

SummerSoc
Service Oriented Computing



Dynamically adapting the deployment topology of
cloud services for minimizing resources
consumption an still meeting service level objective
and performance goals

Advanced School on Service Oriented Computing 30
June – 5 July, 2014, Hersonissos Crete Greece

Cataldo Mega,

IBM Research & Development Laboratory Böblingen, Germany,
cataldo_mega@de.ibm.com

Contents

- **Motivation - will Cloud 3.0 emphasize B2B ?!**
- **Managed Data Services**
- **Challenges in operating Software Solutions in the Cloud**

Software Solutions as a Services in the Cloud

- Cloud implies the **industrialization** of **delivery** for managed **services** “ implying a new consumption and **B-2-B** delivery model inspired by consumer (B-2-C) Internet services.

Cloud aims at:

- Economies-of-scale
- Ubiquity
- Security

In B-B the expectation is:

- **Enterprise ready, replicable standard SaaS offering**
- **Pay per use business model**
- **Security and Privacy i.e. protection of data**



CMaaS – Content Management as a Services

*Using Content Services in the Cloud
for
B-B Scenarios

* IBM Journal of Research and Development Issue 2 • Date March-April 2014

[Dynamic cloud service topology adaption for minimizing resources while meeting performance goals](#)

Mega, C. ; Waizenegger, T. ; Lebutsch, D. ; Schleipen, S. ; Barney, J.M.

Solution Outline Enterprise Content Management



Web, mobile, desktop tools, dashboards and solution applications

Role-based user interface services



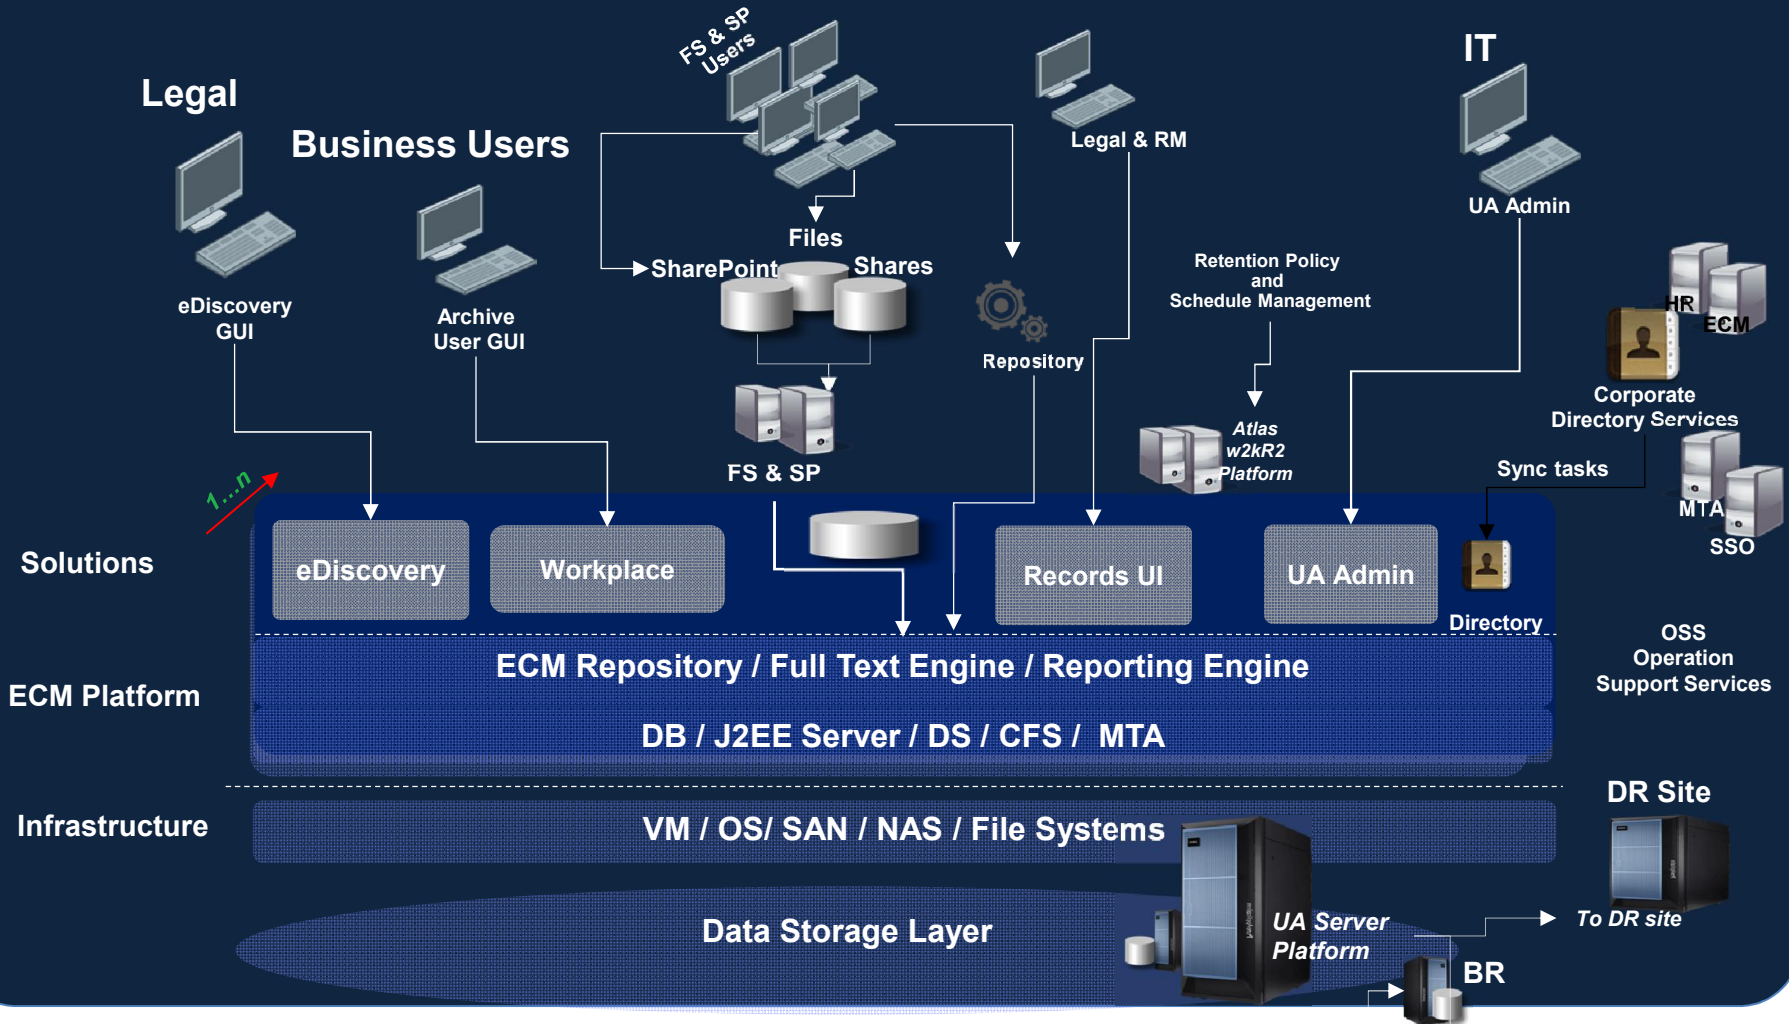
ECM platform services



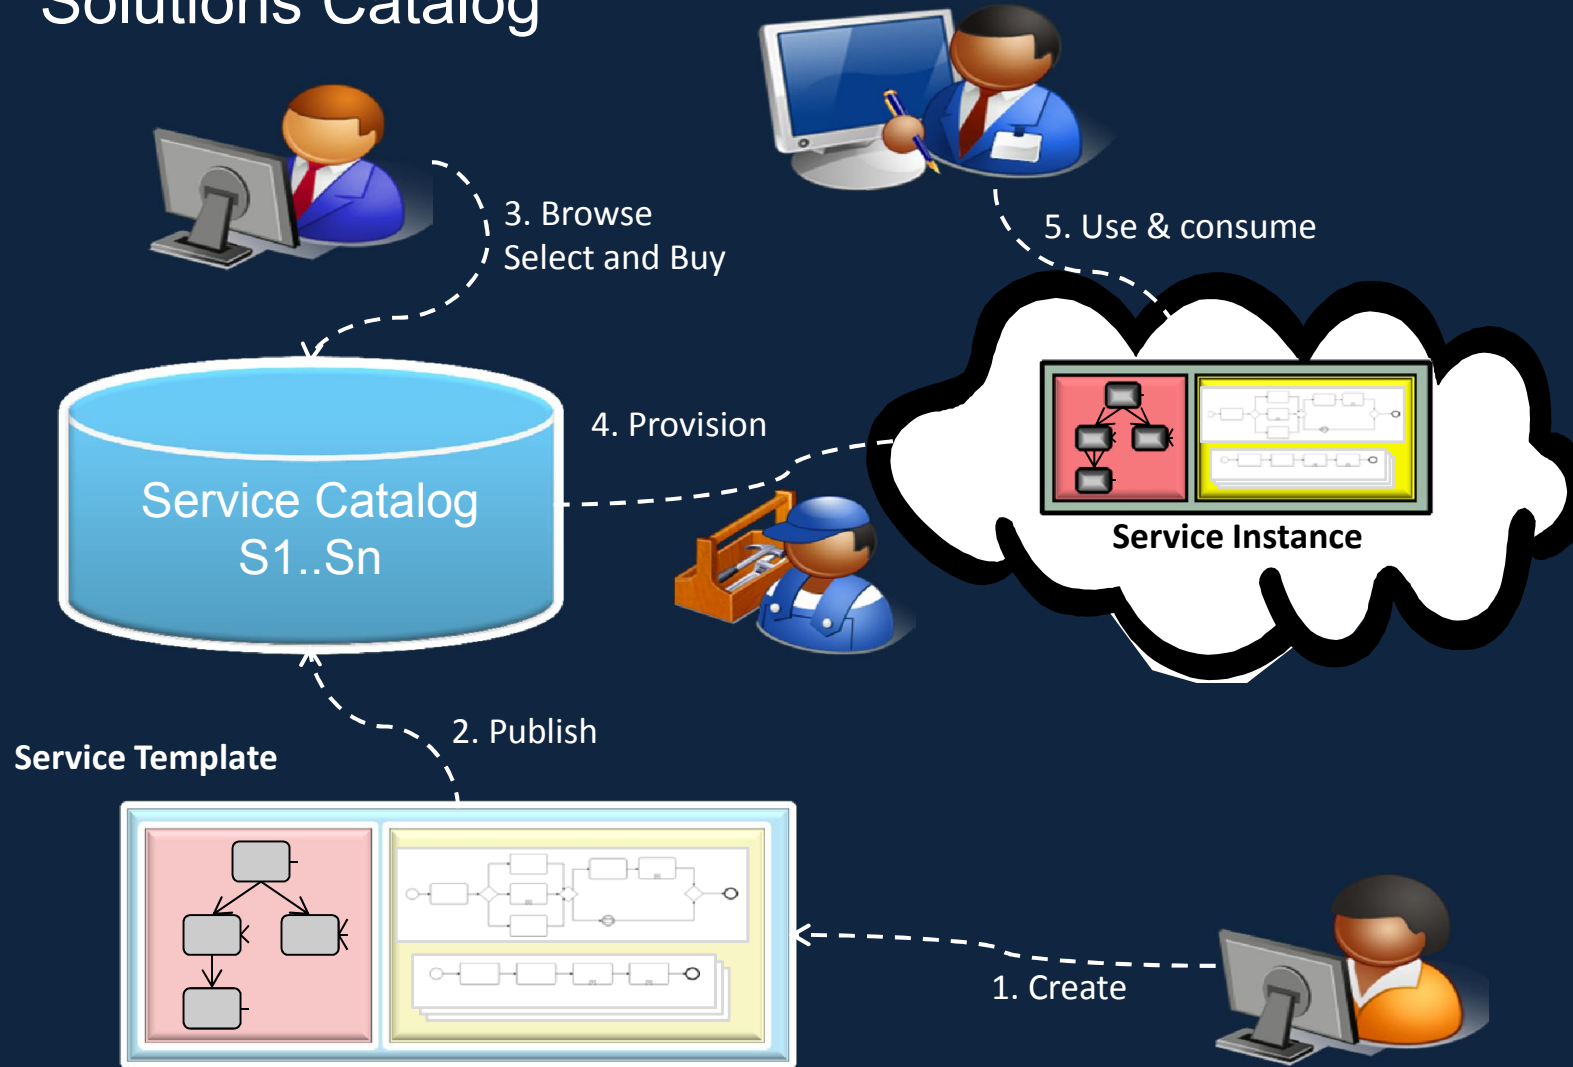
On-premise, appliance, cloud

Unstructured Archive Solution Deployment Architecture Example

Governance & Records Management

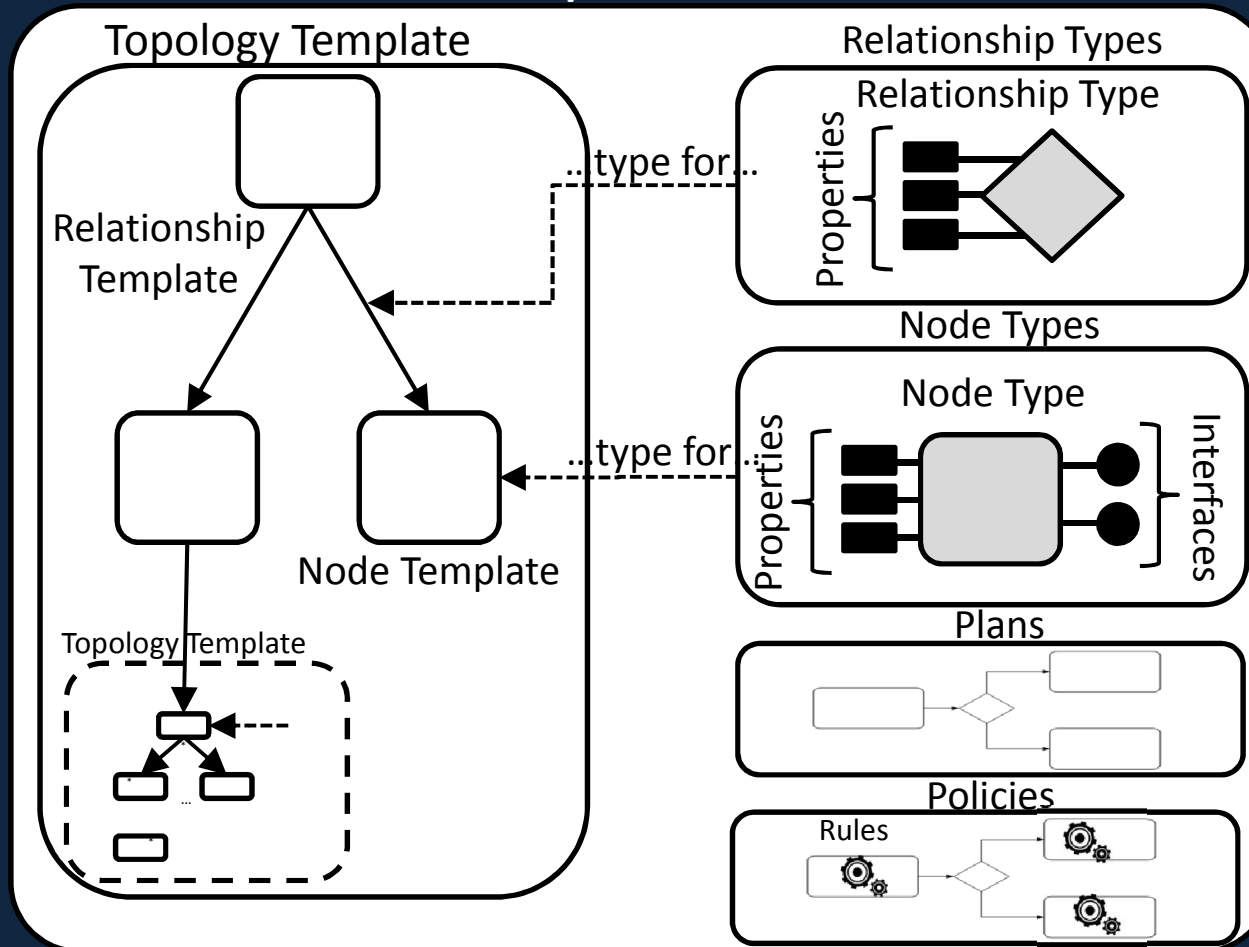


Solutions Catalog



Ingredients of a Software Service Template

Service Template Definition



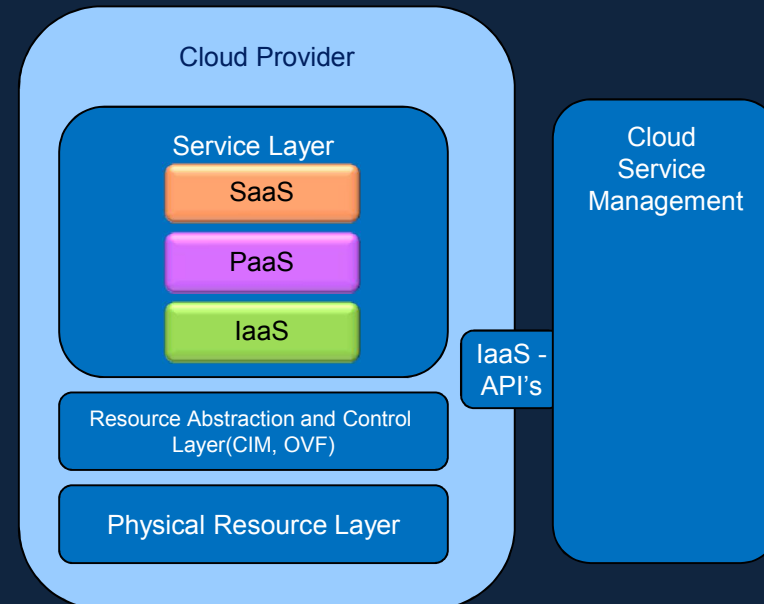
... think about Standards ...

TOSCA - Topology and Orchestration Services for Cloud Applications

How would you ensure the portability of a complex cloud application running on complex software and hardware infrastructure?

Software as a Service →
Platform as a Service - DSL →
Infrastructure as a Service - DSL →

Domain Specific Languages that support automated deployment in the cloud



Steps towards Managed Data Services in the Cloud

Design time:

1. Design the workload model for a given Data Service e.g. ECM solution
2. Develop a service specific workload and resource topology model
3. Define the workload pattern using a formal service definition language
4. Define the plans to provision and de-provision resources as required
5. Define the policies for being able to trigger provisioning & de-provisioning plans

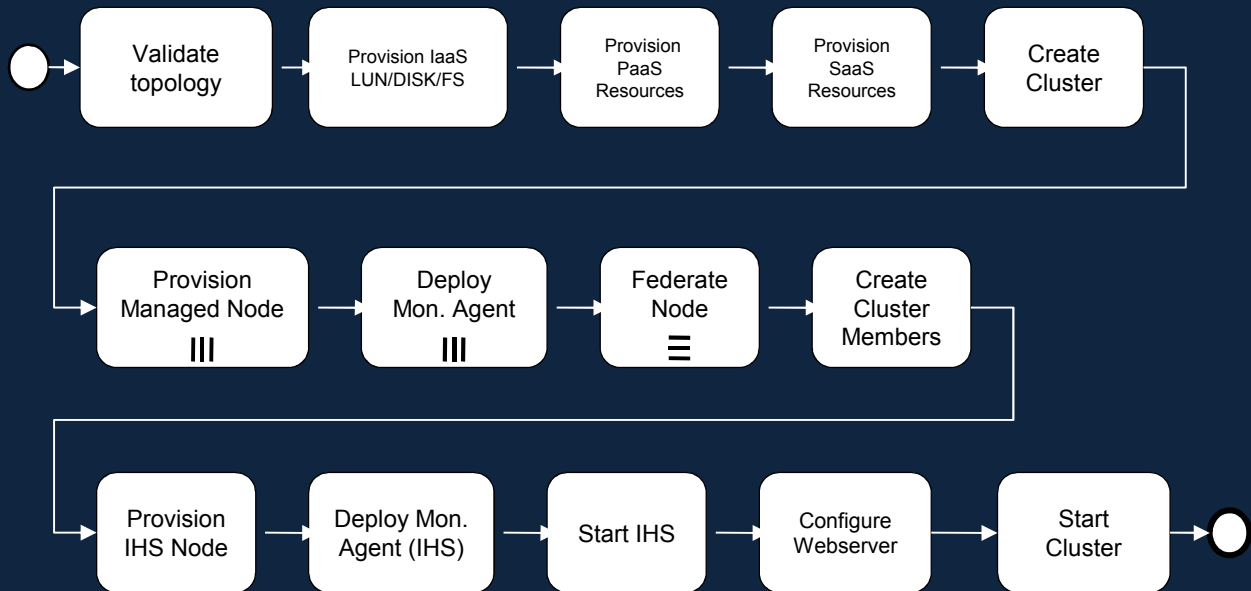
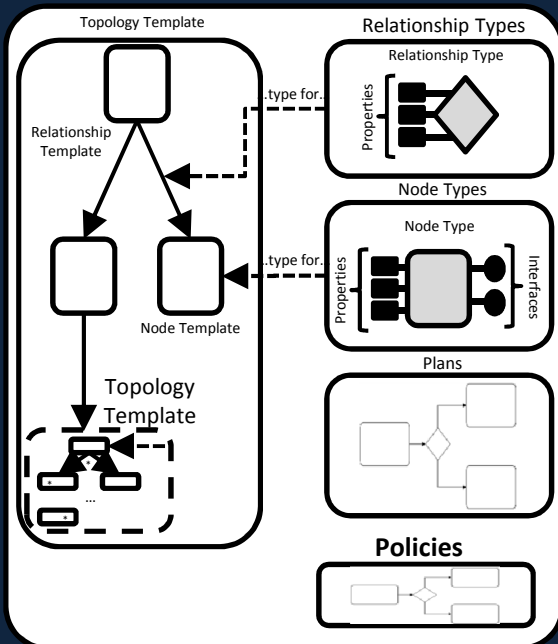
Runtime:

1. Monitor and collect data into a data mart
2. Perform data analysis
3. Predict resource consumption
4. Dynamically explode and implode the resource topology required to sustain the workload at any given time

Service Plans & Workload Patterns

Provisioning


Service Template




Gathering Heuristics

Solution specific Workload

Initial Topology at t1



Change Topology at t2



Policy Catalog



Predictive models


Topology change:
Add Node
Orchestrator



Monitor

Collect

Data Mart






Analyse

Derive Rules

Execute

Rules Engine



Note: Derive rules that drive provisioning flows.
Initial prov 20% -> 80% ..-> 40% dynamic & elastic ...

Service Plans & Workload Patterns

Rules Engine




Measure
Analyse
Plan
Execute


Heuristics
Optimizer

Single tenant topology changes based on current workload

Add Nodes




Monitoring




Add Node

Remove Nodes


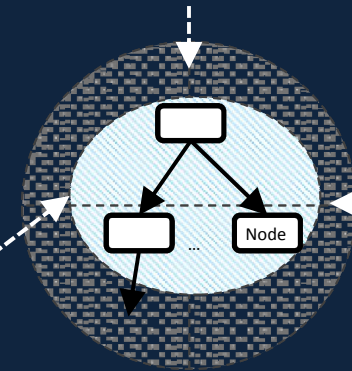


Metering



Remove Node


Elasticity



Service Plans & Workload Patterns Examples

Multi Tenant Topology Changes


Rules Engine




Measure
Analyse
Plan
Execute

Heuristics
Optimizer

Add Tenant


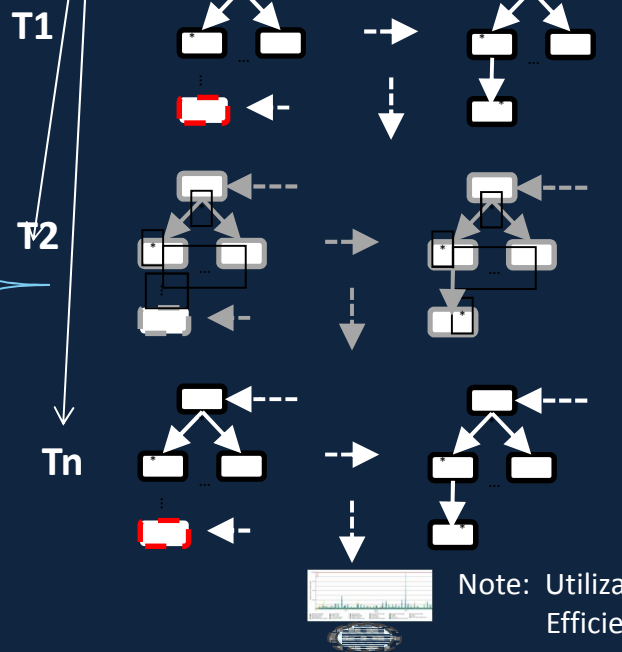


Remove Tenant



BPL Workflow Engine
To orchestrate provisioning flows

Workloads



Note: Utilization = System Resources Utilized / Max System Resources available
Efficiency = Max possible Utilization

ECM Workload Pattern a real world example!

ECM Production Systems Definition

We want to design an ECM production pattern having the following characteristics:

Functions & Capabilities:

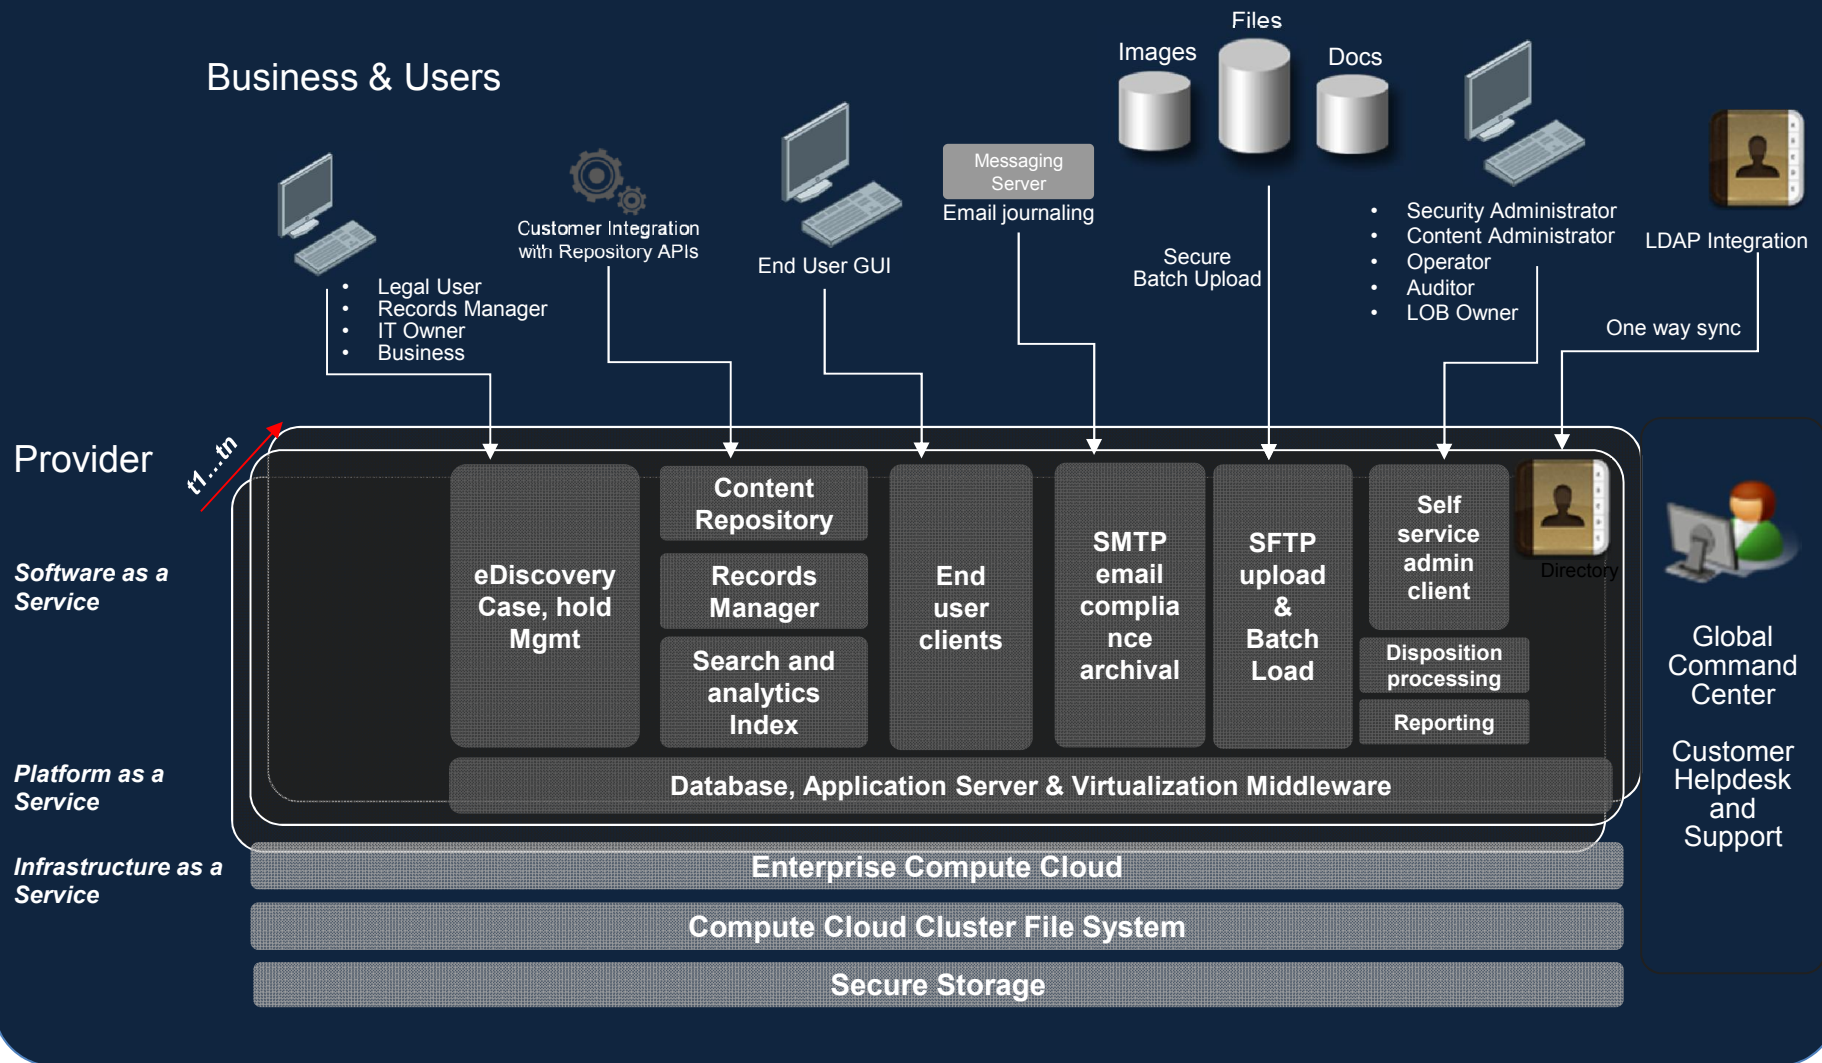
Records and eDiscovery Content Management System

Catalog Data Base, Full Text Index, Persistent Storage

Performance & Scale:

- **Small** is defined in terms of -> **300 concurrent users ; 25K documents per day** with an estimated resource consumption of :
- **Medium** is defined in terms of -> **3000 concurrent users, 250K documents per day** with an estimated resource consumption of :
- **Large** is defined in terms of: **10000 concurrent users, 500K documents per day** with an estimated resource consumption of :
- **Business Continuity:**
 - **Backup&Restore (BR) 3 month, incremental**
 - **High Availability(HA) i.e. 99.9%**
 - **Disaster Recovery (DR) RPO = zero data loss; RTO = 1 day**

Records and eDiscovery Content Management System



ECM PROD Systems Capacity & Sizing Examples

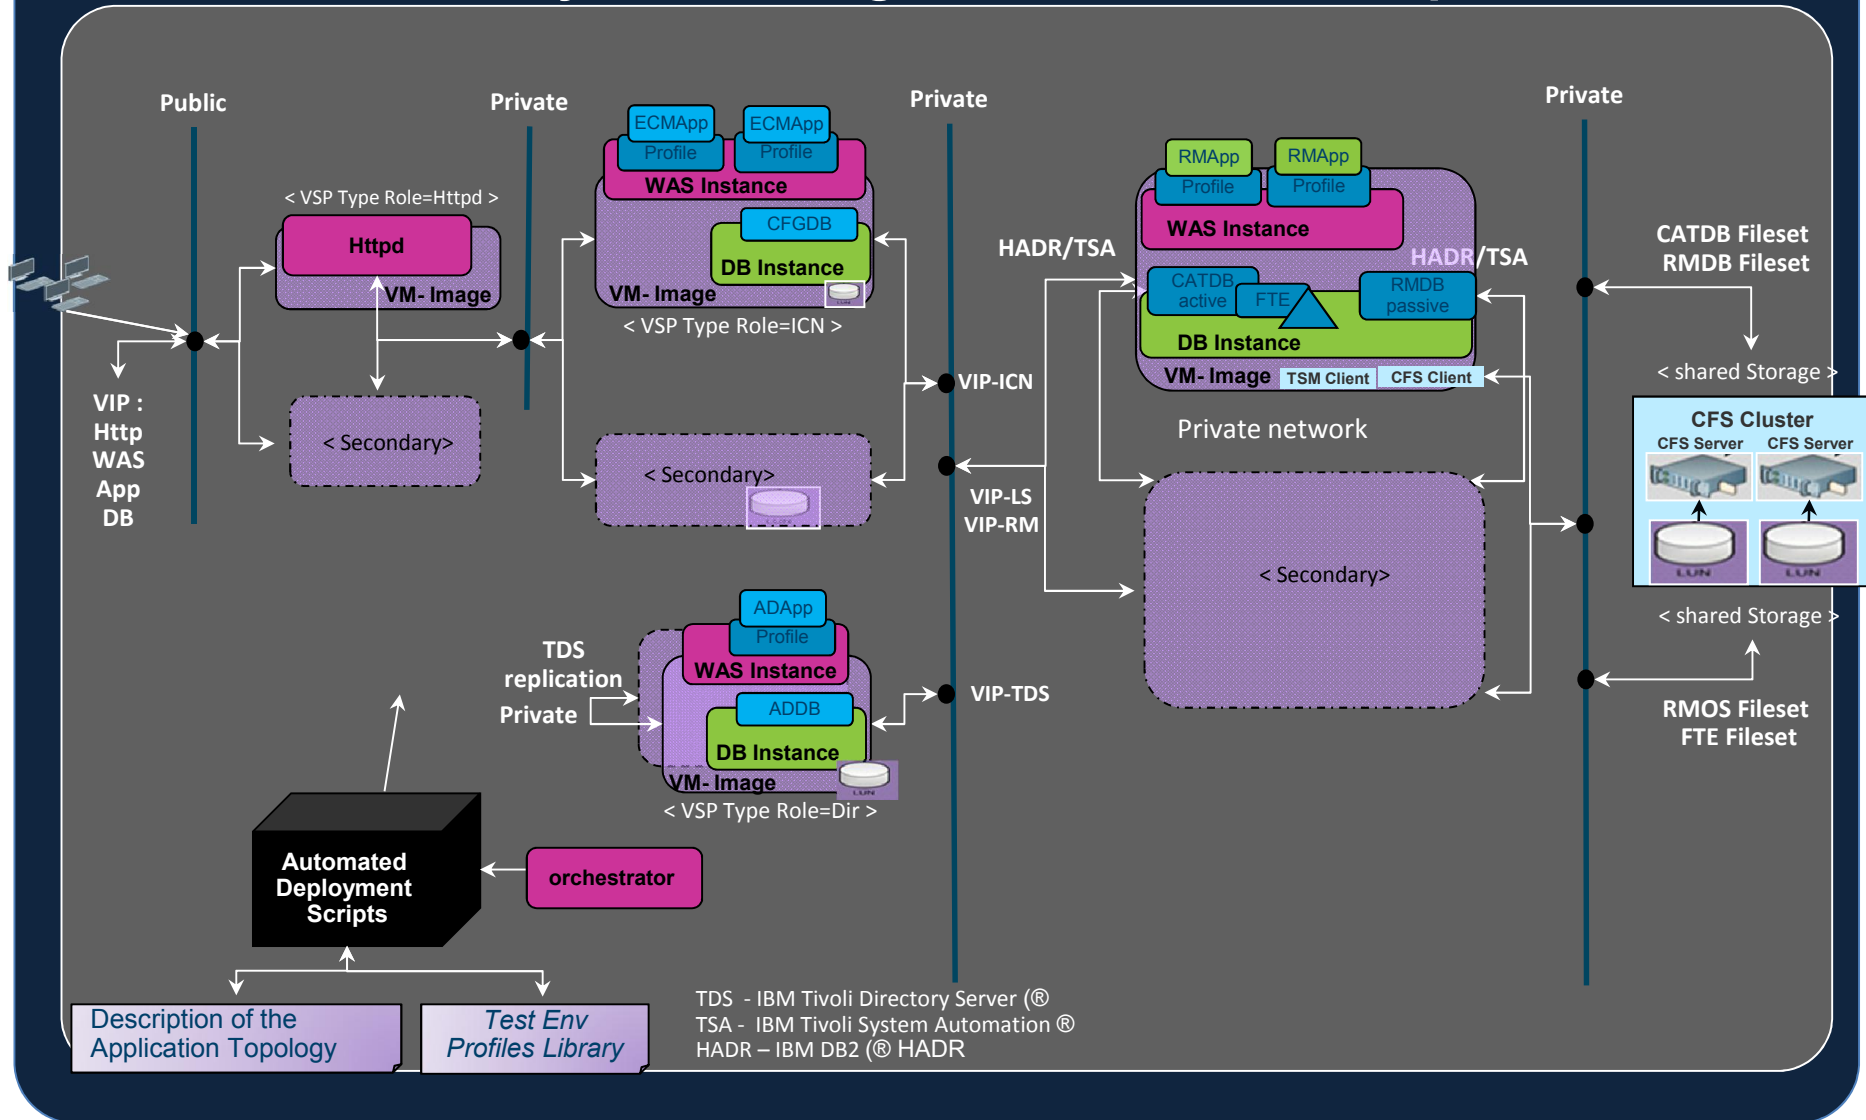
Workload Pattern (VSP)	# VI	# Crs	# RAM	System Capacity & Sizings Figures
PROD – Small w/o HA	4	14	40	A Small PROD System w/o HA is defined in term of: 300 concurrent users; 25,000 documents per day: Configuration: VI: 4x; Capacity: 14xCores, 40RAM, Storage: OS: 64GB, Data: 64GB; Content: 32GB Service level Objectives (SLO): Archiving : 25K doc/d ; 25k doc/d FT indexing; RespoFTE time: CRUDS: 1-3 sec
PROD – Small w HA	8	28	80	A Small PROD System w/o HA is defined in term of: 300 concurrent users; 25,000 documents per day: Configuration: VI: 8x; Capacity: 28xCores, 80 RAM, Storage: OS: 64GB, Data: 64GB; Content: 64GB Service level Objectives (SLO): Archiving : 25K doc/ d : n/a, 25K doc/ d FT indexing; RespoFTE time: CRUDS: 1-3 sec
PROD – Medium	10	24	72	A Medium PROD System is defined in term of: 3000 concurrent users; 250,000 documents per day: Configuration: VI:10x; Capacity: 24xCores, 72xGB RAM, Storage: OS: 64GB, Data: 64GB; Content: 64GB Service level Objectives (SLO): Archiving : 100K docs/ d , 100K docs/ d FT indexing; RespoFTE time: CRUDS: 1-3 sec
PROD – Large	12	24	72	A Large PROD System Large - is defined in term of: 10000 concurrent users; 500,000 documents per day: Configuration: VI:12x; Capacity: 24xCores, 72xGB RAM, Storage: OS: 64GB, Data: 64GB; Content: 64GB Service level Objectives (SLO): Archiving : 0.5m doc / d ;, 0.5m docs/d FT indexing;; RespoFTE time: CRUDS: 1-3 sec
CFS	2	4	8	

Target ECM System Pattern Sizes

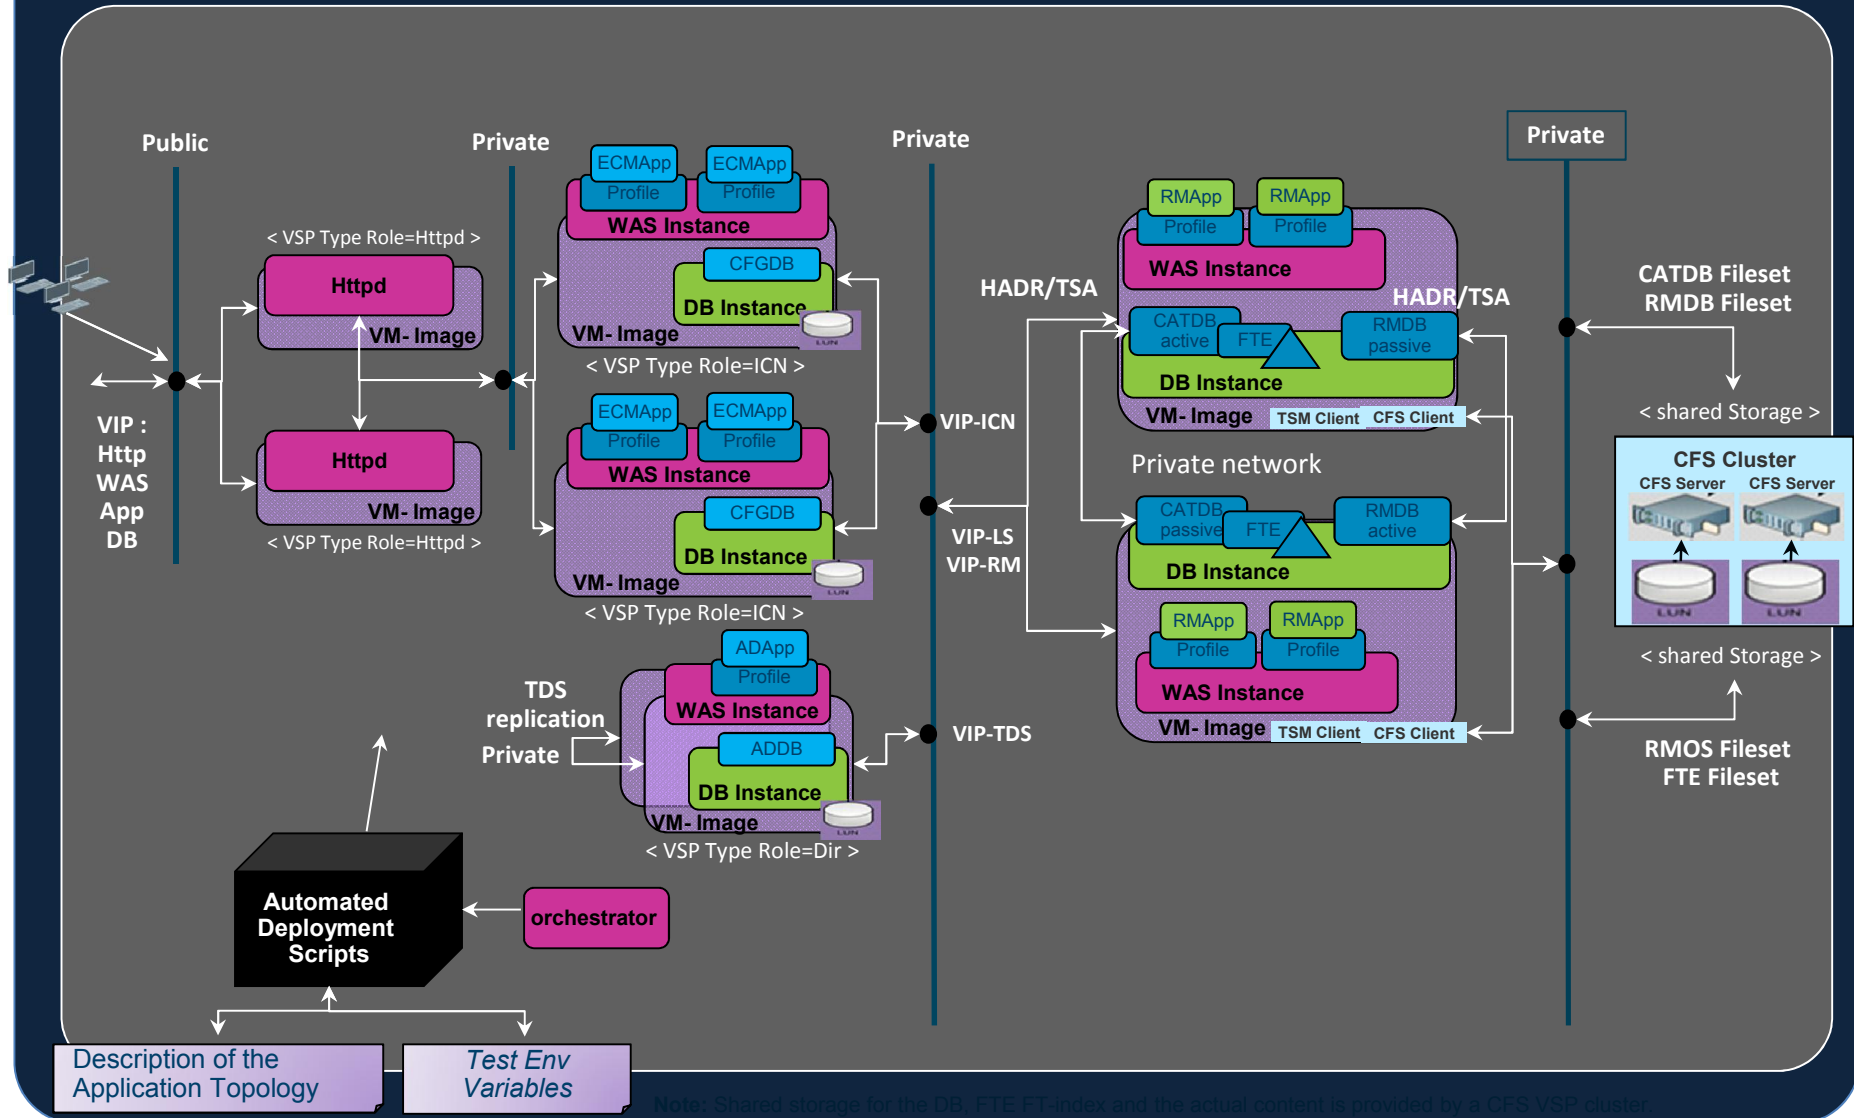
Table shows sample resource consumption figures by ECM system environment.

	Cores	RAM	Storage (GB)					VI	PVU	Private Cloud System			
Custom VI	#	GB	OS-HD	VAR-HD	DB Tbl	Min FT-HD	Min Content-HD	#	#	96	192	384	608
CFS-Storage SRV	4	8	64	64	0	0	500	2	280	24	48	96	152
DEV Sys	4	8	64	64	64	32	32	1	280	24	48	96	152
FVT/SVT Test	10	32	64	64	64	32	32	4	700	10	19	38	61
UAT Test w/o HA	12	36	64	64	64	32	32	5	840	8	16	32	51
UAT Test w HA	24	72	64	64	64	32	32	10	1680	4	8	16	25
SIT Test HAP	24	72	64	64	64	32	32	12	1680	4	8	16	25
Small Prod w/o HA	14	40	64	64	64	128	512+	4	980	7	14	27	43
Small Prod w HA	28	80	64	64	64	128	512+	8	1960	3	7	14	22
Medium Prod w HA	24	72	64	64	64	128	5000+	10	1680	4	8	16	25
Large Prod w HA	24	72	64	64	64	128	10000+	12	1680	4	8	16	25

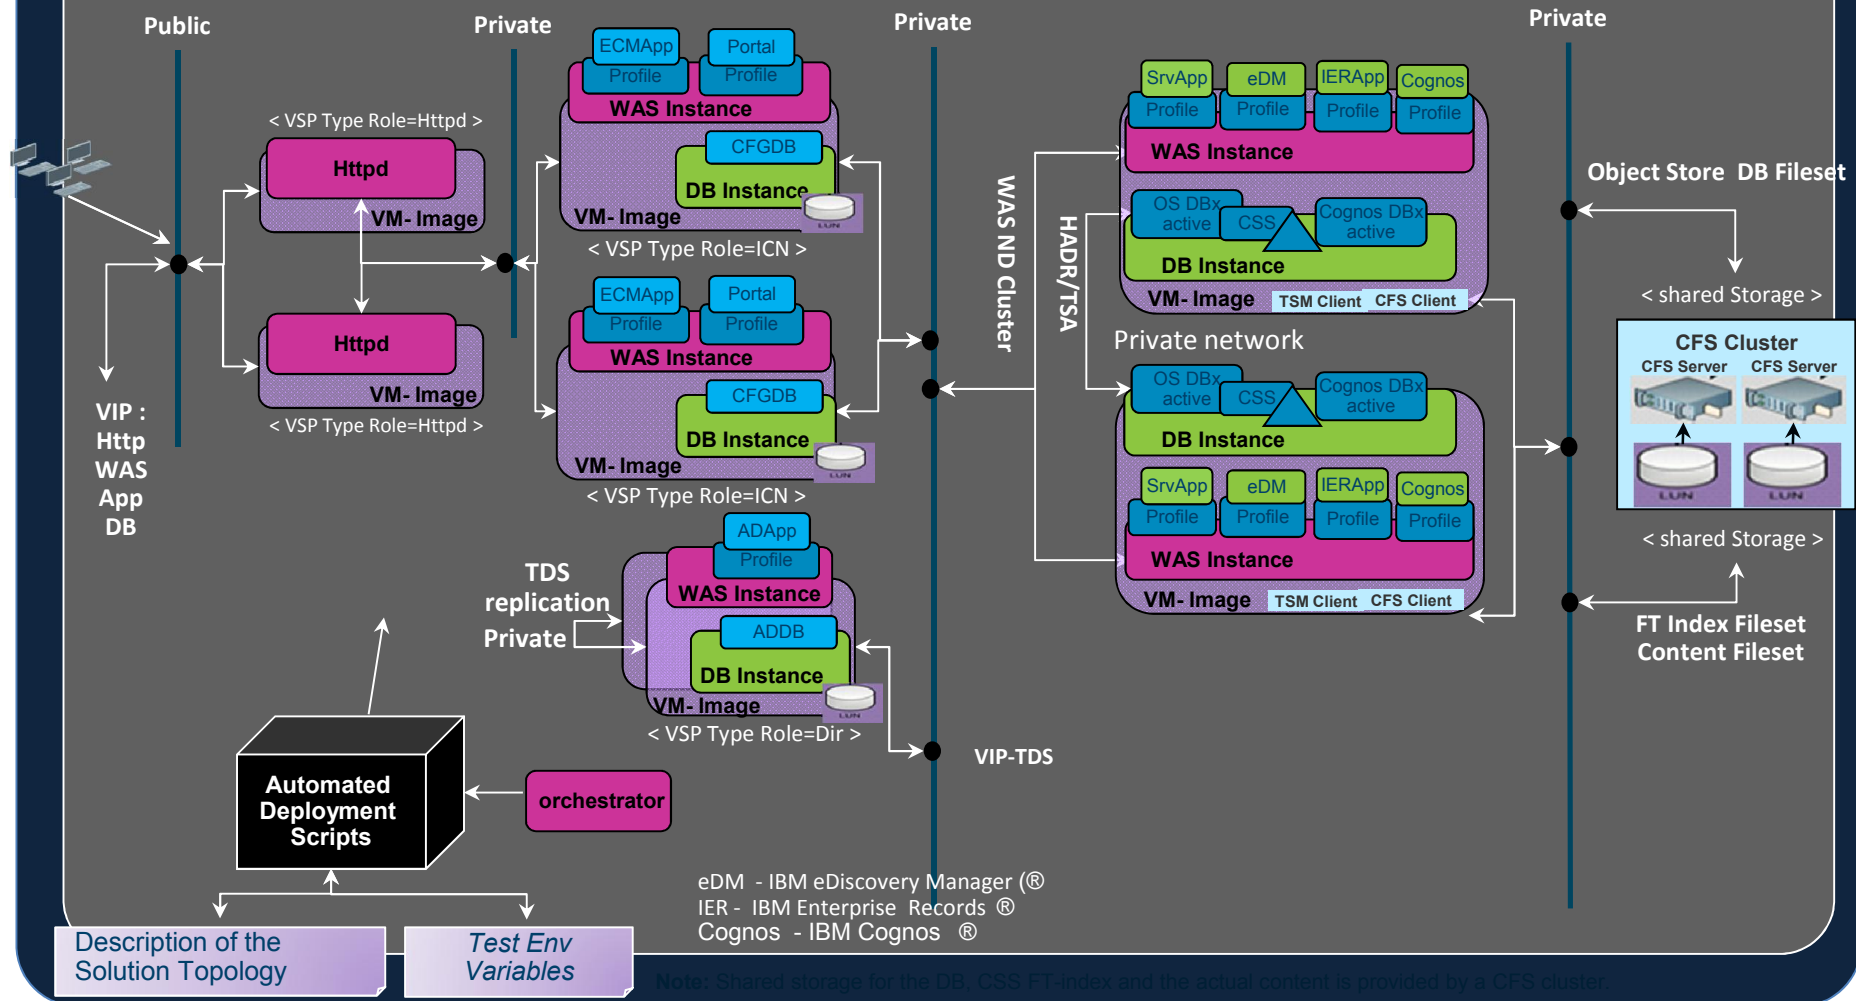
ECM Small PROD System Configuration w/o HA Example



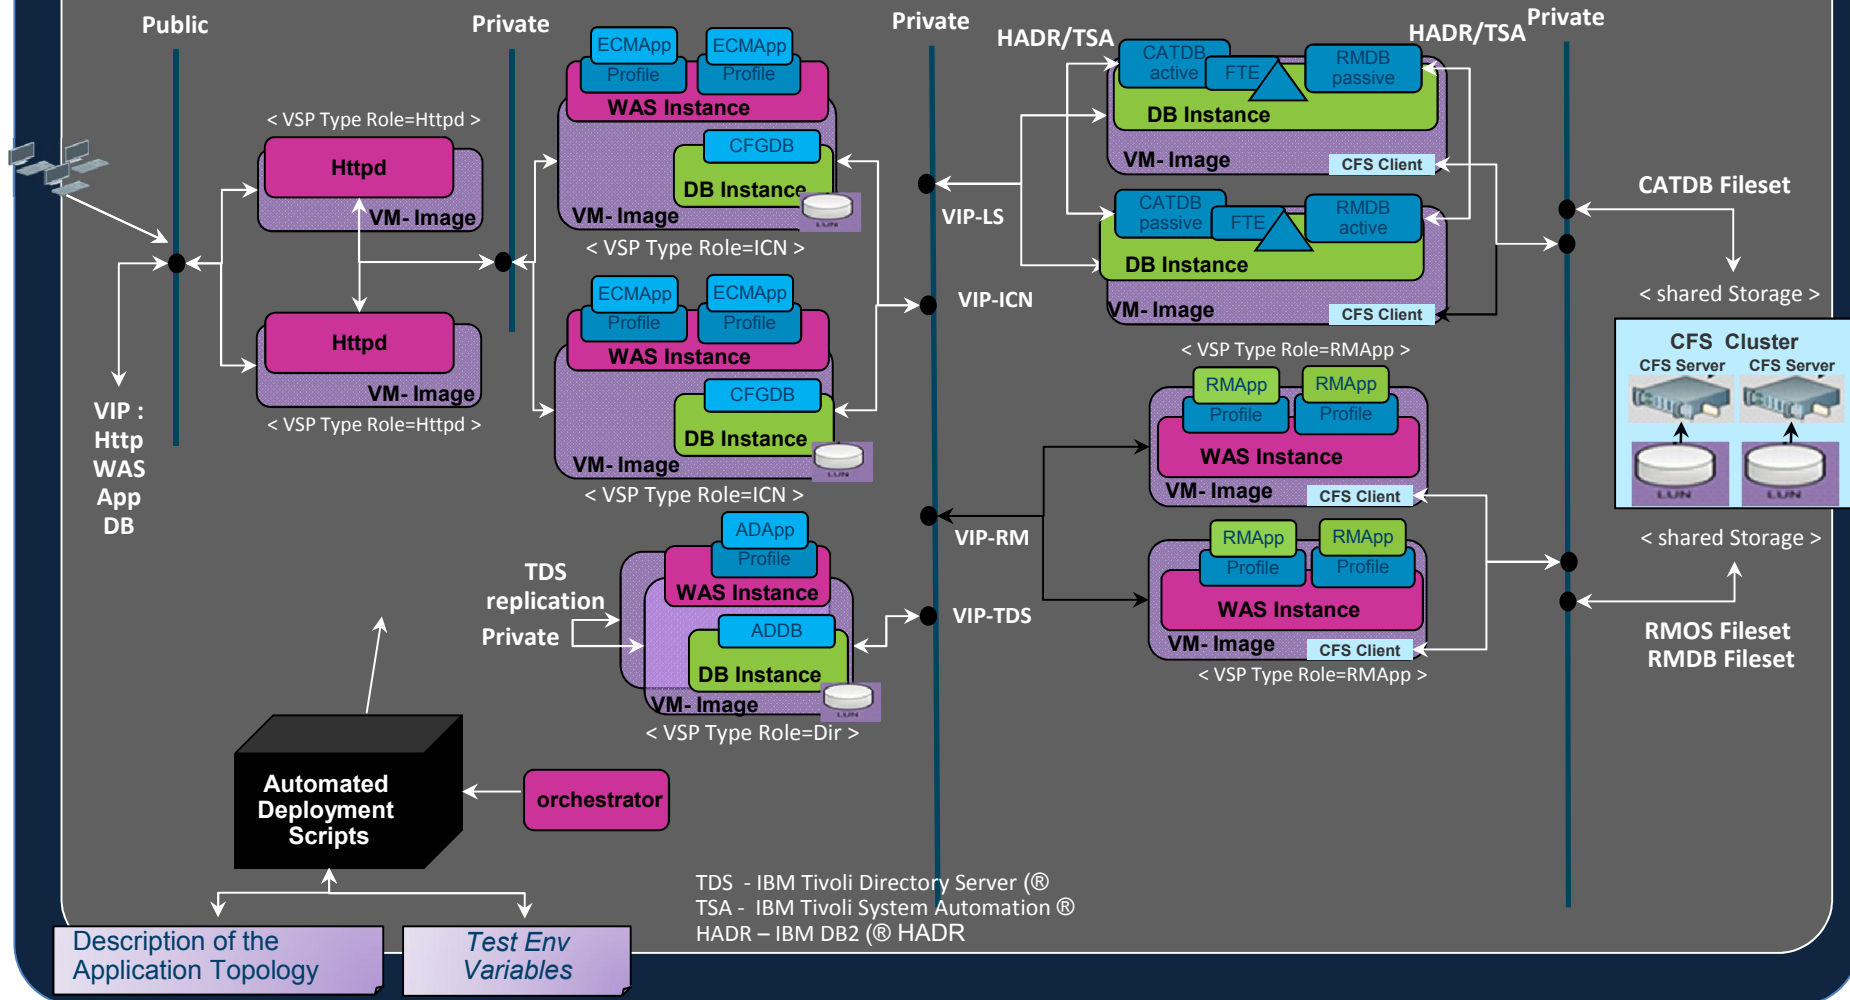
ECM Small PROD System Configuration w HA Example



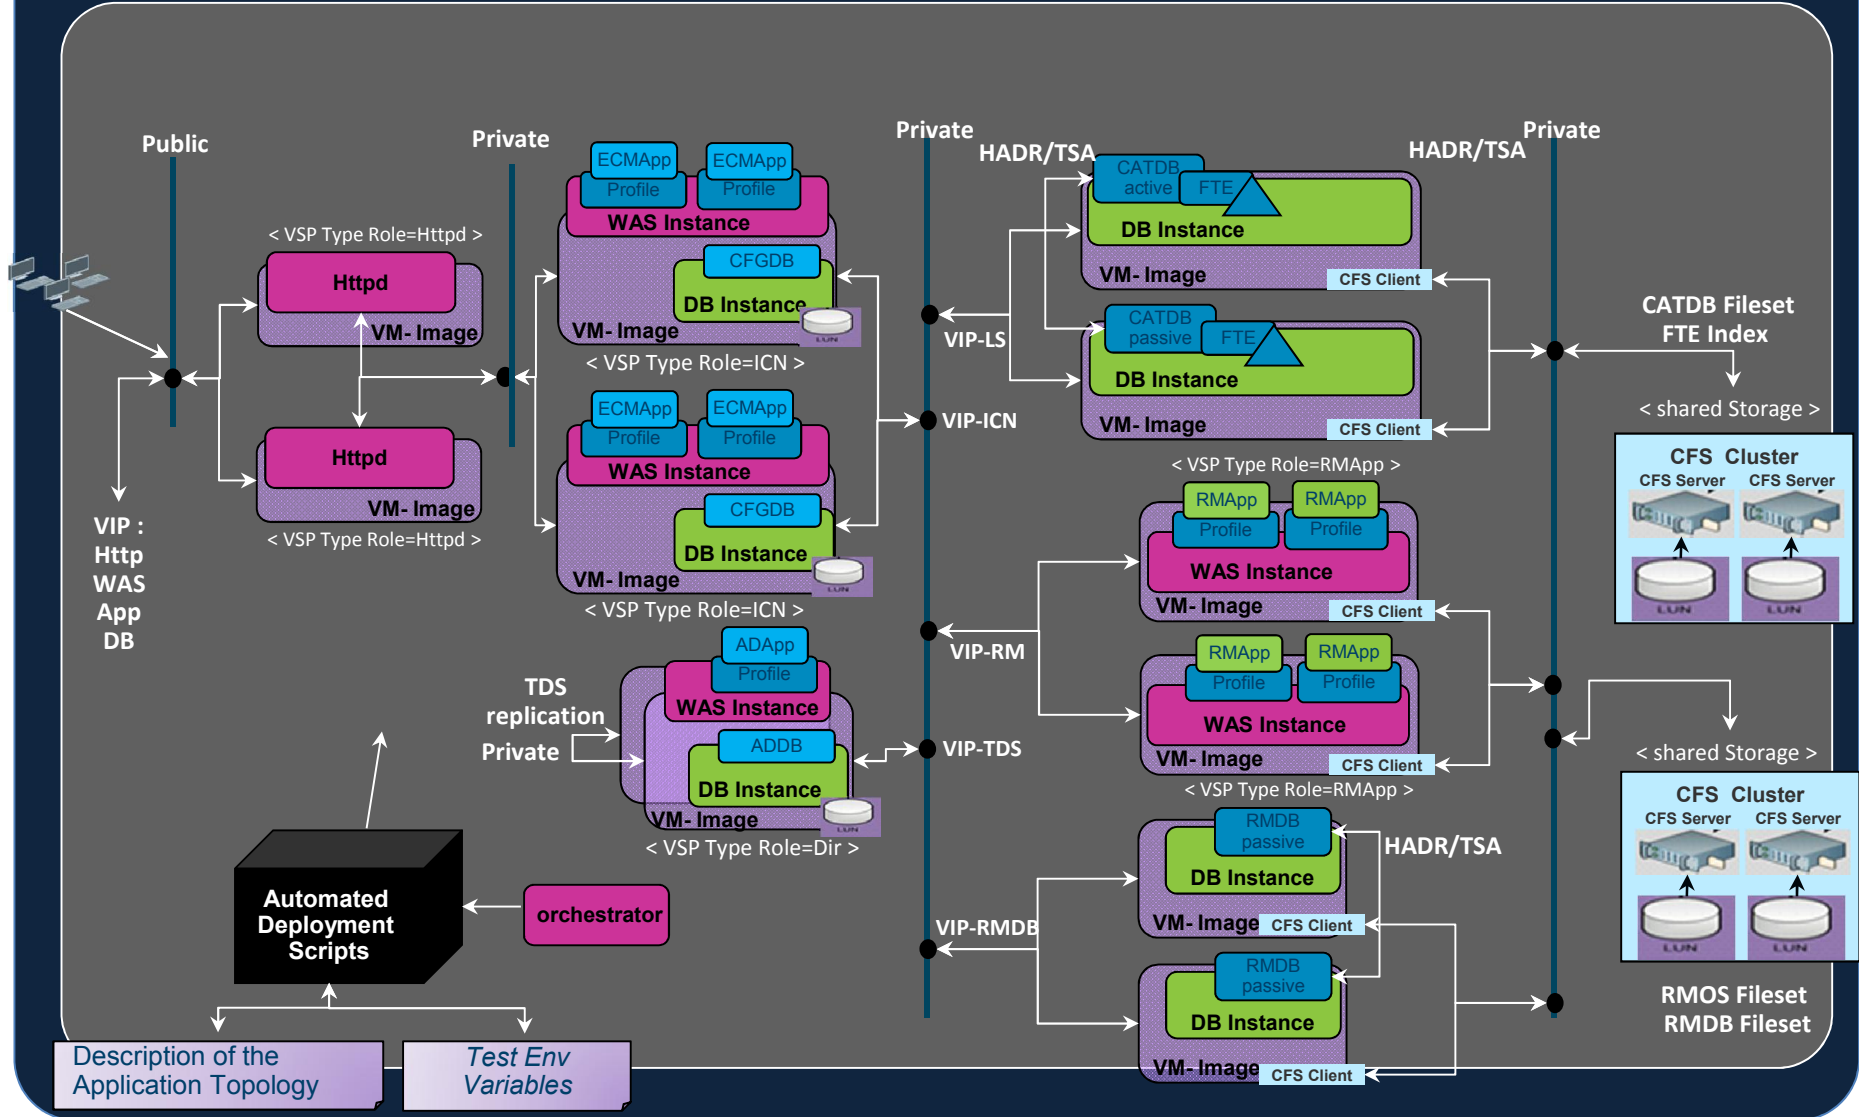
ECM Medium PROD System Configuration single profile server example



ECM Medium PROD System Configuration split profile server example



ECM Large Scale PROD System Configuration Example

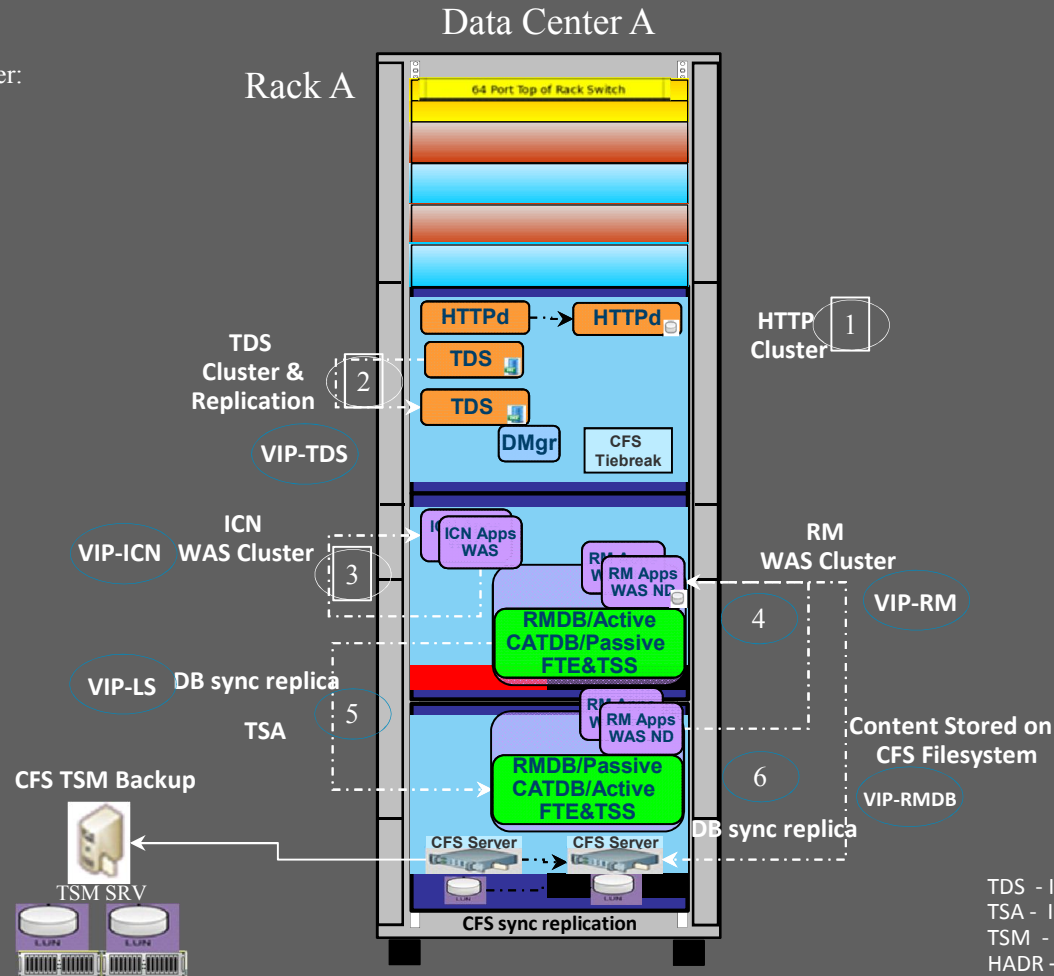


Single Rack / Single Medium Production System with HA Example



Legend HA Cluster:

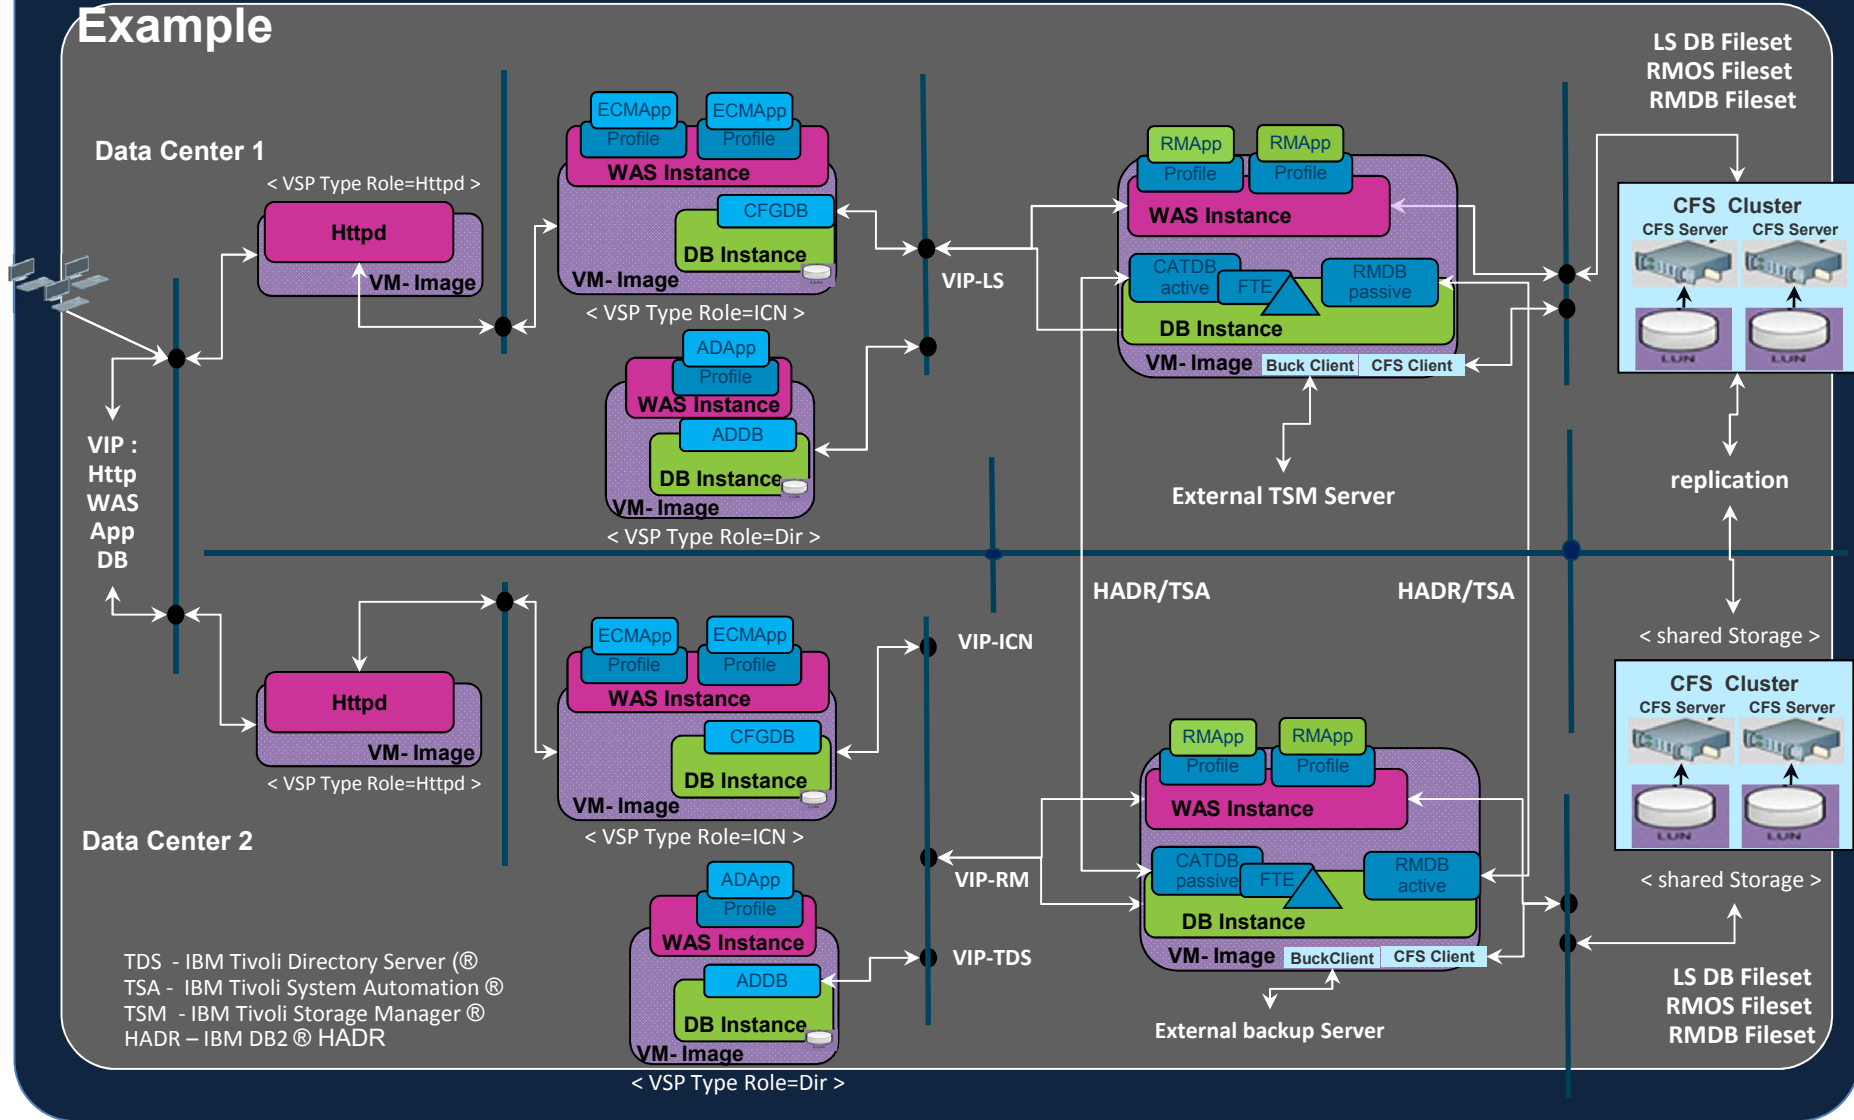
- 1: Http
- 2: TDS
- 3: ICN Apps
- 4: RM Apps
- 5: LS
- 6: RMDB



TDS - IBM Tivoli Directory Server (®)
 TSA - IBM Tivoli System Automation (®)
 TSM - IBM Tivoli Storage Manager (®)
 HADR - IBM DB2 (®) HADR

ECM PROD System Configuration combining HA&DR over 2x DC

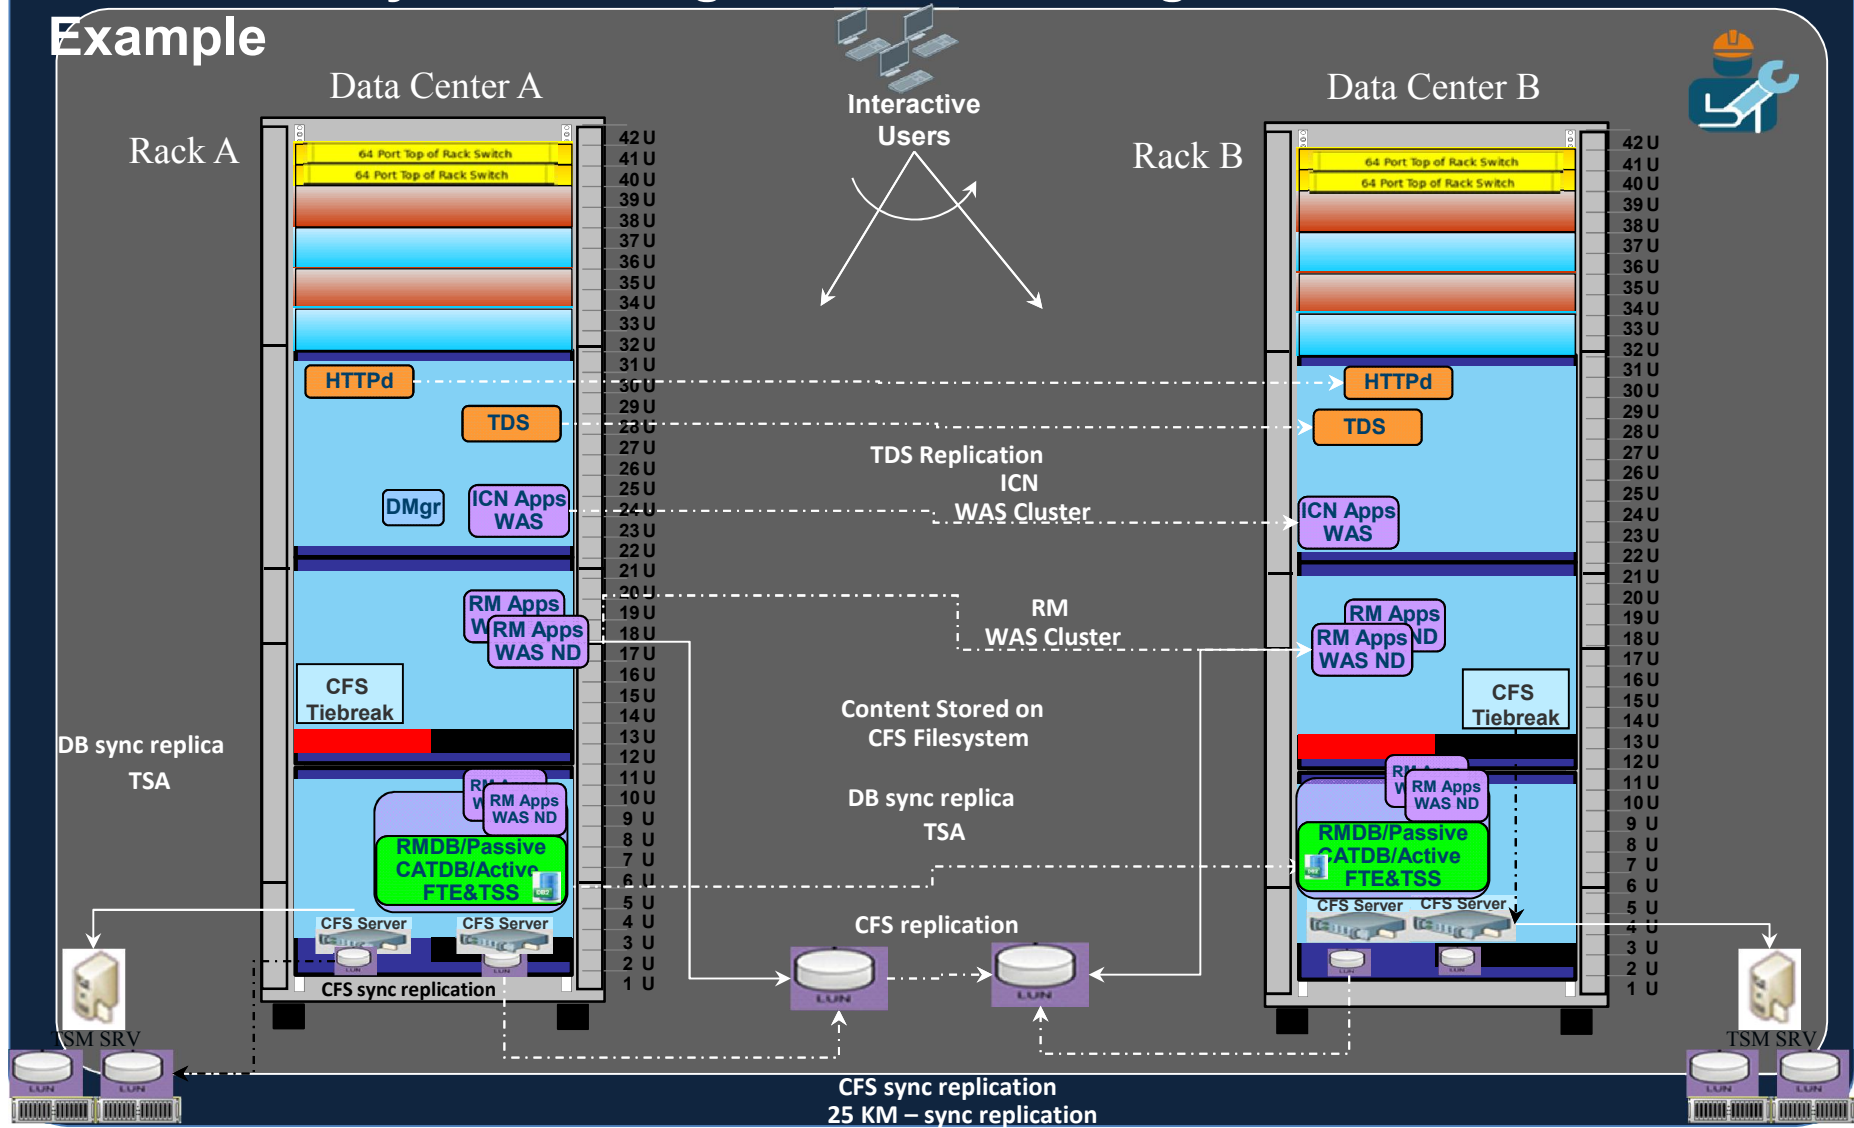
Example



TDS - IBM Tivoli Directory Server (®)
 TSA - IBM Tivoli System Automation (®)
 TSM - IBM Tivoli Storage Manager (®)
 HADR - IBM DB2 (®) HADR

ECM PROD System Configuration combining HA&DR over 2x DC

Example



Summary

Software Solutions must be enhanced before they can run as fully managed services in the Cloud. Key aspects to consider are:

Business model

- Fully managed service instances
- Self-service model
- Pay per use cost model

Service Delivery

- Fully managed service instances
- Integrated holistic solution approach
- Data availability
- Service continuity

Data residency

- Cross data center (cross – national borders) operations and federate access is among the first things to consider

Data security and privacy

- No data co-mingling across tenants boundaries
- Reliable and secure access
- Encryption of data in transit
- Encryption of data at rest
- Tenant specific key management

References

2014 Dynamic cloud service topology adaption for minimizing resources while meeting performance goals
Mega, C. ; Waizenegger, T. ; Lebutsch, D. ; Schleipen, S. ; Barney, J.M.

2012 Thomas Ritter, Bernhard Mitschang, Cataldo Mega: Dynamic Provisioning of System Topologies in the Cloud. I-ESA 2012: 391-401

2013 Tim Waizenegger, Oliver Schiller, Cataldo Mega: Datensicherheit in mandantenfähigen Cloud Umgebungen. BTW 2013: 477-489

2013 Kai Liu – MA Thesis - University of Stuttgart
Development of TOSCA Service Templates for provisioning portable IT Services

2013 - Ashraf, Umair, MA Thesis - University of Stuttgart - Securing Cloud Applications Using Two-Factor Authentication

2011 - Andreas Börner, MA Thesis - University of Stuttgart -
Orchestration and provisioning of dynamic system topologies

2011 - Frank Wagner, PhD Thesis - University of Stuttgart -
A Virtualization Approach to Scalable Enterprise Content Management

2011 Thomas Ritter, PhD Thesis - University of Stuttgart
Dynamic Provisioning of Multi-Tenant System Topologies and Security in the Cloud

2011 Tim Waizenegger, MA Thesis - University of Stuttgart
Data security in multi-tenant environments in the Cloud

References

2010 - Florian Fritz, MA Thesis - University of Stuttgart

Maximization of resource utilization through dynamic provisioning and de-provisioning in the cloud

2009 Mega, Cataldo; Krebs, Kathleen; Wagner, Frank; Ritter, Norbert; Mitschang, Bernhard: Content-Management-Systeme der nächsten Generation. In: Keuper, Frank (ed.); Neumann, Fritz (ed.): Wissens- und Informationsmanagement ; Strategien, Organisation und Prozesse, Wiesbaden: Gabler Verlag, 2009.

2008 Moosbrugger, Alexander: Evaluation of Scale-Out Characteristics of an ECM System using Overlay-Techniques to Support Dynamic Topologies, Diploma Thesis No. 2711,

2008. Wagner, Frank; Krebs, Kathleen; Mega, Cataldo; Mitschang, Bernhard; Ritter, Norbert: Towards the Design of a Scalable Email Archiving and Discovery Solution. In: Atzeni, Paolo (ed.); Caplinskas, Albertas (ed.); Jaakkola, Hannu (ed.): Proceedings of the 12th East-European Conference on Advances in Databases and Information Systems,

2008. Wagner, Frank; Krebs, Kathleen; Mega, Cataldo; Mitschang, Bernhard; Ritter, Norbert: Email Archiving and Discovery as a Service. In: Badica, Costin (ed.); Mangioni, Giuseppe (ed.); Carchiolo, Vincenza (ed.); Burdescu, Dumitru Dan (ed.): Intelligent Distributed Computing, Systems and Applications; Proceedings of the 2nd International Symposium on Intelligent Distributed Computing: IDC 2008; Catania, Italy.

2005 . Mega, Cataldo; Wagner, Frank; Mitschang, Bernhard: From Content Management to Enterprise Content Management. In: Gesellschaft für Informatik (ed.): Datenbanksysteme in Business, Technologie und Web, 2005.

Thank You!

Please note

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Acknowledgements and Disclaimers

Availability. References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates.

The workshops, sessions and materials have been prepared by IBM or the session speakers and reflect their own views. They are provided for informational purposes only, and are neither intended to, nor shall have the effect of being, legal or other guidance or advice to any participant. While efforts were made to verify the completeness and accuracy of the information contained in this presentation, it is provided AS-IS without warranty of any kind, express or implied. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this presentation or any other materials. Nothing contained in this presentation is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.

All customer examples described are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual environmental costs and performance characteristics may vary by customer. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

© *Copyright IBM Corporation 2014. All rights reserved.*

• ***U.S. Government Users Restricted Rights - Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.***

IBM, the IBM logo, ibm.com, and IBM SmartCloud Content Management are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or ™), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at www.ibm.com/legal/copytrade.shtml

Other company, product, or service names may be trademarks or service marks of others.