Publications Office
of the European Union

INFORMATION SYSTEMS

Technical Environment and Standard Operating Procedures of the Publications Office

Type	Annex	Status	
Version	1.0	Date	15-06-2020
Reference		Authors	
Revised by			

Table of Content

Table of Figures

Table of Tables

Purpose of the Document

This document provides an overview of the technical environment of the Publications Office as well as some general rules linked to the technical organisation of the Publications Office and applicable to all applications hosted at the Publications Office.

Disclaimer

The information contained in this document reflects the situation in force at the Publications Office at the time of writing and is subject to change.

The Publications Office cannot be held liable for the consequences of any reliance on the information provided or for any inaccuracies in such information and it does not commit the Publications Office regarding the future evolution of its data processing and network environment.

The content of this document may vary in relation to any particular project. In particular, the environment to take into consideration for a specific project/purchase order – especially the exact software versions – will be determined at the very beginning of the project. This also includes the rules to apply by both parties in order to modify this environment.

The Publications Office strongly advises the contractor to ask for clarification should there be any doubt about the contents of this document. If requested by the contractor, a meeting can take place at the very beginning of the project to answer questions and provide examples of the expected documents.

These specifications describe the environments, as it is at the launch of the procurement procedure. During execution of the contract, some technologies may be added or modified, whereas others may be phased out. This may happen for:

- Improving quality of services to Publications Office' Users
- Applying cost reduction
- Following technologies lifecycle
- Applying new technologies (e.g. increase PaaS approach, implementing auto-scaling, etc.)
- Improving compliance with standards and methodologies (i.e. improve ITIL approach, apply DevOps approach, increase in-house SW development)
- Work Optimization
- Etc.

In any case, this will not change the overall scope of the contract, but would require a capacity of adaptation to changes from the contractor.

The contractor is required to adapt to these changes for the entire duration of the contract.

1 Technical Environment of the Publications Office

1.1 Introduction

The Publications Office makes a distinction between systems used for office automation and administrative information systems on the one hand and systems used for production on the other hand. The quality of service and the constraints of availability are tighter for the production systems, since external partners with contractual agreements are already in place. Another important difference between these two types of information systems is linked to their architecture. The production information systems are usually spread over several servers and include complex production chains with processing on all nodes, whereas administrative and office automation systems are simpler and frequently use a one-to-many relationship between a server and its clients.

To run its applications, the Publications Office is using different architectures: Datacentre Housing and Hosting services offered by DIGIT (EC Directorate) and PaaS based on Cloud Technology. The following figure depicts sharing of responsibility between Publications Office and its providers.

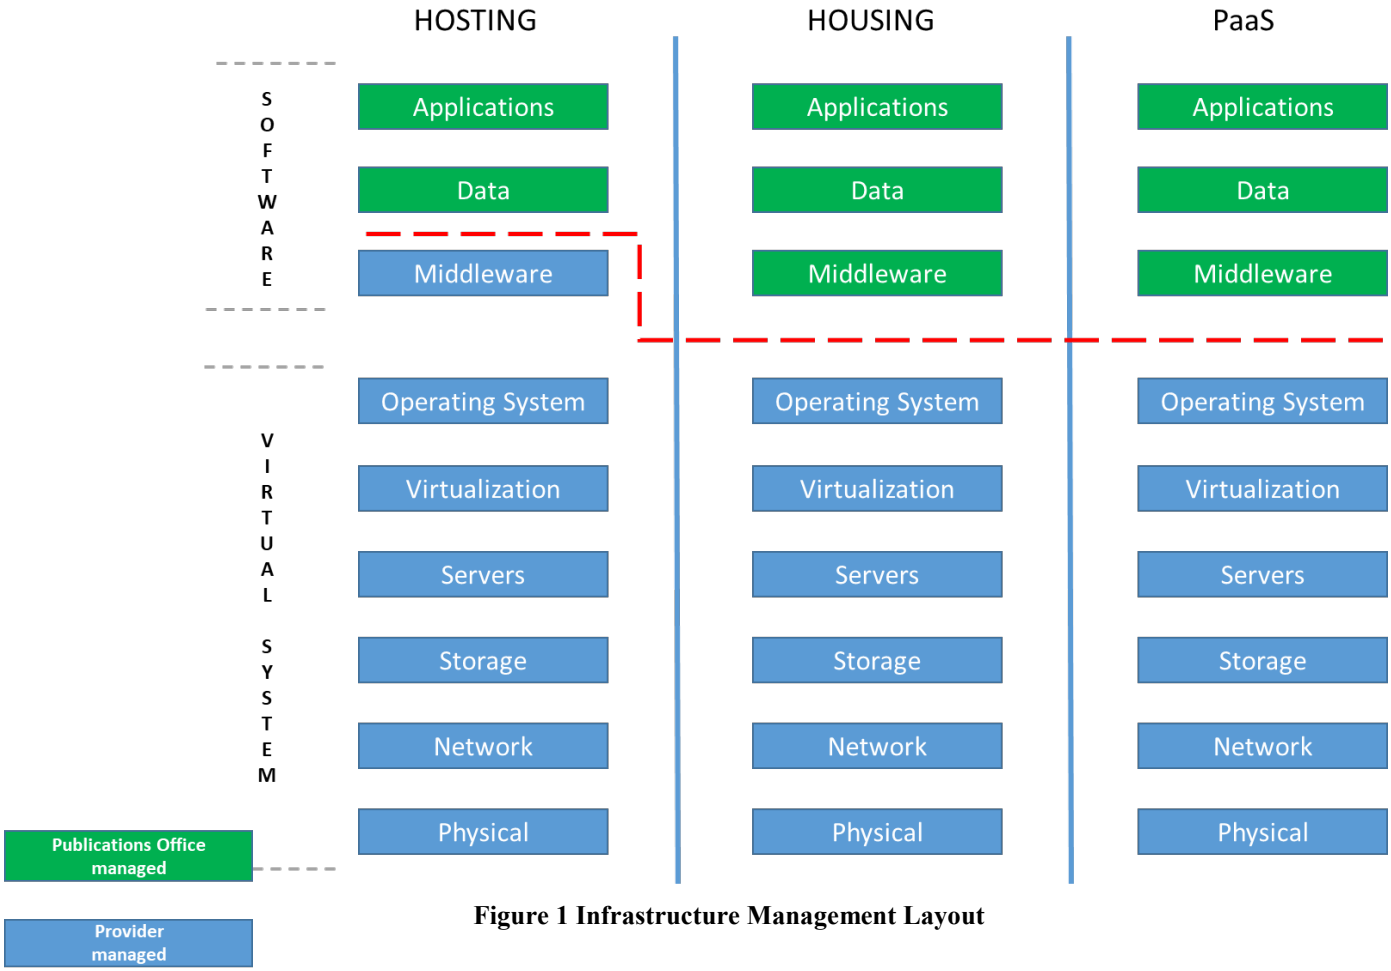

Figure 1 Infrastructure Management Layout

1.1.1 *DIGIT HOSTING*

Applications making use of this architecture use standard HOSTING services (e.g. Server, Storage, Application Server, Web Server, Database, etc.) on which the full control is managed by the provider. The Publications Office is installing, configuring and operating the Information Systems without direct access to the servers.

1.1.2 *DIGIT HOUSING*

This architecture provides the Publications Office access to the required Virtual Machine. This is provided with resources, storage and network connectivity as a specific service (e.g. Unix Container). It is then the responsibility of the Publications Office to install all packages (e.g. Application Server, Web Server, Database, Software, etc.) accessing the servers with provided account that is not granted with admin rights.

1.1.3 *AWS PaaS*

This Cloud configuration provides the Publications Office access to the required Virtual Machine. This is provided with resources, storage and network connectivity as a specific service (e.g. Unix Container). It is then the responsibility of the Publications Office to install all packages (e.g. Application Server, Web Server, Database, Software, etc.) accessing the servers with a provided account that is granted with admin rights. Deployment of entire Applications (including OS, Middleware, Software, etc.) is done using the Ansible playbook approach.

1.2 Application Architecture

The hardware and software architecture to use within a project is generally proposed by the contractor and should be compatible with the technical environment described in this document. The proposed architecture must take the total cost of ownership into account and ensure the sustainability of the application architecture. This architecture has to be validated by the Publications Office before implementation.

For the design of this architecture, the contractor has to take the following aspects into consideration:

- Hosting environment is the preferred solution (use of Housing environment has to be justified by the developer and approved by the Publications Office)
- Linux is the recommended OS.
- Recommended Database type is Oracle

For the information on supported platforms, databases, application servers and COTSs, please consult the Providers' Service Catalogue and the relevant documentation available:

- DIGIT Service Catalogue
- AWS Service Catalogue
- DIGIT Service Level Agreement

The Publications Office fosters professional methods of managing systems and therefore implements monitoring and measuring tools for systems and produces

statistics on the use of computing resources and on the quality of service provided. Developers should support this initiative with implementing an interface (as defined by OP) for additional monitoring of the application's status and health using the OP's monitoring tools.

The Publications Office promotes the implementation of a three-tier architecture, using thin clients, DIGIT's 'hosting' Linux application servers and mainly Oracle databases for reasons of performance, scalability and flexibility.

1.3 Data Exchange

Data exchange between applications is managed by a dedicated infrastructure supported by a set of tools called MFT (Managed File Transfer). MFT implements two different mechanisms:

- Internal data exchange between applications – This is managed via shared repositories. All relevant applications are configured to pick-up or drop files in specific folders.
- External data exchange with external entities – This is managed via SFTP interface via file push mechanism.

Managed File Transfer is the exchange of data (files) between applications – possibly running on distinct servers; the tools in use allow the triggering of processes based on the arrival of a file in a predefined directory (pre- and post-processing). Due to the asynchronous character of file exchanges, the order of files exchanged is not guaranteed; if sequencing is an issue, it must be managed at application level.

The Publications Office strongly advises the contractor to ask for practical implementation guidelines before starting any development that could require integration or interaction with the MFT tools.

Files that have to be transferred between disparate systems should follow the **POSIX.1-2017 portable filename character set**; incompatibilities between different platforms and operating systems are to be avoided, e.g. avoid using reserved keywords of operating systems, the space character and the ampersand in files transferred between systems.

1.4 Security Provisions

A separate document describing the OP Minimum Security Requirements is attached to each Call for Tender for information processing system and/or services issued by the Publications Office. Any application, system or service that is introduced at the Publications Office has to comply with these requirements. The contractor is requested also to comply with all EC security requirements established in the Security standards applying to all European Commission information systems that are in force in every moment (and which can be found at https://ec.europa.eu/info/publications/security-standards-applying-all-european-commission-information-systems_en).

1.5 Other Provisions

The Publications Office promotes the virtualisation of services and the use of abstraction layers in order to increase flexibility. This implies in particular that:

- web-based applications must allow the deployment and the correct operation behind any http reverse proxy chain;
- applications must allow virtual hosting i.e. the binding of the application to only some of the IP addresses/hostnames of a multi-homed server;
- applications must allow easy integration in the DRP of the Publications Office;
- applications must be compatible with an MS Windows 2003/2008 terminal server architecture;
- if the application depends on network services like DNS or LDAP, the server name or the IP address should not be hard-coded but should remain configurable.

Before being put into production, all of the Publications Office's core business applications, which are often interdependent, are tested for incompatibilities on dedicated machines.

Authentication mechanisms must use the available centralised directory server infrastructure (ECAS server of the Commission, Active Directory server of the Publications Office, …)

Data archiving and purging mechanisms have to be foreseen and implemented so that data volume growth does not degrade application performance nor backup/restore operations. Data management should in general be separated from modifications to be done of the application binary code or configuration files.

On the hardware side, the technical data processing infrastructure is currently made of several components that can be grouped into the following categories:

- Network
- Servers Platforms (Data Centre and Cloud based)
- Applications servers (Hosting & Housing)
- Databases servers (Hosting & Housing)
- Storage & backup (Data Centre and Cloud based)
- Workstations

Details on these categories can be found in the relevant sections later in this document.

2 Network

2.1 General network topology

The full network stack (infrastructure, security, subnets, DMZ and Proxy, interconnection and public access) for OP applications is provided as a service by the relevant infrastructure provider. The provider has thus the full control of all components and configurations and OP Applications must comply with security requirements and best practice.

Proxy Internet access is the default Internet access service for all applications. It uses proxy servers with authentication, which enable a higher degree of control over the communications originating from the Commission network and going towards the Internet.

All internal subnets (VLANs) assigned to OP are in general fully trusted each other and all systems and applications are reachable and available for direct connection.

There are specific cases where subnets cannot be directly connected (i.e. DMZ or specific subnet) to the internal OP subnets. Here, there is firewall service filtering the traffic and by default everything is closed. So any connection to be implemented in this case shall be requests as a change and is granted (if duly justified and approved) applying point-to-point concept.

Network connectivity is provided to all containers for the Applications using a fully redundant network infrastructure. In case of problems on one of them, all network traffic is transparently routed through another path.

OP network is also integrated within the network of the European Commission that is connected to the sTESTA network (Trans-European Services for Telematics between Administrations). The sTESTA network allows the Publications Office to establish private connections with most of the national Administrations of the EU Member States and most of the EU Institutions. Dedicated SFTP servers hosted at DIGIT/Commission are used for sending and receiving files over the public Internet or via the sTESTA network.

Application can be also implemented with the Load Balancing availability option.

2.2 File Transfer

Any file transfer with the Publications Office shall be managed only using SFTP. Proper access will be granted on the specific OP SFTP interface (Proxy SFTP server with proper username and password). File transfer towards the OP is permitted only with push option. Pull of files is not allowed from OP systems.

2.3 Mapping & Proxy

Web-based Information Systems (IS) can only be accessed via a reverse proxy. ISP will implement a reverse proxy mapping between the external URL (visible to the IS users) and the internal URL. The mapping must comply with ISP's mapping guidelines.

The reverse proxy service dispatches in-bound network traffic to a set of servers, presenting a single interface to the caller.

All accesses that make use of the network of the European Commission (e.g. Internet accesses) have to comply with the general security rules of the Commission. This also implies that all mentioned networks are interconnected through stateful inspection firewalls. In particular, outbound direct internet access is NOT allowed. Applications that require internet access have to be "proxy aware".

Please refer to the documentation of the Reverse proxy mapping service in the DIGIT Service Catalogue.

2.4 Remote Access

Remote Access is to grant OP service providers with access to OP infrastructure hosted in DIGIT Data Centre for remote maintenance on servers and applications. There is always a strong authentication from the company and each authenticated source has its own set of Firewall rules allowing access to a restricted list of destinations until a specific end date. The following picture depicts the high level architecture in place.

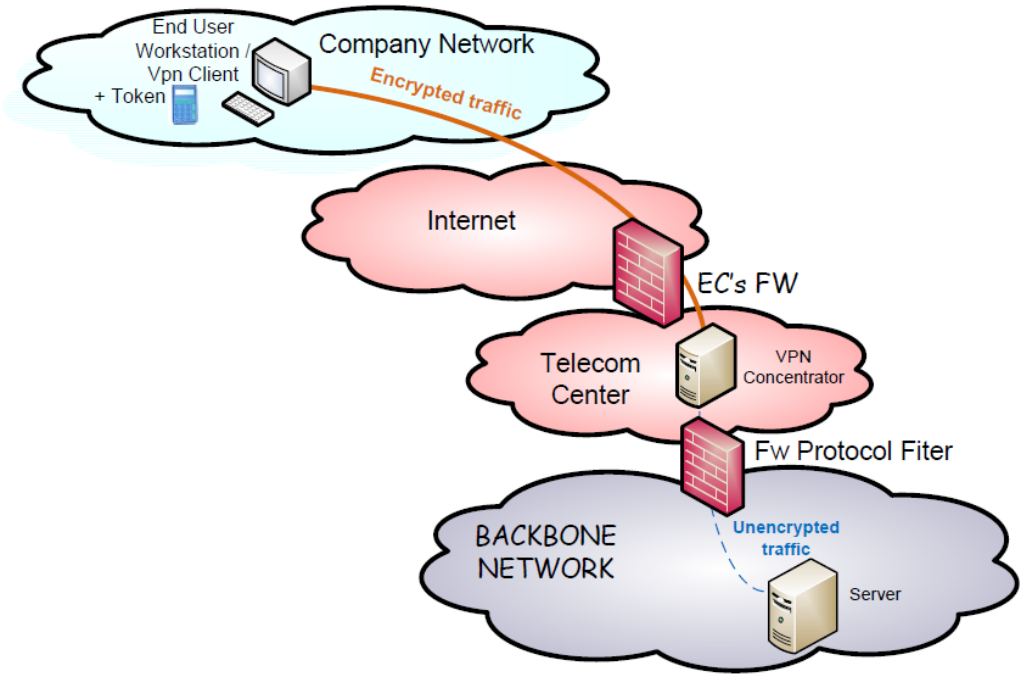

Figure 2 Remote Access Architecture

There is always a strong authentication from the company via a "token" device. The actual

authentication is carried out from inside the EC telecom centre. The network traffic between the remote-end workstation and the VPN concentrator is tunnelled and encrypted (AES-256) once the remote user is successfully authenticated4. Any TOKEN device assigned to a specific company (security convention) can be used for the remote-end authentication. All those Tokens are linked to the authorised list of accesses. The protocol filtering is performed by firewall equipments configured with a set of rules. Each rule has an expiry date that matches the expiry date of the security convention to which it refers to. In addition - each firewall rule has a set of parameters which are:

- Source IP address / Source port number
- Destination IP address / Destination port number
- IP protocol number
- End Date
- Maximum session duration = 10h
- Idle time out is 30 minutes

3 Servers Platforms

3.1 Application Hosting Servers

This section describes the environment of Applications Servers provided as Hosting Services. Hosting services are controlled and managed by service providers on their own platform. The Publications Office installs applications and applies relevant configuration following agreed processes and procedures and using standard and official interfaces (either automatic installation tool or ticket raised to provider). The Publications Office has access to the Applications Servers' log files , so they can be collected in case of investigation. The version of the Application Servers and the version of the platform running it (Operating System) are controlled by the providers and are following the standard Lifecycle process (patching, upgrade/update). The contractor shall take the Lifecycle into proper consideration and align its development process to it. Each applications server runs on its own virtual server and the contractor has to provide indication on resources to be allocated (CPU, RAM), specific performances requirements if any (i.e. access to storage systems), etc.

3.1.1 *Tomcat*

This hosting applications service is providing standard Tomcat instances. As an open-source implementation of the Java Servlet, JavaServer Pages, Java Expression Language and Java WebSocket technologies, it provides a "pure Java" HTTP web server environment for Java code to run in. Bundled with an optional Tomcat Native component, Tomcat can use Apache Portable Runtime (APR). Red Hat Enterprise Linux is used as the hosting platform for this service.

This service manages standard java applications deployed in war files and having configuration in Tomcat/Java properties files. War files and relevant properties files are uploaded on specific repository and following deployment (as well as stop/start operation if required) is managed via proper interface made available by provider. Applications (and relevant configuration). To be noted that as this is a hosting service, there are limited rights granted to Publications Office:

- Stop/start of Tomcat can be managed via specific interface
- No access to relevant servers
- Log files are visible via specific interface
- Application shall run without admin privileges (including start/stop)

The following picture provides an example of tasks that can be managed by the Publications Office (in some cases they are fully automatic).

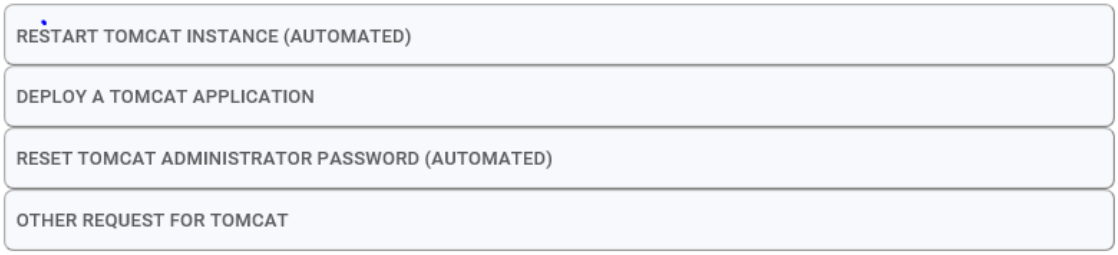

Figure 3 Tomcat Applications Server Service Requests

3.1.2 *LAMPT*

This service provides a dedicated LAMP/T stack (Linux Apache, MySQL, PHP, Perl, Python, Tomcat). This is particularly useful to have all components deployed in one single VM. On top of what included with a Tomcat service (see previous chapter) this services is also offering MySQL database, PHP, Perl and Python server-side scripting languages. The service is operated in the same fashion as Tomcat. Proper interfaces is provided to deploy war file to be run and relevant configuration/properties files. Limitations on admin rights apply also to this service, but here there is the possibility to have ssh access to the server (with user owning the application) to manage PHP, Python and Perl part of the service. Concerning Database part instead, MySQL can be accessed for limited DBA tasks vie phpMyAdmin interface. Again admin rights are not granted as MySQL managed as hosted DB.

The following picture provides an example of tasks that can be managed by the Publications Office (in some cases they are fully automatic).

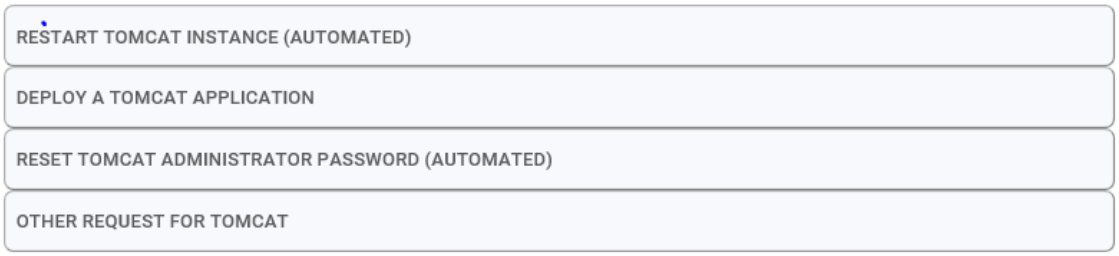

Figure 4 LAMPT Application Server Service Requests

3.1.3 *WebLogic*

This hosting service provides standard Weblogic Instance. It is including the following components:

- WLS: (WebLogic Server)
- OSB: (Oracle Service Bus)
- SOA / BPM: (Oracle SOA Suite)
- Customer portal and an API

The service is managed via provider interface. The following picture provides an example of tasks that can be managed by the Publications Office (in some cases they are fully automatic).

Figure 5 Documentum Application Server Service Requests

No Access is granted to the server.

3.1.4 *Documentum*

This service is just implemented for the Publications Office, it is still under definition and is still not formally integrated in the DIGIT service catalogue. The service is fully managed by DIGIT; however it is not integrated in the system used to manage services and any task is managed via ticket. This service provides the following component:

- Documentum Content Server
- Documentum DocBroker

No access to the server is granted.

3.1.5 *Coldfusion*

This Application server is delivering a full and standard Coldfusion instance that is running on top of JEE application server (Tomcat) installed on a RedHat Enterprise system. Lifecycle process is fully controlled by provider and is also dependant on relevant lifecycle process applied for underlying component (JDK/Tomcat). No rights are granted to administer Coldfusion Instance. Any change or task related to administration has to be requested via standard interface. The service comprises the following product ad options:

- Adobe ColdFusion Enterprise Edition
- ColdFusion J2EE cluster
- Tomcat application server
- Java JDK
- DataDirect JDBC driver
- Apache Solr full text search engine
- Web services
- Flex integration
- Event Gateways
- Sandbox security

Each application is installed on a dedicated ColdFusion instance. Each instance has separate settings and runs in its own Java Virtual Machine. This way, performance problems on one application are less likely to impact others. Furthermore, for security reasons, a ColdFusion instance should only provide one type of content.

The following picture provides an example of tasks that can be managed by the Publications Office (in some cases they are fully automatic)

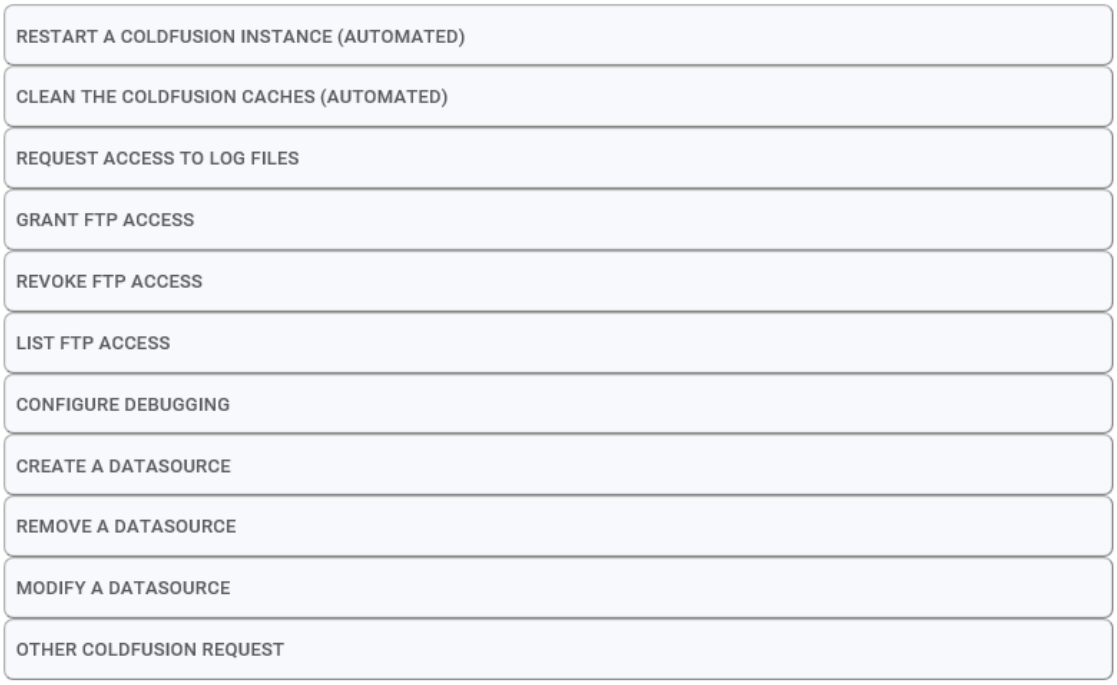

Figure 6 ColdFusion Application Server Service Requests

3.2 Application Housing Servers

The preferred model of infrastructure is for the Publications Office is Hosting; thus, the adoption of Housing infrastructure shall be justified (not manageable constraints) and approved by the Publications Office.

3.2.1 *LINUX Servers*

This chapter describes constrains that apply to the OP Linux systems installed in Datacentre Housing mode or in PaaS mode.

3.2.1.1 Operating System part

3.2.1.1.1 Storage

All necessary local and shared storage is provisioned by the provider following the specifications requested (either NAS or SAN). Access to storage is always done via standard filesystem (e.g. / /etc /var etc.).

3.2.1.1.2 Network

By default, each VM is supplied with one virtual network interface. Additional interface can be asked if connections to second subnet is necessary.

The main IP address of the server is associated in the DNS with its name; one or multiple CNAME records can be defined on the same IPs, to be used by the application to bind on it when exposing some services.

3.2.1.1.3 Operating System

The operating system is installed and updated regularly by RH Satellite. By default the last revision of RHEL present on RH Satellite is deployed (refer to Service Catalogue). The provider might be forced to immediately apply Security patches depending on the severity of the vulnerability.

3.2.1.2 Application part

3.2.1.2.1 Folders

The folder /ec/<environment>/app/<application-name>/ is used to host applications files.

access rights	owner	group owner	folder
drwxr-xr-x	i_<application-name>	<application-name>	/ec/<environment>/app/<application-name>/data
drwxrwxr-x	i_<application-name>	<application-name>	/ec/<environment>/app/<application-name>/data/<component>
drwxr-xr-x	i_<application-name>	<application-name>	/ec/<environment>/app/<application-name>/logs
drwxrwxr-x	i_<application-name>	<application-name>	/ec/<environment>/app/<application-name>/logs/<component>
drwxr-xr-x	i_<application-name>	i_<application-name>	/ec/<environment>/app/<application-name>/users
drwxrwxr-t	i_<application-name>	<application-name>	/ec/<environment>/app/<application-name>/users/<application-name>
drwxr-xr-x	i_<application-name>	i_<application-name>	/ec/<environment>/app/<application-name>/users/i_<application-name>

access rights	owner	group owner	folder
drwxr-xr-x	i_<application-name>	i_<application-name>	/ec/<environment>/app/<application-name>/users/system

Table 1 Linux Server' Filesystems layout

If an application is exchanging data with other applications or parts of the Datacentre infrastructure, a proper folder structure must be supported by the application:

owner	group owner	folder
i_<application-name>	woodsys	/ec/<environment>/app /<application-name>/xchange

Table 2 Linux Server Data Exchange Repository

3.2.1.2.2 Users and groups

As not root privileges is granted and as **SUDO**, privileges are rarely permitted, system users shall be able to:

- Modify configuration files of all the application components
- Read the log files, both system's and application's
- Locate scripts/whatever in predefined locations under the /applications/<application-name>/ directory hierarchy
- Stop and start application components

At least two users must be created: the <application-name> and the i_<application-name> users.

Note that these functional accounts can only be 5, 6 or 8 characters long (7 characters are reserved for individual accounts).

- <application-name> user

This user is used only for running the application or its manually-installed components; it needs r/w access to all the data directories but won't be able to modify other files.

- i_<application-name> user

The i_<application-name> user is introduced to implement the above requirements and enhance security. This user will:

- own all the binary (static) files of components (not delivered by means of an RPM package)
- have read/write access to the configuration files and read access to the log files
- have the rights to start/stop all the applications components
- have the rights to switch and <application-name> users and group (su and sg commands) with no password.

Concerning groups, two groups shall be created

- <**application-name**>
- i_<**application-name**>

and the following user configuration shall be applied:

- i_<application-name> user is member of i_<application-name> and <application-name> groups.
- <application-name> user is member of <application-name> and i_<application-name> groups.

3.2.1.2.3 Application Start/Stop

Because applications must start at server boot time, application start procedures are integrated to Systemd (preferred on RHEL7).
OP uses these mechanisms to manage stop and start the applications as we need.

Startup and Stop scripts will be launched by i_<applications-name> user, with the `systemctl` command.

For systems installed in PaaS infrastructure, the above is managed via Ansible playbook.

3.2.2 *Windows Servers*

This chapter describes constrains that apply to OP Windows systems installed in Datacentre Housing mode. For Windows servers, OP is granted with local admin rights to named AD users.

Type	Product/Version recommended for all new developments
Operating system	MS Windows 2016 Server standard edition/enterprise edition
Storage	C: D: (on request)
User and Group	DIGIT Microsoft AD is used to provide Users and Groups.

Table 3 Windows Servers Configuration

3.3 Database

This chapter provides a description of the DB model currently in place at Publications Office.

3.3.1 *Database Hosting Datacentre*

This Type of Database is offered as a service by DIGIT. The overall Database infrastructure (called Container Database) is fully managed by provider and this include Lifecycle process (patching, upgrade/update). The contractor shall take the Lifecycle into proper consideration and align its development process to it. Provider deploys virtual Database instances (so called Pluggable Database) as per Publications Office requirements and implement the required settings (including data import/export). No admin rights are granted on the Database. Publications office is granted with users and rights to manage data and to manage reduced set of Database structure. Database is then made accessible to users and applications via proper ODBC link. Supported databases are:

- ORACLE
- MySQL
- MSSQL

All DBs are managed via standard Interface. The following picture provide an example of tasks that can be managed by the Publications Office (in some cases they are fully automatic)

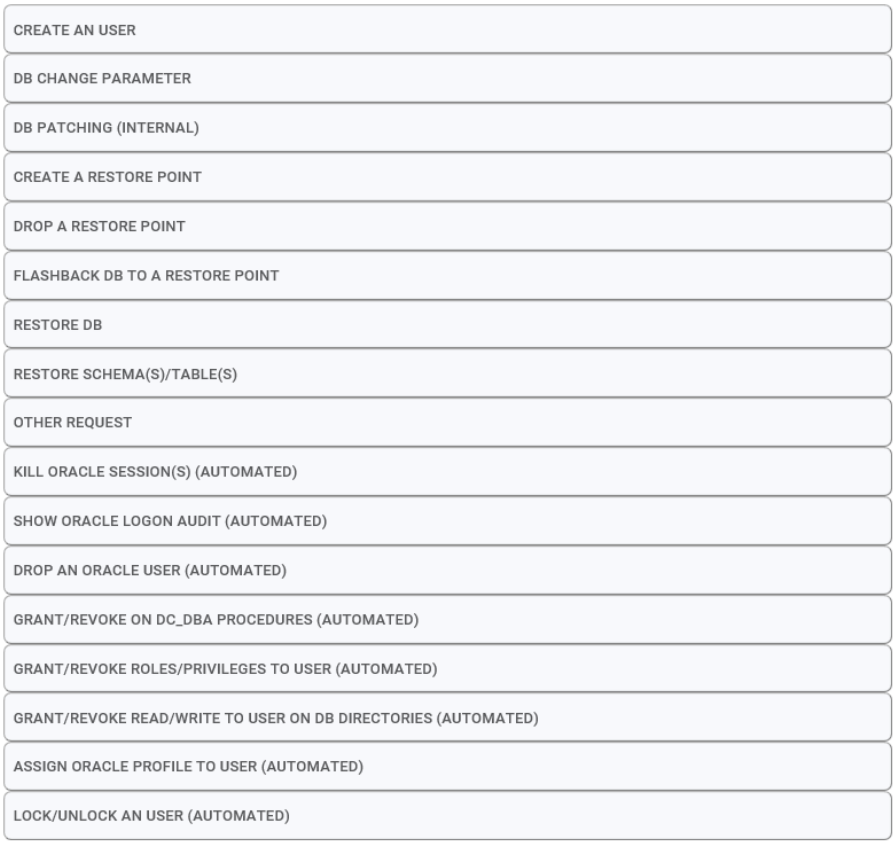

Figure 7 Database Service' Requests

3.3.2 *Database Hosting Cloud*

This Type of Database is offered as a service by Cloud provider (currently RDS by Amazon). The basic building block of Amazon RDS is the Database instance. A Database instance is an isolated database environment in the AWS Cloud, can contain multiple user-created databases and can be accessed by using the same tools and applications used with a standalone database instance. Database instances are created and modified by using the AWS Command Line Interface, the Amazon RDS API, or the AWS Management Console.

Each Database instance runs a Database engine (MySQL, MariaDB, PostgreSQL, Oracle, and Microsoft SQL). Each Database engine has its own supported features, and each version of a Database engine may include specific features. Additionally, each Database engine has a set of parameters in a Database parameter group that control the behaviour of the databases that it manages. As fully managed service Amazon control Lifecycle process (patching, upgrade/update). The contractor shall take the Lifecycle into proper consideration and align its development process to it. Amazon RDS does not provide shell access to Database instances. It also restricts access to certain system procedures and tables that require advanced privileges.

3.3.3 *Database Housing*

As all other Housing solution, adoption of Housing DB shall be duly justified and finally approved by the Publications Office. This kind of DB is fully managed by the Publications Office. It is installed on Housing Virtual Machine by the Publications Office who is, thus, granted with full Admin rights on these instances. Lifecycle process in this case if fully controlled by the Publications Office; however, this will be following the process designed by DIGIT in order to have an homogenous DB environment.

3.4 Storage & Backup

3.4.1 *Storage*

There mainly two types of Storage used in the infrastructure:

1. SAN

The Storage Area Network Service is a service, which consists of delivering and managing high performance and high capacity data disk storage across a storage network.

This service covers SAN storage for Application Housing Storage. Covered technology: SAN shared storage for datacentre. This Storage is available in two Service classes:

   a. Replicated (for Disaster Recovery purposes), backup retention 35 days
   b. Non-replicated (no Disaster Recovery), backup retention 35 days

2. NAS

The service consists in provisioning, operating and maintaining the underlying file sharing infrastructure, providing data disk storage access to computer systems, and ensuring that appropriate security controls are in place in regards of availability, integrity and confidentiality.

This service covers Data Centre File Service. The technology used to provide this service is NAS shares for datacentre. This Storage is provided in 3 service classes:

   a. Service level GOLD: the storage is mirrored with site redundancy, backup retention 35 days with site redundancy.
   b. Service level SILVER: the storage is mirrored with site redundancy, backup retention 35 days without site redundancy.
   c. Service level BRONZE: the storage is not mirrored, backup retention 35 days without site redundancy.

3.4.2 *Backup*

Backup is part of the provided services. This is provided at systems level and at storage level via proper tools and methodology (i.e. file based backup, storage snapshot, etc.). Standard retention time is 35 days.

Backup is in place also for Databases (making use or Oracle Recovery Manager for example) and is executed as per following description:

a. Full backups are performed once per week.

b. Incremental backups are performed once per day.

c. Archives are backed-up every 4 hours outside the backup window.

The contractor has to specify data (and specific process if any) that have to be backed up to recover the Application to a healthy status.

4 Workstations

In terms of software, the standard configuration for the workstations is the following:

Type	Product/Version
Operating system	MS Windows 10 64b
Office automation suite	MS Office 2016
Web browser	Microsoft Edge 41(+) Firefox 60.9(+)
Mail client	MS Outlook 2016
Connection middleware	Oracle NET 10G
Anti-virus	McAfee Endpoint Security
Application Locker	Applocker: Software Restriction & Policies feature; Soft should be installed and operate from C:\Program Files or if not possible it should contain a digital signature.
BitLocker	Hard Drive encryption on Laptop.
Application Virtualization	Microsoft Application Virtualization (App-V) Client 5.0 SP1 + ThinApp
Miscellaneous	Adobe Flash player Shockwave 12 Java V 1.6 & 1.7 MS Dot NET 4 PowerBuilder client 2.0

Table 4 Windows Workstation Configuration

Some other local production tools (Visio, Microsoft Project...) could also be found on the workstations as well as more specialised tools which are used for publishing (Adobe Photoshop, Adobe Acrobat Pro and Reader, Adobe CS, QuarkXPress ...).

For security reasons, the Publications Office must sometimes apply hotfixes to ALL workstations.

In terms of hardware, the workstations are ranging from Pentium E6300 Dual Core 2.8 GHz with 8 to 32 GB of memory

5 Standard Operating Procedures

This section provides information on the operational procedures and standards of applicable for software maintenance. Management of Bug Reports and Change Requests

Contractors and suppliers must propose procedures that specify how to manage bug reports and change requests. These procedures must allow the unambiguous

identification of every bug report and change request. The Publications Office hosts an instance of the Jira issue tracking product that is accessible from the internet and may be used for this purpose.

5.1 Software Deliveries

The Publications Office reserves the right to measure the quality of software deliveries it receives from contractors and suppliers, based on agreed-upon criteria, and reject any software delivery that does not fulfil these quality criteria. These criteria should be defined and validated in the project kick-off meeting by the different parties.

The Publications Office monitors the compliance of software deliveries with contractual obligations that arise as part of this contract; non-compliance with contractual obligations concerning the delivery of software may lead to rejection of the delivery. The compliance with standards and procedures will be measured and recorded and feedback will be given to the contractor, who is obliged to take these remarks into account for subsequent deliveries.

The Publications Office reserves the right to request information from contractors and suppliers concerning the development and delivery process. To this end, the contractor may be required to fill in a set of templates that will enable the Publications Office to gain insight into how the contractor complies with industry standard best practices concerning software development processes.

The Publications Office will measure metrics of the source code delivered in a process based on automated tools and on request deviating results must be justified or corrected by the contractor.

A software delivery can concern either a full installation or a partial installation that needs to be installed over an existing installation, for example as part of a patch or a hotfix. It should be clearly indicated whether a delivery concerns a full or a partial installation to be able to estimate the effort required. The installation instructions must be customised and targeted to the particular installation.

Application software delivered to the Publications Office by the contractor shall comply with the following rules:

- The contractor must define a detailed and **unambiguous numbering scheme** and use it for each software delivery.
- The delivery should contain **the source code of the application and the executable binary code** that can be deployed and installed by following the installation instructions.
  - Please note that if the contractor makes the source code available in a way that the Publications Office is able to compile the application binary files from it, the resulting binaries will be used for the installation instead and the contractor does not need to deliver the executable binary code.
- Each delivery that contains the source code of the application shall include a **build procedure** detailing how to build the executable code from the source code. The build procedure should contain a default target so that the build process can be automated.

- The first delivery of an application should be as complete as possible and include all required components; subsequent deliveries should only contain the updated components required for the current installation request.
- The delivery has to be uploaded to the **revision control system** of the Publications Office, which is currently based on Subversion.
- Each delivery shall include **installation instructions** (in electronic, editable format); the minimum contents of the installation instructions are further detailed in section *9.5 Installation Instructions*.
- Each delivery shall include a **release note** (in electronic format) containing the following information:
  - project/application name
  - unambiguous identification[1] of the version of the delivered software
  - for partial deliveries, the version of the application that is a prerequisite for the installation of the delivered package
  - for full deliveries, whether or not the installation has to take place on a clean environment
  - details on the changes or enhancements that are implemented with the delivered release
  - approximate uncompressed size of the delivery
  - a reference to the installation instructions (as defined in 9.5 Installation Instructions)
  - information about test results and the test procedure
- Each delivery shall include a **TEST folder** containing all the necessary information to be able to verify that acceptance testing has been performed by the contractor before the software was delivered. This folder will contain:
  - test procedures and test cases executed
  - test data used
  - test results and/or execution report

5.2 Technical Tests

Before an application is put into production, the Publications Office will conduct specific technical tests to assess whether the application adheres to the operational requirements of being run and operated in its data centre.

Depending on the results of these technical tests, the Publications Office reserves the right to reject a software delivery.

The contractor must foresee in the planning of the project a dedicated period of time for the execution of technical tests.

The technical tests will be performed by the Publications Office's staff in close collaboration with the contractor and based on the test procedures and test cases prepared by the contractor.

The contractor will prepare a report of the execution of the technical tests. The contractor will deliver this report to the Publications Office, together with the test data used. The test procedures and test cases will first be validated by the Publications Office.

[1] In accordance with the procedure defined in the approved quality plan

The test procedures/test cases must allow validating at least the following elements:

- application start and stop procedure
- application backup and restore procedure (including consistency checks after restore)
- disk space usage
- location of the log files
- periodic operational tasks (data reorganisation, purging, archiving, indexing...)
- correct working of the interfaces
- virtualisation capability (suitability of the application to be moved from one server to another) in the context of disaster recovery (DRP) and high availability (HA)

The effectiveness of the procedures but also their efficiency will be tested. Their impact on the overall performance of the system will be evaluated too.

The test procedures/test cases must refer to procedures described in the System Operation Manual in order to validate this manual. The minimal contents of the **System Operation Manual** are described in section 9.6 System Operation Manual, the contractor is free to add any other information deemed useful.

For the execution of the technical tests the Publications Office uses monitoring and measurement tools listed in section 7 UNIX/LINUX Servers and section 8 Windows Servers.

The technical tests will be conducted with a significant amount of data in order to evaluate the impact of data volume on effectiveness, performance and efficiency of the application.

Special attention will be placed on the CPU, I/O and memory intensive and time consuming tasks like:

- data archiving
- data purging
- data reorganisation
- data consistency check
- data synchronisation
- data indexing
- statistics

5.3 Installations

For each application, two different environments are set up at the Publications Office, a test environment and a production environment. No software installation in the production environment will be allowed without prior validation in the test environment.

The Publications Office strongly advises the contractor to set up at their premises a development environment similar (e.g. same OS version, same RDBMS version…) to the target production environment. It remains the responsibility of the contractor to make sure that the delivered software will run correctly in the technical environment of the Publications Office.

The installation of hardware, the operating system and other low-level software remains the responsibility of the Publications Office. The installation of application software is done by a dedicated team of integrators. If the complexity of an installation requires special expertise, the support or assistance of a technical expert of the contractor may be formally requested by the Publications Office; this support may be provided off-site (remotely by email or telephone) or on-site, depending on the specific situation.

Apart from the test and the production environments, the Publications Office may decide to set up an additional environment in order for contractors/suppliers to demonstrate that the software delivered can be installed and run correctly within the technical environment of the Publications Office while conforming to the requirements as described in the technical and functional specifications. In this case, the contractor will perform all installation and configuration tasks, with the assistance of technical staff of the Publications Office. The contractor will be granted the necessary access rights so that all necessary installation and configuration tasks can be performed. The granted access rights will be limited to those allowed by the security standards of the Publications Office. If the Publications Office requests the contractor to perform this demonstration, no associated costs will be reimbursed.

In order to ensure smooth installation and to evaluate the need for assistance or support by the contractor, the installation instructions should be delivered to the Publications Office 5 days before the official software delivery date, i.e. 5 days before the start of the official installation period.

The Publications Office uses a tool (Atlassian JIRA) to manage installation requests. Atlassian JIRA is a web-based application that can be accessed from the internet, which makes it possible for the contractor to monitor the progress of an installation request.

5.4 Application Documentation

5.4.1 *Installation Instructions*

This section details the minimum contents of the installation instructions. The Publications Office will provide a template to help with creating the installation instructions on request.

5.4.1.1 General Conditions and Limitations

The installation document is a compulsory requirement for each delivery and has to provide step-by-step instructions to be executed, organised in a sequential and logical manner.

The installation instructions should contain all information necessary to perform the implementation tasks of the installation by an experienced system integrator without knowledge of the application.

Variables and parameters of the application that have dependencies with other software applications or the operating system should be made configurable. Variables and parameters that need to be changed or adapted for a specific environment have to be clearly marked and if their values are unknown, sufficient information has to be provided that makes it possible to identify the required value.

The installation instructions have to clearly indicate who will be responsible for the execution of a particular task; complex actions should be broken down into smaller tasks; as far as possible, where different actors are involved in the execution of tasks, these tasks should be grouped together by actor to minimise a back-and-forth between different teams and to make the installation process more seamless.

5.4.1.2 Specific Conditions and Limitations

- the internal standards and procedures of the Publications Office pertaining to the installation and operation of software applications must be respected
- use of root user is prohibited; installations are to be done using a user identification that is specifically created to install and run the application; where root access is required during installation, this must be explicitly stated
- use of sysdba is prohibited where database access is required
- kernel and system parameters should not be changed
- hard-coded IP addresses have to be avoided

5.4.1.3 Pre-requisites

The installation instructions describe in detail the hardware configuration (agreed at the beginning of the project) and the software configuration that has to be in place before the installation is started.

- Minimum hardware requirements (CPU, memory, disk space, network throughput, etc.) in compliance with the standards of the Publications Office.
- Software requirements (version, patches, specific configuration parameters, required modules, etc.) for software required to install the software delivery, for example the operating system, tools, and other pre-requisite software. For some software, for example Oracle, some supplementary modules might be required. In case of specific patches or service packs required, the exact version of the concerned application or component that has to be in place before the application can be installed should be specified.

5.4.1.4 Application Interface/Data Flow

The installation instructions should give an end-to-end description of each data flow.

For each flow, the following elements should be described:

- origin (source server name and directory)
- destination (target server name and directory)
- protocol to be used (including user identification and password required if applicable)
- estimate of volume of data exchanged
- how data transfers are scheduled or triggered
- description of pre- or post-processing commands to be executed
- error handling (including users and email distribution lists to inform)

5.4.1.5 Application Installation

This section concerns the information that needs to be provided for the installation of the application software. In case the application architecture follows the client/server model, the installation instructions for client and server have to be kept separate.

Preparation

- Description of the tree structure of the installed files
- List of all compressed and uncompressed files included in the delivery
- List of file systems to create with sizing
- Description of the logical and physical layout of the Oracle Database (if any), taking into account the following rules:
  - each database schema must contain at least two table spaces, one for the data and one for the index (e.g. if there are 2 oracle schemas, 4 table spaces are created: 2 table spaces for the data and 2 for the indices)
  - extra table spaces can be foreseen to address special needs (e.g. in case of partitioning)
  - for each table space, the initial size of the corresponding data file must be specified
- List of specific users, groups, and roles to be created and, for the database, the privileges to grant (the DBA role is not allowed). Generally, for web-based applications, the user accessing database objects through the web interface is not the owner of these objects in order to avoid accidental deletion of objects. This implies that the required privileges should be granted to the user accessing the database objects.
- Environment variables to define; variables will be used in order to avoid hard-coded values in the application source code or scripts; variable names should be meaningful.

Installation Procedure

- The installation procedure should be organised in clearly identified steps. Each step should have a sequence number, an application level description and technical comments.
- The installation procedure should be based on scripts to launch commands rather than on sequences of commands to type to avoid typing or cut-and-paste mistakes.
- The installation procedure should produce an installation log file.
- The installation procedure should be able to cope with the standard configuration of the host system.
- The installation instructions must offer the opportunity to arbitrarily and independently choose the installation, execution and data directories.
- The installation instructions should contain the location of the configuration files and a list and description of useful parameters.
- The configuration parameters should be grouped in a minimum of configuration files; global configuration files should be used to avoid multiple definitions of the same parameters or variables. A section of the installation instructions should list all configuration parameters, their description, and the specific values used for the installation in the environment of the Publications Office. If technically possible,

application parameters should be set in configuration files that are outside of application files (ear, war, jar, …).

- For easier installation, administration and trouble-shooting, the log files produced by the application should be contained in as few folders as possible; log4j or a similar logging framework should be used.
- The log files produced by the application and underlying systems must be properly managed; in particular, log files should be rotated daily and a mechanism should be foreseen to limit the number of log files.
- Oracle scripts should appropriately use the "commit" statement, together with the "whenever sqlerror exit rollback" and "whenever oserror exit rollback" statements in order to ensure application data consistency in case of errors.
- Ideally, all scripts needed to build the elements of the database (i.e. creation of the table spaces, tables, indices, triggers, users and roles including the privileges to grant access and to load data) should be delivered. An alternative is to deliver a script to create the table spaces and the users and the export (dump) of the data.

Post-installation

The installation instructions should contain a check list that details:

- the list of all files modified during the installation
- the list of periodic jobs to schedule
- a procedure to check the correct installation/working of the application: basic checks to be performed by the person in charge of the installation should allow to check if the application is behaving correctly without requiring a full functional validation.

System Uninstall

A detailed procedure how to uninstall the application should be provided that follows the same general remarks as the ones for the installation instructions.

5.4.2 *System Operation Manual*

This section describes the minimum contents of the system operation manual.

5.4.2.1 Hardware and Software Architecture

This section should contain:

- schemas that depict the hardware and software architecture
- a list of installed software including:
  - name of the product/tool
  - version
  - installation parameters (e.g. installation path, users, groups, environment variables...)
- a description of all file systems used by the application with their content and specific access rights (including any temporary space used)
- a description of the specific users, groups, and roles used by the application
- details of the network configuration including:

- interfaces used
- IP addresses
- virtual hosting
- IP and port bindings
- specific routing
- name servers
- LDAP servers used
- local name resolution

5.4.2.2 Configuration

- Application Start-up and Shutdown

  The sequence of steps to start and stop the application and any dependencies will be described. It must be possible to automate the start/stop procedure; the contractor is requested to deliver start/stop scripts for the application that can be integrated in the operating system service management framework.

- Configuration Files

  This section should specify:

- the location of the configuration files on the hard disk
- a list of useful parameters and their description
- how to modify these parameters

- Log Files

  This section should specify:

- a description of the contents of the log files
- the location of the log files on the hard disk
- how to set different log levels
- clean-up and archiving (regular purge and rotation)

5.4.2.3 User Management

- User Management

  This section should describe:

- how to create and delete users
- how to manage access rights and privileges
- information on the authentication mechanisms used by the application

5.4.2.4 Backup and Restore

- Backup Procedure

  A detailed backup procedure that includes at least the following elements must be provided:

- list of file systems/directories to backup (including pattern of file names to backup)
- scheduling/triggering schema
- pre- and post-processing commands to execute
- specific techniques to use (e.g. snapshot, hot backups...)

- Restore Procedure

  A detailed restore procedure covering the most common disaster situations should be provided. Special attention should be put on the following aspects:

- sequencing of the restore operations in order to minimise the downtime
- consistency checks to execute
- repair/resync procedures to execute
- system operation checks after restore

- Copy Procedure

A detailed procedure on how to copy an existing installation from one environment to another must be delivered. This procedure should allow creating a copy of for example the production environment to a test environment. This procedure must clearly indicate the parameters to modify in order to have a fully operational system in the test environment after completion of the copy procedure. The procedure must require as few manual interventions as possible.

5.4.2.5 Monitoring

- Monitoring

The contractor must deliver instructions on how to check the availability and the response time of the application; it should be easy to integrate these checks in an automated monitoring system (for example, a set of URLs to check); currently, the monitoring system at the Publications Office is based on Centreon/Nagios. The instructions and checks to be performed should cover all major components of the application, including components that the application depends on (like database server, middleware and other computer software that provides services to the software application beyond those available from the operating system).

This section should contain:

- a list of file systems to monitor with thresholds and critical values
- a list of processes to monitor with thresholds and critical values
- the description of monitoring and alert mechanisms included in the application
- a list of resources (e.g. URLs and relevant keyword/pattern to be checked in the case of a web application) that can be monitored in an automated and unattended manner by a system and network monitoring application

5.4.2.6 Application Management

The system operation manual should contain at least the following section with the suggested minimum contents.

- Administration Interfaces

  A detailed description of all available application administration interfaces will be provided. This section will at least include:

- how to access the administration interface
- a description of the functionality and features
- instructions for use

- Periodic Tasks

The tasks that have to be executed on a regular basis should be described in this section. It should be possible to automate these tasks as much as possible, for example by providing a script. For each task, the following information will at least be provided:

- description of the task
- procedure to execute
- how to schedule or trigger the task

Special attention should be put on the description of resource intensive and time consuming tasks like:

- data archiving
- data purging
- data reorganisation
- data consistency check
- data synchronisation
- data indexing
- statistics

- Database Management Tasks

  All specific database-related tasks that are not already described above should be described here.

- Best Practices and FAQs

  A description of best practices for dealing with common issues, troubleshooting procedures, hints and tricks, and a list of frequently asked questions should be included here.

6 Acronyms

AD	Active Directory
API	Application Programming Interface
APR	Apache Portable Runtime
AWS	Amazon Web Services
CNAME	Canonical Name
DB	Database
DBA	Database Administrator
DEVOPS	Development and Operations
DMZ	Demilitarized Zone
DNS	Domain Name System
DRP	Disaster Recovery Plan
EC	European Commission
ECAS	European Commission Authentication Service
FAQ	Frequently Asked Questions
HA	High Availability
HTTP	Hypertext Transfer Protocol
IS	Information System
ISP	Information System Proxy

ITIL	Information Technology Infrastructure Library
J2EE	Java 2 Platform, Enterprise Edition
JDBC	Java Database Connectivity
JDK	Java Development Kit
LAMPT	Linux, Apache, MySQL, PHP/Python/Perl, Tomcat
LDAP	Lightweight Directory Access Protocol
MFT	Managed File Transfer
NAS	Network Attached Storage
OP	Publication Office
OS	Operating System
OSB	Oracle Service Bus
PAAS	Platform as a service
RDBMS	Relational Database Management System
RDS	Relational Database Service
RH	Red Hat
RHEL	Red Hat Enterprise Linux
SAN	Storage Area Network
SFTP	Secure File Transfer Protocol
SQL	Structured Query Language
S-TESTA	Trans-European Services for Telematics between Administrations
SUDO	Super User Do
URL	Uniform Resource Locator
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPN	Virtual Private Network
WLS	WebLogic Server