

ANNEX T5

TECHNICAL ENVIRONMENT AND STANDARD OPERATING PROCEDURES OF THE PUBLICATIONS OFFICE

1.	PURPOSE OF THE DOCUMENT	3
2.	DISCLAIMER	3
3.	TECHNICAL ENVIRONMENT OF THE OFFICE	4
a.	<i>Introduction</i>	4
b.	<i>Network and telecommunications</i>	6
c.	<i>Storage and backup systems</i>	8
d.	<i>Workstations and peripherals</i>	9
e.	<i>Unix servers</i>	10
f.	<i>Windows servers</i>	12
4.	STANDARD OPERATING PROCEDURES	13
a.	<i>Management of software releases, bug reports and change requests</i>	13
b.	<i>Software deliveries</i>	13
c.	<i>Installations</i>	14
d.	<i>Technical tests</i>	15
e.	<i>Access to the Office's environment</i>	16
	APPENDIX A: CURRENT BACKUP POLICY AND PROCEDURES	17
	APPENDIX B: STANDARD FILE SYSTEM ORGANIZATION ON UNIX SYSTEMS (DIRECTORY STRUCTURE)	19

1. Purpose of the document

This document gives a general overview of the technical environment of the Publications Office - hereafter referred to as the Office – as well as some general rules linked to the technical organisation of the Office and applicable to all applications hosted at the Office.

2. Disclaimer

The information contained in this annex reflects the situation in force at the Office at the time of publishing of this document.

It does not commit the Office for the future evolution of its data processing and networking environments. It is mainly given here for clarity reasons. Changes can happen at any time and without any prior notification from the Office.

However, the environment to take into consideration for a specific project - especially the exact software versions - will be fixed at the very beginning of the project. This also includes the rules to apply by both parties in order to modify this environment.

The Office strongly advises the contractor/supplier to ask for clarification in case of doubt about the contents of this document. A meeting will anyway take place at the very beginning of the project to answer questions and, to a certain extent, examples of the expected documents could be provided.

3. Technical environment of the Office

a. Introduction

The 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 more tying on 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-one relationship between a server and its clients.

However, the same basic infrastructure is made available for both types of information systems, as described hereafter.

The Office recently deployed a DRP (Disaster Recovery Plan) making use of 2 different geographical sites and based on the following principles:

- The DRP is conform to the Contingency Plan of the Office
- The data replication between the 2 sites is synchronous
- Both sites are hosting "active" applications

Unix is the recommended environment for production systems while office automation systems are normally hosted on Windows servers.

The Office fosters professional methods of managing systems and therefore implements monitoring tools for critical systems and produces statistics on the use of resources and on the quality of service provided.

The Office promotes the implementation of the three tiers architecture, using thin clients, application servers and mainly Oracle databases because of performance, scalability and flexibility reasons.

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

- All web-based applications must allow the deployment and the correct operation behind any http reverse proxy chain.
- All applications must allow virtual hosting i.e. the binding of the application to only some of the IP addresses/hostnames of a multihomed server
- All applications must allow easy integration in the DRP of the Office
- All applications must be compatible with an MS Windows 2003 terminal server architecture

The Office's core business applications are tested on dedicated machines, before the production is spread over several production servers tightly interconnected.

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

- Network and telecommunications
- Storage and backup systems
- Workstations and peripherals
- Unix servers
- Windows servers

The hardware/software architecture to use within the framework of a project is generally proposed by the contractor/supplier. However, the Office will always first validate this architecture in order to evaluate the consistency of the proposal with the environment of the Office.

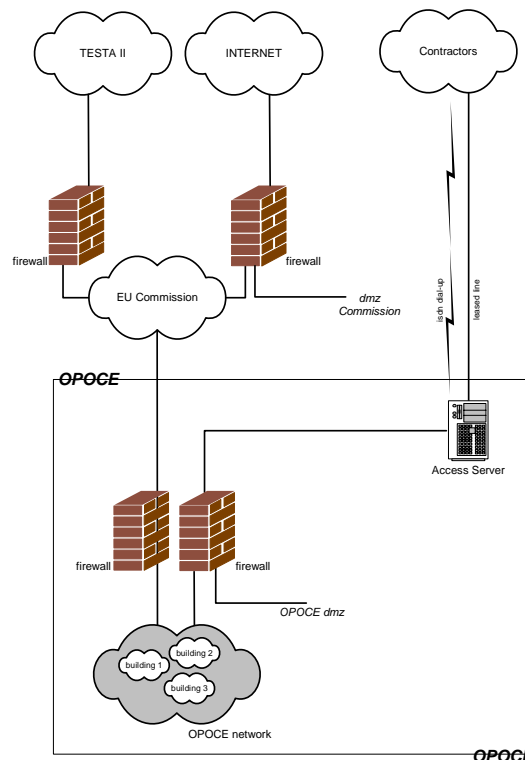
b. Network and telecommunications

The Office's staff is spread over several buildings. The Office's buildings use a unique cabling system for the telephone as well as for the TCP/IP network for data. This structured cabling system uses copper connections category 5 (or higher) to desktop computers, and optical fibre connections for backbones and to high-end servers with a throughput up to 1 Gbps. The patching mechanism relies on the AT&T 110 standard and the RJ45 standard. Both networks deal with about 1000 telephones and 1000 Ethernet devices, and will continue to grow with potential new sites.

The telephone system is based on three Siemens Hicom telephone exchanges, carrying out all dial-up services, regular dial-up connections to contractors, and even backup connections for leased lines. The Office telephone network is part of the European Commission telephone network, but it's fully managed by the Office network team. Videoconferencing is also made available to the end-user.

The Wide Area and the Metropolitan Area Networks (between buildings) use leased lines, either over dark fiber high speed connections (1 Gbps) or over 2 Mbps connections to the PSTN¹. IP multiplexers are deployed to share high speed lines between voice and data.

The following picture shows the interconnection of the TCP/IP network of the Office.



In particular, connections with the contractors working for the Office use 2 Mbps leased lines or the EURO ISDN network for lower speed data transmission. These dial-up services are available through Basic Rate or Primary Rate interfaces of a TCP/IP router. The ISDN dialup service is in many installations also serving as a backup solution for the leased line connection.

¹ Public Switched Telephone Network

The most common application for the Office remote access service is file transfer (FTP). Files are transferred via a FTP gateway installed in the Office DMZ.

The TCP/IP network is also interconnected with the network of the European Commission that is connected to the Internet and to the TESTA II network (Trans-European Services for Telematics between Administrations). The TESTA II network allows the Office to establish private connections with most of the national Administrations of the EU Member States and most of the EU Institutions. FTPStore is a service offered by DIGIT/Commission for sending and receiving files over the public Internet or via the TESTA network. FTPStore is not offering any notification or monitoring of file movements.

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 sub-mentioned networks are interconnected through stateful inspection firewalls.

Three Windows servers running the MRS software offer fax services and are part of automatic production chains in the daily publishing process. The fax servers are running MRS of Cycos. These systems also integrate the voice mail system with the Office e-mail server in the computer telephone integration.

c. Storage and backup systems

The Office has decided to strengthen the quality and reliability of the data by centralising the administration of the data storage service on storage systems such as EMC² (mainly for production systems) and Sun Microsystems T3, 6130 and 6140 (mainly for non-production systems).

Generally RAID-1 is used for production systems while RAID-5 is used for non-production systems. Around 20 Terabytes of data are currently made available to end-users and applications. This amount is steadily increasing together with the number of users and with a wide usage of electronic documents in the Office; the increase rate is about 30% per year.

The backup system is currently under revision to implement an improved backup strategy based on LAN-less and server-less procedures used in a SAN¹ context.

The Office will focus its storage and backup strategy on:

- Online storage consolidation (SAN)
- Storage on demand
- Advanced centralised backup
- Archiving

The current backup policy and procedures are described in **Appendix A**.

¹ Storage Area Network

d. *Workstations and peripherals*

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

Type	Product/Version
Operating system	MS Windows XP SP2
Office automation suite	MS Office 2003
Web browser	MS Internet Explorer 7 (with 128 bits encryption)
XML tools	XMLSpy (development) XMetal (Authoring)
Reporting tools	Business Objects 5.1 – 6.x
Mail client	MS Outlook 2003
Connection middlewares and runtimes	SQL*Net 2.3.4 NET 8.1.7 - NET 9.x ODBC
Anti-virus	McAfee VirusScan 8.1

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

The Office can sometimes have to apply hotfixes to ALL workstations without any prior notification if deemed necessary (e.g. for security reasons).

A small number of workstations are Unix-based or Macintosh-based (for DTP).

In terms of hardware, the workstations are ranging from Pentium III 850 MHz to Pentium 4 2.5 GHz with 128 to 512 MB of memory.

Various types of peripheral equipment are also installed, such as local and remote printers, faxes, scanners, CD burners, etc.

All workstations are clients of the Landesk Workstation Manager server, which allows remote control and software/hardware inventory.

Local data on the workstations is not backed up and users are therefore encouraged to use only shared resources for storing valuable/persistent information.

e. *Unix servers*

The Office computing centres host more than 30 Sun Microsystems servers mainly based on the Sunfire platform, ranging from the SF V210, V890 up to the SF25000, on T5220/T5240 Enterprise servers and on M5000 Enterprise servers. The CPU's are mainly SPARC 64-bits Ultra Sparc III+ (1200 MHz), Ultra Sparc IV (1500 MHz) and Ultra Sparc VI (2.15 GHz) and T2 (1.2 GHz). These servers host more than 20 virtual servers (container/zone).

In terms of software, the following table gives an overview of the main products used:

Type	Product/Version installed	Product/Version recommended for all new developments
Operating system	Sun Solaris 2.6 - 8 – 9 and 10	Solaris 10 in zone/container
DBMS	Oracle from 8.0.5 to 10g Rel 2 ADABAS 5.1	Oracle 11gR1 Character set: AL32UTF8
Text retrieval	Oracle Intermedia/Context	Oracle Intermedia/Context
Web servers	Sun ONE Web Server Apache 2.x	Apache 2.x
Application servers	Adobe ColdFusion MX6.1 Apache Tomcat 4.x, 5.x JBoss 2.4 Oracle iAS 9.0.2	Adobe ColdFusion MX 8 Apache Tomcat 6 JBoss 4.x BEA WebLogic 10 ¹
ERP	Oracle Financials 11i	Oracle Financials 11i
Programming languages	Java 1.4.x/1.5.x Natural 6.2	Java 1.6.x
Scripting languages	Perl 5.6 sh/ksh	Perl 5.8 sh/ksh
Workflow/Document management systems	OT Livelink 9.1 SP3 DORIS 3.51SP3 Documentum Content Svr 5.2.5 (incl CRS/DTS module)	Documentum Content Svr 6.5 or above Other third party products
Reporting tools	Business Objects WebI 2.7 (hosted at DIGIT/DC)	Business Objects WebI 6.5 (hosted at DIGIT/DC)
XML related tools	XSV (XSMML Schema) RXP (DTD) SAXON (XSLT)	XSV/XERCES (XML Schema) RXP (DTD) SAXON/XALAN (XSLT)
Middelware	WOOD ²	WOOD ³
Monitoring tools	Nagios TeamQuest	Nagios TeamQuest
Backup	Networker 7.2	Networker 7.5

More than 70 Oracle instances (production + test) are currently installed.

Monitoring and reporting tools have been built around small third-party products and in-house developments.

¹ The use of WebLogic must be justified and must be validated by the Office

² proprietary tool based on Perl

³ proprietary tool based on Perl

For critical applications, Sun Cluster installations have been built.

In general, application processes exchange data either by mail or by FTP using a dedicated and in-house developed proprietary tool (WOOD – *Worldwide Object Dispatcher*).

Besides the exchange of data (files) between processes – eventually running on distinct servers – this tool allows the triggering of processes based on the arrival of a file in a predefined directory. The tool is written in Perl and uses normally FTP as underlying protocol but could theoretically use whatever standard file transfer protocol (e.g. scp). Due to the asynchronous character of this tool, the WOOD cannot guarantee the eventual sequencing of the data exchanged. If sequencing is an issue, it must be managed at application level.

The Office strongly advises the contractor/supplier to ask for practical implementation guidelines before to start any development that could require integration/interaction with the WOOD.

The standard file system organization (directory structure) for the Unix servers is described in **Appendix B**.

Direct dependencies between servers (e.g. NFS mounts, DB links, ...) are generally prohibited.

f. *Windows servers*

The Office computing centre hosts about 60 MS Windows servers in a Windows 2003 Active Directory domain. Some servers are virtual machines running on VMware. All new physical servers are 2 or 4 CPU servers. Some are installed in cluster mode (for critical data and processing).

Besides standard functions like user authentication, roaming profiles, home directories, shared drives (public/group) and print services that are spread over several servers in order to improve reliability and performance, the Windows servers also host services like email, fax gateways, standalone applications and other small information systems (in-house developments, automated tasks using Microsoft Office products, small SQL server DB, IIS servers with Coldfusion, etc.).

Type	Product/Version installed	Product/Version recommended for all new developments
Operating system	MS Windows 2000 Server SP4 MS Windows 2003 Server standard edition/enterprise edition	MS Windows 2008 Server standard edition/enterprise edition
OS virtual servers	VMware ESX Server 2.51 – 3	VMware ESX Server 3.5
DBMS	MS SQL Server 2000	MS SQL Server 2005
Web servers	IIS 5 – 6	IIS 6 or +
Application servers	ColdFusion MX6.1 -7	ColdFusion MX8
Programming languages	Visual Basic .Net	Visual Basic .Net
Mail servers	MS Exchange 2003	MS Exchange 2007
Reporting tools	Business Objects WebI 2.7	N/A
Backup	Networker 7.2	N/A
Anti-virus	McAfee VirusScan 7.1 Trend Micro Scan Mail	N/A

For business critical reasons, Windows clusters host the Exchange server with about 1000 mailboxes, the shared drives and the disk spaces accessible to all end-users. All production data is stored on the SAN.

Regarding security on the server side, anti-virus checking is performed on the Microsoft Exchange mailboxes and regularly on file systems.

4. **Standard operating procedures**

a. *Management of software releases, bug reports and change requests*

Within his quality plan, the contractor/supplier must define a detailed and unambiguous numbering scheme for the software deliveries.

The contractor/supplier must also propose procedures for the management of bug reports and change requests. These procedures must allow to unambiguously identify every bug and every change request.

b. *Software deliveries*

The contractor/supplier has to include in any software delivery submitted to the Office a **release note** (in electronic format) containing following information:

- Project/application name concerned.
- Unambiguous identification¹ of the version of the delivered software.
- Version of the project/application on which the delivered package has to be installed
- Unambiguous identification² of the bugs/change requests concerned (if applicable i.e. mainly in case of patch)
- approximate uncompressed size of the delivery
- Reference to the installation instructions (as defined in the annex “*Technical documentation guidelines*”)

The software will be delivered in a standard archive format and with compression (tar, gzip, compress, zip, ...)

The compressed archive files will be created using relative path in order to allow decompression in whatever directory.

The delivery will include all the information necessary to check that acceptance testing has been performed by the contractor/supplier BEFORE delivering the software. This includes :

- Test procedures/test cases executed
- Test data used
- Test results (execution report)

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

² In accordance with the procedure defined in the approved quality plan

c. Installations

In principle, for each application, two distinct environments are set up at the Office's premises: a **test environment** and a **production environment**. This does not preclude the fact that several applications can share the same production or test environment.

The hardware installations (including OS installations) remain the exclusive prerogative of the Office.

No software installation in the production environment will be allowed without prior validation in the test environment.

No development environment will be set up at the Office's premises. The Office strongly advises the contractor/supplier to set up at its 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/supplier to make sure that software deliveries will run correctly in the Office's technical environment.

Any software installation (including patch installations) requires the software delivery to be in accordance with the rules mentioned in the previous paragraph (software deliveries).

The minimal contents of the **Installation Instructions** are described in the annex "*Technical documentation guidelines*".

The contractor/supplier is free to add any information deemed useful.

The installation instructions must offer the opportunity to arbitrarily and independently choose the installation, execution and data directories.

All software installations are normally done by the Office's staff with the assistance/support of the contractor/supplier, but depending on the complexity of the installation to perform (as evaluated by the Office), either it does not require the support of the contractor/supplier or the support of the contractor/supplier is formally required.

A script allowing checking the availability and the response time of the application - as seen by the end-user - must be delivered. The triggering mechanism of this script must allow its easy integration in whatever integrated monitoring system.

A detailed procedure allowing copying the production environment to a test environment must be delivered. 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.

In order to ensure smooth installations and to evaluate the need for assistance/support, the installation instructions will be - unless otherwise stated - delivered to the Office, 5 days before the official software delivery date i.e. 5 days before the start of the official installation period.

The Office strongly advises the contractor/supplier - when possible (e.g. web application) - to setup as soon as possible a remote access to one of his environments in order to give the Office's staff the opportunity to have a first impression of the application being developed before any installation at the Office's premises.

d. Technical tests

Prior to be put in production and prior to any official acceptance, the Office will conduct specific technical tests in order to evaluate the *"manageability readiness"* of the application.

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

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

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

- Application start and stop procedure
- Application backup and restore procedure (including consistency checks after restore)
- Disk space usage
- Localization of the log files
- Operational periodic tasks (data reorganization, purging, archiving, indexing...)
- Correct working of the interfaces
- Virtualization capabilities (application ability to move from one server to another one) in a Disaster Recovery Plan context

Not only the effectiveness of the procedures but also their efficiency (duration) must be tested. Their impact on the overall performance of the system must be evaluated too. The technical tests must be conducted with a significant amount of data in order to evaluate the impact of volumes on performance, efficiency and effectiveness.

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 the annex *"Technical documentation guidelines"*.

The contractor/supplier is free to add any information deemed useful.

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

e. Access to the Office's environment

No direct access (telnet, ftp...) to the Office's environment (production or test) will be granted to the contractor/supplier.

Specific interfaces (e.g. Web/CGI) have to be developed for the administration of the application (e.g. periodic follow-up) and/or the production follow-up. Especially the access to interfaces (data exchange directories) must be controlled by the application in order to validate the contents of the data exchanged and to allow the "replay" of data transfers.

APPENDIX A: CURRENT BACKUP POLICY AND PROCEDURES

1. Categories of backup jobs

The following backup jobs are distinguished:

- full backup: backup of all data.
- level backup (differential backup): backup of all data changes since the last lower level backup (full backup corresponds to level 0). Only one level is currently defined. Thus, level backups are always performed with respect to a full backup.
- incremental backup: backup of data changed since the last backup, independently of its type.

2. Backup policy:

- All Unix and Windows servers are clients of one Unix central backup server running Legato Networker. This server stores the backup sets on one tape library. This library is able to handle up to 1200 tapes and is connected to the SAN.
- The jobs are scheduled between 8 p.m. and 2:30 a.m. The effective backup might start only after the scheduled pre-processing jobs (e.g. database snapshot) have been finished.
- Full backup jobs are generally run once per week and are distributed between all days of the week.
- Level jobs (i.e. differential jobs) are run 4 times per week between Monday and Friday every day the full backup job is NOT run.
- Incremental jobs are run once a day, except the day of full backup.
- The browse policy (direct access to file and directory information of the backed up file systems) and the retention policy are generally set to 3 months.
- Software compression is used on the client.
- Hardware compression is used at tape device level.
- No separation is made between full, level (differential) and incremental job tapes.
- The backup tapes remain in the tape library. The reusability of these tapes is controlled by the software.
- The tapes are cloned.
- Cloned tapes are removed regularly by the computer center operators and placed in the safe.
- The backup software provides appropriate features to assure that the data of a tape volume can be safely overwritten.

3. Backup of Oracle databases:

One or a combination of the following techniques is used:

- **Logical database backup:** Oracle export dump files are generated while the database is in restricted access mode. These files are written on the file system of the corresponding servers. The database is then stopped and backups of the server file systems are made within the global backup framework.
- **Physical cold backup:** The database is stopped and backups of the server file systems are made within the global backup framework.
- **Physical hot backup:** An open database backup is done (alter tablespace ... begin backup ; alter tablespace ... end backup).
- **RMAN** (oracle recovery manager): All standard features of RMAN are used.

- **"Freeze"/Snapshot techniques:** In order to limit the unavailability of the database during the physical cold backup, the database is stopped only during the time needed to make a "freeze" or a snapshot of the file systems hosting the database. The "freeze" or the snapshot is then available for tape backup or for restore purposes.
 - A **snapshot** consists in attaching an additional mirror volume for the desired directories and to make a synchronisation between the additional mirror and the mirrored volume.
 - A **"Freeze"** consists in keeping the state of a file system at the given "Freeze" time by means of a second file system with exactly the same directory structure. Prior to any modification of the "frozen" file system, the unmodified data is copied to the second file system. By combining both file systems, the "Frozen" state is always available.

4. *Implementation:*

It is the responsibility of the Office to define the objectives to reach in terms of availability/unavailability of the application.

Based on these requirements, the contractor/supplier must define the backup procedures to implement in order to fulfil the requirements. Ideally, the contractor/supplier must base these backup procedures on techniques the Office's staff is familiar with (see contents of this appendix).

If the wishes of the Office in terms of availability cannot be satisfied with techniques the Office's staff is familiar with, the contractor/supplier must provide a full description of the procedures and techniques to use so that the implementation by the Office's staff could happen smoothly.

APPENDIX B: STANDARD FILE SYSTEM ORGANIZATION ON UNIX SYSTEMS (DIRECTORY STRUCTURE)

In order to ease the co-existence of applications on the same server and to ease the potential move of an application to another server, applications are generally installed under **/applications/application_name/** where *application_name* refers to the name of the application.

This level is then subdivided into:

- **/applications/application_name/users** which is itself subdivided into:
 - **/applications/application_name/users/system**: directory simulating the root directory for the application. Specific products used by the application are installed here (e.g. the web server is installed under **/applications/application_name/users/system/apache/**, the application server is installed under **/applications/application_name/users/system/tomcat/**).
 - **/applications/application_name/users/oracle**: Oracle application, oracle environment, and oracle admin directory.
 - one or more directories **/applications/application_name/users/user_name**: home directories of the users *user_name* used by the application. These directories are linked to **/home/user_name**. Generally, there is only one directory **/applications/application_name/users/user_name** and *user_name* is identical to *application_name*.
- **/applications/application_name/xchange**: root of interfaces (in case of exchange of data with remote applications). The following specific structure is used:
 - **/applications/application_name/xchange/remote_application_x/(sublevel if necessary)/in** for **incoming data** and
 - **/applications/application_name/xchange/remote_application_x/(sublevel if necessary)/out** for **outgoing data**

where *application_name* refers to the name of the application and *remote_application_x* refers to the name of the remote application.

e.g. **/applications/eub/xchange/eudor/in** is used for the data flow exchange from eudor to eub.

/applications/eub/xchange/gescom/gcb/out is used for the data flow exchange from eub to the gcb part of gescom.

- **/applications/application_name/oradata**: Oracle datafiles.
- **/applications/application_name/oraexp**: Oracle exports
- **/applications/application_name/oralog**: Oracle online archive logs.

Deviations from this description are possible but the Office must first validate the deviations.

```

/applications
|
-- /application_name
|
-- /users
|
-- /system (to install by Publications Office)
|
-- /init.d                               Start/stop scripts
|
-- /product_1 (e.g Tomcat - Apache - ...)
|
-- /...
|
-- /product_n
|
-- /oracle                               oracle binaries
|
-- /user_name_1 (link to /home/user_name_1)
|
-- /... (if required - link to /home/...)
|
-- /user_name_n (if required - link to /home/user_name_n)
|
-- /oradata                             oracle datafiles
|
-- /oraexp (if required)                 oracle export
|
-- /oralog                              oracle logs
|
-- /data_1 (e.g. Documentum filestore)   appli data
|
-- /data_... (if required)              appli data
|
-- /data_n (if required)                appli data
|
-- /xchange
|
-- /remote_appli_1
|
-- /in
|
-- /out
|
-- ... other interfaces ...

```

(The names of the different filesystems are in bold)