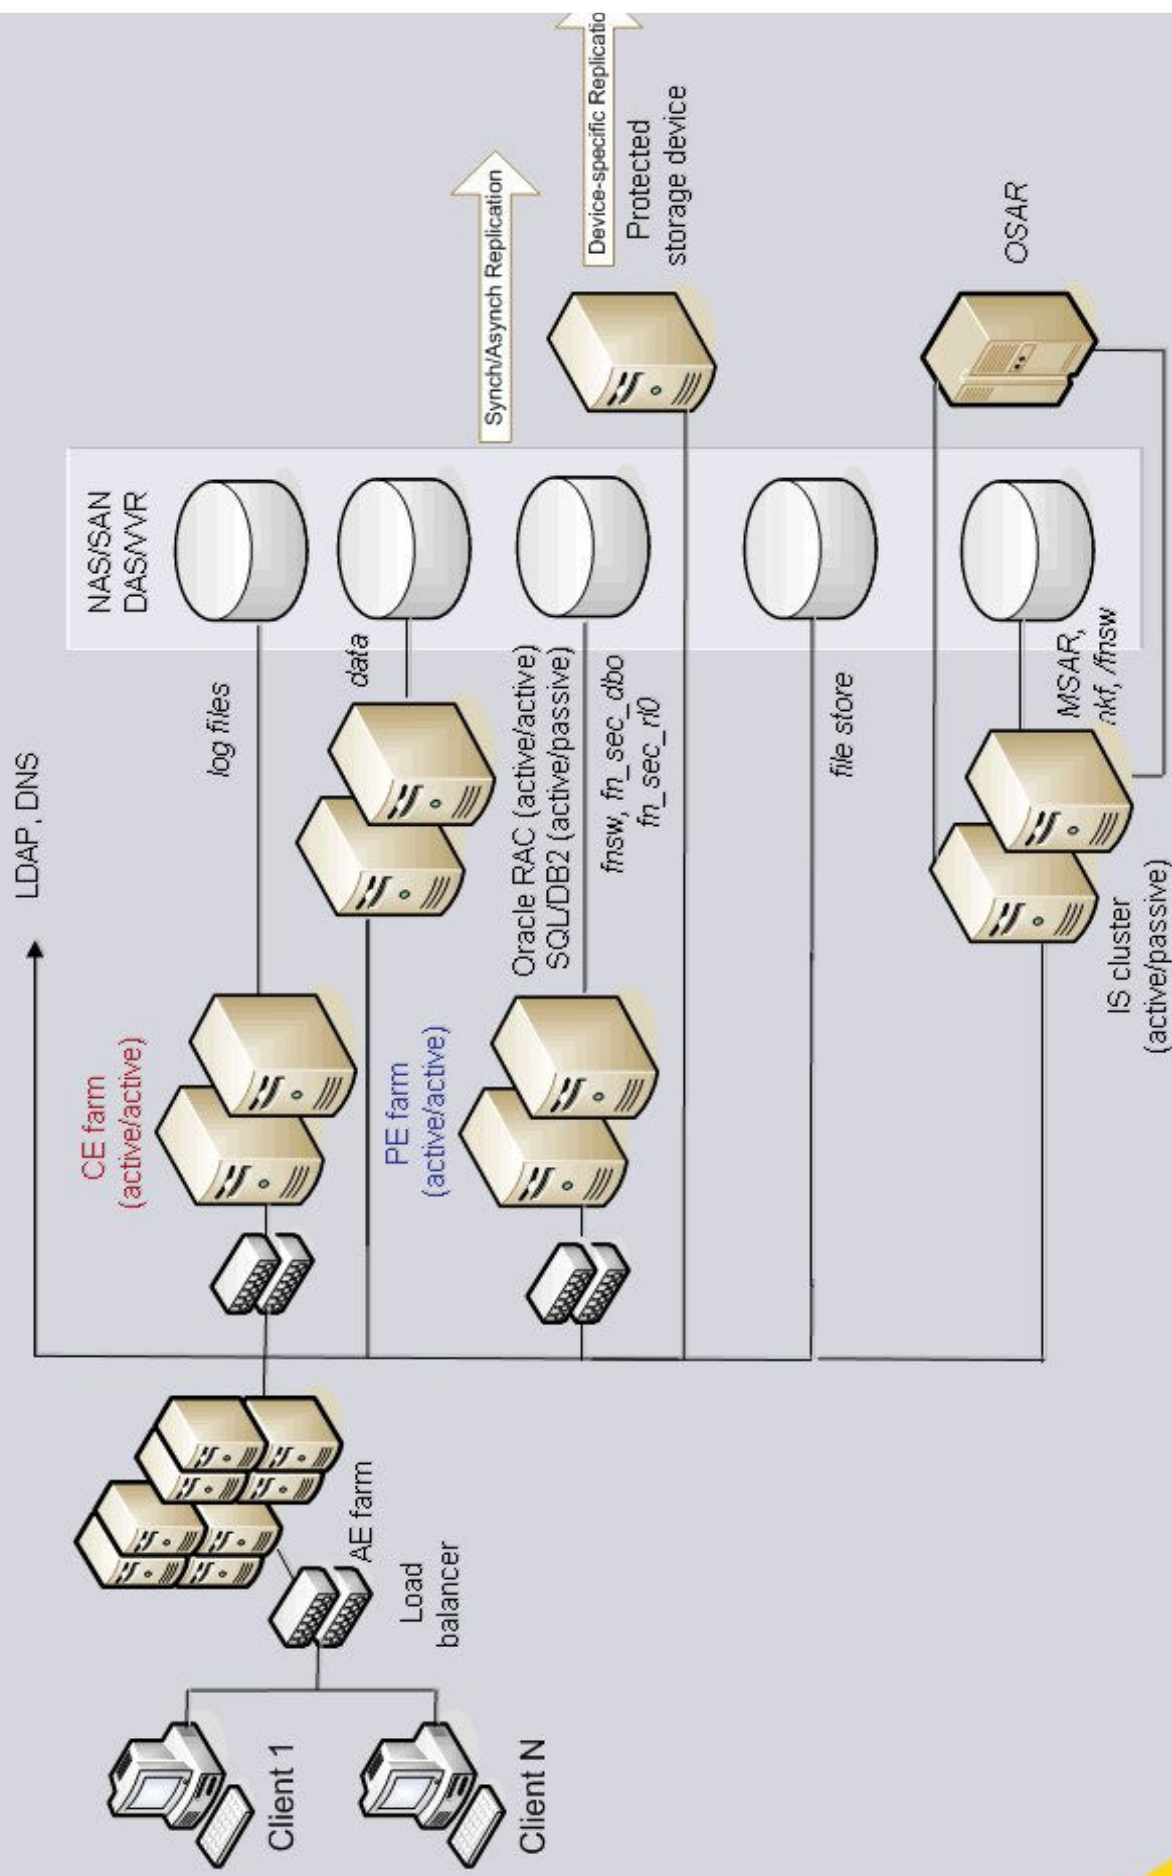


# P8 3.5 HA System - Before Upgrade





**After**





## 3<sup>rd</sup> Party Product Support for P8 4.0 GA

- Working on more flexible support policies for 3<sup>rd</sup> party products
- Load balancing
  - Cisco Local Director
  - WebLogic and WebSphere Load Balancing
- Clustering
  - Veritas Cluster Service (VCS) and custom agent certification for HP-UX, AIX, Solaris, Itanium for PE
  - VCS with database service on Oracle and SQL
  - Microsoft Cluster Service (MSCS) for PE
  - MSCS with database service on SQL server and Oracle
- DB farming
  - Oracle RAC 9i and 10g
- Replication
  - Veritas Volume Replicator (VVR)