



EUROPEAN COMMISSION
DIRECTORATE-GENERAL
INFORMATICS
Infrastructure Directorate
Data Centre



European Commission

Overview of the usage of the Information System Hosting Services of the Data Centre

Date: 25/11//2005
Approved by: TE KOLSTE Georges-Eric
Reference Number: 3771V.3

TABLE OF CONTENTS

1. INTRODUCTION	4
1.1. Overview	4
1.2. How to read this document.....	4
1.3. Mission and Services of the Data Centre.....	5
1.4. Organisation of the Data Centre	6
1.5. Contact.....	6
2. ARCHITECTURE.....	8
3. MANAGEMENT PROCEDURES.....	10
3.1. Introduction	10
3.2. Application Lifecycle	10
3.3. Hosting Request and the Change Management Process	11
3.4. Incident Management for ISHS	12
3.5. Product Management.....	16
3.6. Data Protection	17
3.7. Security.....	17
3.8. Service Level Agreement	18
3.9. Preventive maintenance	19
3.10. Configuration data.....	19
3.11. Acceptance testing.....	19
3.12. Reception Procedure - Production check list for new IS and major releases	20
3.13. Publishing guidelines on Europa and IntraComm	21
3.14. Monitoring.....	22
3.15. User Statistics – Data Centre Resources Usage.....	23
3.16. User statistics – Web Consultation.....	23
4. INFRASTRUCTURE.....	24
4.1. Introduction	24
4.2. Servers	24
4.3. Time Servers.....	25
4.4. Domain Name Servers.....	25
4.5. ORACLE Name Servers.....	25
4.6. DMZ	26
4.7. Storage.....	26
4.8. Contingency.....	28
4.9. Load Balancing.....	28
4.10. Backup Service for the local servers in the DGs	28
4.11. Environments.....	30

5. TECHNOLOGICAL ENVIRONMENTS	31
5.1. Introduction	31
5.2. Access Layer	31
5.3. Web Servers.....	31
5.4. Coldfusion	31
5.5. WebLogic	38
5.6. Business Objects.....	41
5.7. Oracle	43
5.8. UNIX/LINUX.....	48
5.9. Windows.....	50
5.10. SMTP relays for E-mail applications.....	50
5.11. Mailbox access for E-Mail applications	52
5.12. ECAS & LDAP	54
5.13. Active Directory	54
5.14. Corporate Web Content Management System.....	55
5.15. SAS.....	55
5.16. FTP store	55

Table of Figures

Figure 1 - Organisation Chart of Data Centre -	6
Figure 2 - Data Centre Architecture -	9
Figure 3 - Incident Management -	13
Figure 4 - Organisation of the ISHFO -	14
Figure 5 - ColdFusion Hierarchy -	36
Figure 6 - WebLogic Domain -	40
Figure 7 - Failover Architecture -	45

Document History

Version	Date	Comment	Modified Pages
1.000	01/09/2005	Creation of the Document.	ALL

Contacts

Name	DG/Service	email	Phone Number
General Mailbox	DIGIT/A/3	digit-dc-guidelines@cec.eu.int	N/A

General Disclaimer

This document is the property of the European Commission.

The information provided in this document allows all interested parties to see the guidelines for Information System Hosting. The content of this document may be subject to rapid changes for a variety of reasons. The European Commission can not be held liable for the consequences of any reliance on the information provided or for any inaccuracies in such information.

All or part of the document may not be copied without prior agreement by the European Commission and without specific reference to the latter as well as the source of the extract.

The master document is written in English, if translations of this document exist in other languages, the reference will be this document.

1. INTRODUCTION

1.1. Overview

The Commission communication on IT Governance ([SEC\(2004\)1267](#)) provides a framework for increased collaboration between DG DIGIT and the DGs. Within this context this document contributes to developing a common approach in the specific area of hosting information systems in the Data Centre.

The Data Centre provides a range of services to Directorates General, other European Institutions and Agencies. These services include the hosting of both corporate and specific information systems (Information Systems Hosting Service).

The purpose of these guidelines is to define, for the Information Systems Hosting service, the range of services offered by the Data Centre and the conditions under which they are made available to Clients.

The Data Centre hosts a very large number of information systems. In order to make this manageable it is evident that a certain degree of standardisation and discipline is necessary, particularly in order to ensure that the appropriate performance and availability levels are maintained, and that common resources can be shared conveniently between different systems, without adversely impacting the services offered to others.

The Data Centre is currently adapting its processes so as to implement ITIL recommendations for Service Delivery and Support.

The complete set of guidelines will be reviewed and republished as a whole by the Data Centre every semester.

These guidelines provide the framework for Information System Owners, Customers and their development teams to define and manage their requests for services of the Data Centre. They may be made available to subcontractors on condition that they have signed a “Non Disclosure Agreement” (normally part of the general terms and conditions of the contract).

These guidelines apply to requests for hosting initial and major releases of mission critical Information Systems and will be completed with specific procedures for managing minor releases.

The IT strategy of the Commission is elaborated in detail since 2005 by DG DIGIT in close collaboration with the services as foreseen in the Communication. The Data Centre’s procedures will be aligned with these decisions particularly concerning the annual “Schéma Directeur” exercise and development methodologies. These guidelines will be continuously updated to reflect these changes.

1.2. How to read this document

If you are an Information System Owner or a Supplier, you will be mostly interested in Chapters 1, 2 and 3: Chapter 1 gives a brief overview of the Data Centre. Chapter 2 describes the Data Centre’s architecture. Chapter 3 covers the management procedures to be followed when interacting with the Data Centre (how to introduce a request for hosting, how to report an incident,).

If you are an Information System Developer, you will be mostly interested in Chapter 4 and chapter 5. Chapter 4 explains the infrastructure currently available at the Data Centre. Chapter 5 covers the technological platforms supported by the Data Centre and gives guidelines to be followed during the application development lifecycle.

1.3. Mission and Services of the Data Centre

The Data Centre is a service oriented unit providing the Commission with a corporate, secure, reliable and high performance IT infrastructure to support 24h/24 the information systems and services needed to implement the e-Commission.

The services provided by the Data Centre are:

- Corporate information system hosting particularly for the mission critical systems which support Community policies and which underpin the Commission's internal administration;
- Web dissemination infrastructure particularly for EUROPA, the European Union's presence on Internet and IntraComm (the European Commission's corporate web site);
- Corporate E-mail including calendar and workgroup facilities and with virus protected connectivity to the external world;
- Corporate infrastructure for applications including office information services.

These services include back-up/restore and for a number of critical systems a redundant infrastructure to ensure business continuity in the event of a major incident.

The Data Centre's mandate is to ensure the availability, reliability, performance and scalability of these services with the highest levels of security and to foresee their evolution with cost-effective, state of the art solutions.

The Data Centre provides these services in partnership with the DGs concerned, based on Service Level Agreements with appropriate quality indicators and associated monitoring and measurement procedures.

These services are supplied to the end users with the participation of the other units of DG DIGIT's Infrastructure Directorate, particularly concerning telecommunications, so as to ensure end-to-end performance and availability.

1.4. Organisation of the Data Centre

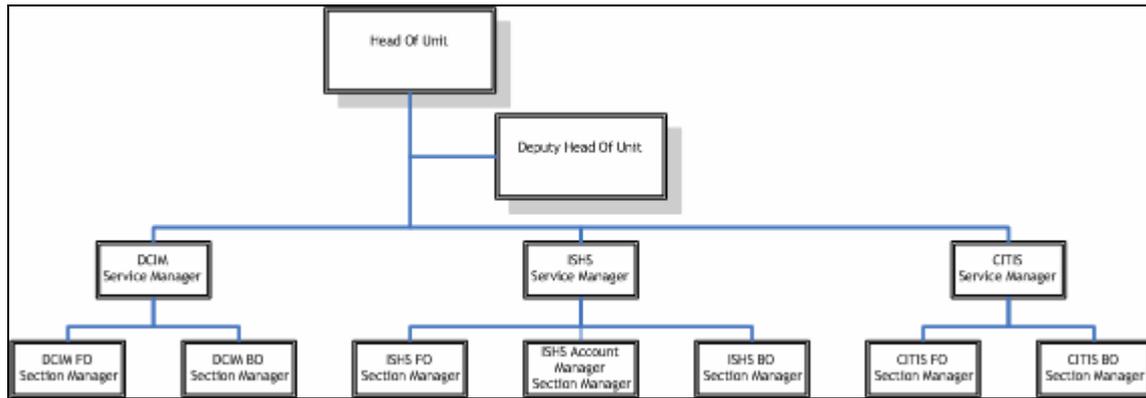


Figure 1 - Organisation Chart of Data Centre -

- DCIM F.O. (Data Centre Infrastructure Management Front Office): focused on day to day operations and Incident Management related to the Data Centre Infrastructure.
- DCIM B.O. (Data Centre Infrastructure Management Back Office): focused on Data Centre. Infrastructure architectures, strategy and planning.
- ISHS F.O. (Information System Hosting Services Front Office): focused on day to day operations and Incident Management related to Information Systems.
- ISHS B.O. (Information System Hosting Services Back Office): focused on Application architectures, strategy and planning.
- ISHS Account Management is responsible for the relationship with the DGs regarding their Information Systems.
- CITIS F.O. (Corporate IT Infrastructure Services Front Office): focused on day to day operations and Incident Management related to Corporate IT Infrastructure Services.
- CITIS B.O. (Corporate IT Infrastructure Services Back Office): focused on Corporate IT Infrastructure Services architectures, strategy and planning.

1.5. Contact

In line with the implementation of a process oriented organisation, the Data Centre has put in place a number of specialised Points of Contacts for processing new requests and for managing incidents.

1.5.1. For introducing a new request

When introducing a new request to the Data Centre, for instance, the creation, modification or transfer of a hosting environment, clients are invited to fill in a request form in the on line system MIRELLA which is the central management tool used by the Data Centre on Information Systems. In case of questions or if assistance is needed, Customers can contact the Account Manager assigned to their DG.

1.5.2. For reporting an incident or requesting information

The Service Desk will ensure follow up of Service Requests or Incidents initiated by creating a ticket in “ Service Centre”, the on line Service Management Tool used by the Commission, either by the Local Helpdesk, the IS owners, or the Central Helpdesk.

2. ARCHITECTURE

The Data Centre offers services which are hosted at two sites with the objective of implementing both failover and emergency backup facilities between them.

The schema overleaf represents the 4 layer architecture (access, presentation, application, data base) based on which all Information Systems will be implemented on a secure, redundant, cost effective infrastructure. The Europa, IntraComm and DG Web Servers follow the same architecture.

The servers are divided between the two sites and, in case of an emergency, a single site may take over the critical production workload and ensure continual operation in degraded mode. The architecture is based on shared storage with both Network Attached Storage and Storage Area Network configurations. The SAN storage is split between the two sites with synchronous remote mirroring between the SAN arrays via high-speed Dense Wavelength Division Multiplexing connections. The NAS storage is also split between the two sites but with asynchronous remote mirroring.

The servers deployed for the hosting service in the Data Centre are mainly UNIX systems (including LINUX). Some hosting on WINDOWS servers is also provided, but this is normally only for very specific projects.

The application architecture may use the application components provided by CITIS (eg: email, terminal services, active directory)

Network connectivity is through the internal Commission LAN (SNET) and WAN (via TESTA or the Internet). Access control is carried out using LDAP (username/password authentication). Stronger authentication is now being implemented with the ECAS project. Access control can also be carried out using SSL via a Remote Access Service.

Backup of the Data Centre systems is carried out using cartridges robots.

All systems are monitored with an integrated call-out system to technical support teams. Support staff is, where required, on call 24/7.

In order to guarantee performance and availability levels, all components of Information Systems must be hosted in the Data Centre.

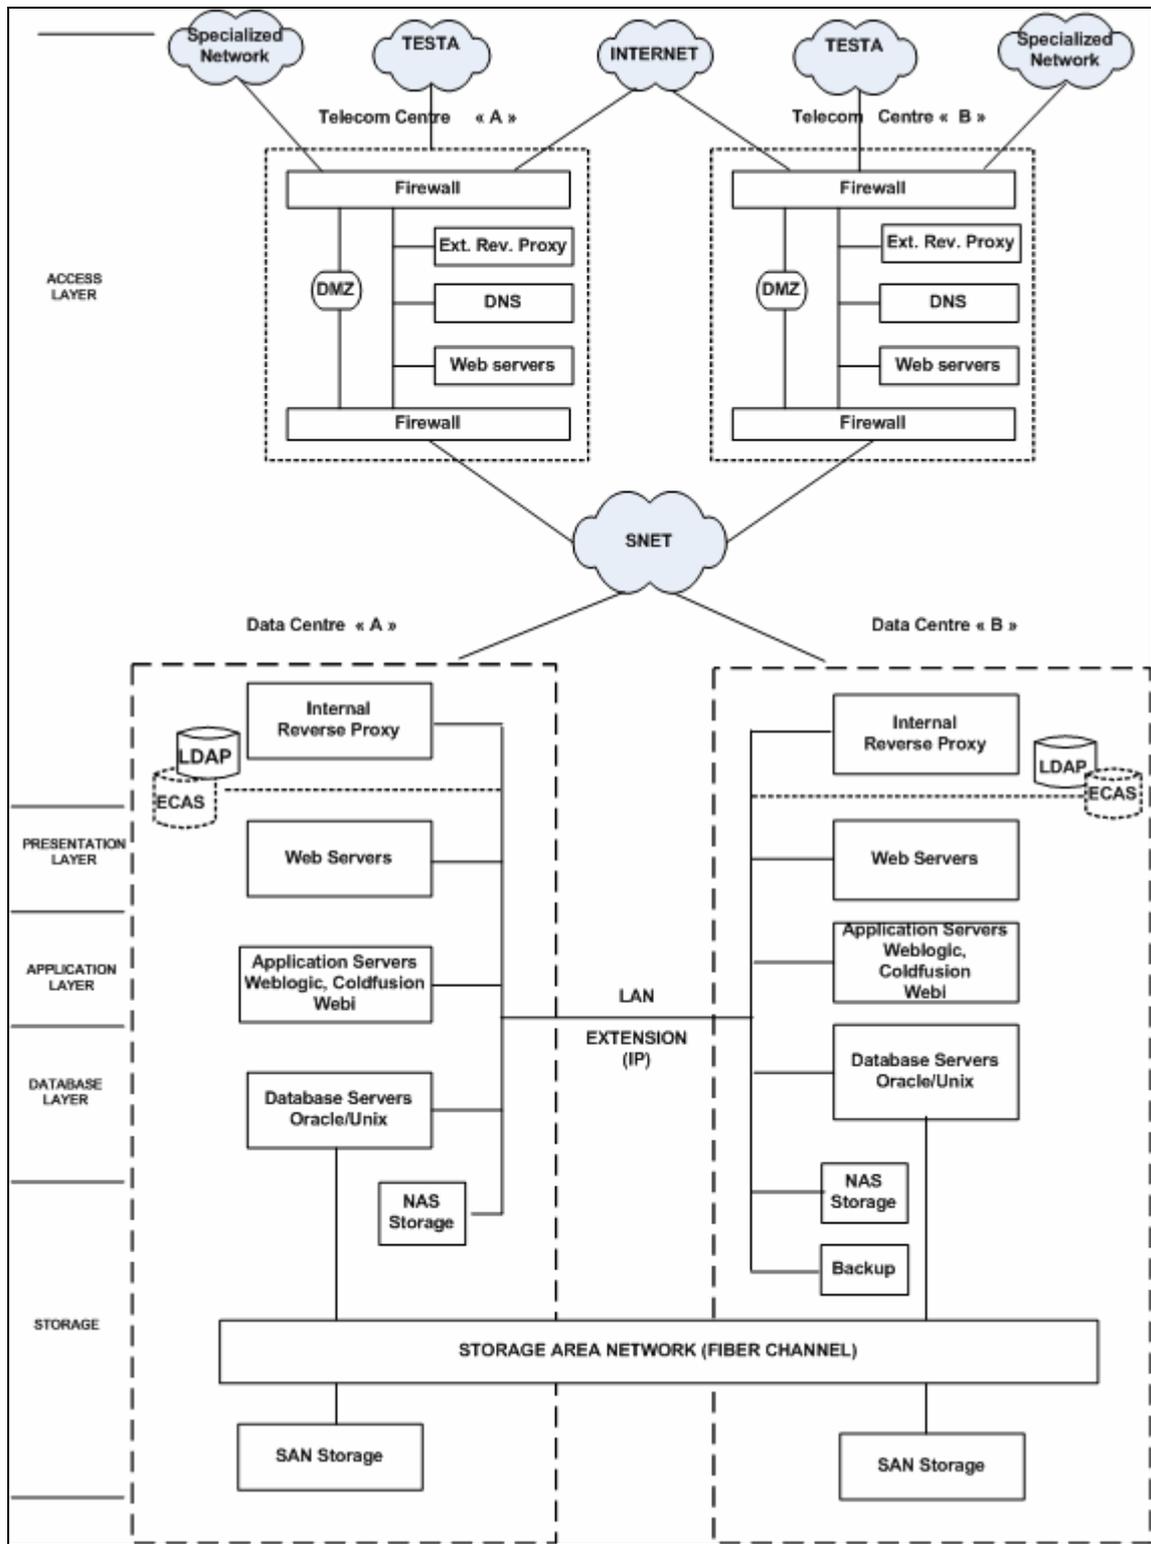


Figure 2 - Data Centre Architecture -

3. MANAGEMENT PROCEDURES

3.1. Introduction

This chapter explains the management processes and procedures that are followed during the application lifecycle so as to ensure that a hosted application is compliant with the strategic guidelines defined by the Commission and DIGIT in different areas including data protection, security, product management and architecture, service level management, configuration management and change management. The chapter is mainly intended for Information System owners in the DGs and Customers.

3.2. Application Lifecycle

Information System development is an engineering process with a lifecycle including several phases.

For major Information Systems, the following phases are important:

- Problem statement;
- Pre-analysis;
- Feasibility study including a proof of concept;
- Development;
- Testing and Training;
- Operation;
- Phase out / renewal.

Once a new IS is hosted, there is a loop between the “Testing & Training “phase” and the “Operation” phase for implementing new releases.

These phases will be adapted in line with the RUP methodology in the next version of this document.

It is desirable that the Data Centre is involved, at the earliest stage of Information Systems projects so that:

- The architecture may be defined in line with these guidelines when the hosting request is presented to the Data Centre;
- The required infrastructure for the Information System may be acquired and installed in time;
- Acceptance procedures may be agreed and actions linked to performance and functional tests be planned;
- The SLA can be signed before the operation phase starts.

3.3. Hosting Request and the Change Management Process

3.3.1. Hosting Request

In line with the Commission decision on IT Governance, DGs should include initial requests for hosting information systems in the Data Centre in their annual “Schéma Directeur” exercise.

The subsequent detailed hosting request should be introduced by the Customer responsible as early as possible in the development life cycle.

An application schema is required so as to understand the application architecture, the software components and data access paths. It is very important for the Data Centre to understand the architectural requirements so that servers can be dimensioned adequately at each level of the architecture (access, presentation, application, and database).

Depending on the request, additional technical information may be necessary to ensure that it is well understood by the Data Centre and can be handled with maximum professionalism. Indeed, it is crucial for the Data Centre to receive information on the architectural requirements and products, on the evolution of disk usage and workload, on the criticality of the application (availability, fail over, contingency) and potential security issues (external accesses, ...) so that appropriate solutions can be implemented and a reliable capacity plan can be prepared in terms of infrastructure and staffing.

A minimum delay of 2 weeks is required by the Data Centre to deliver a new environment in the simplest case. This delay is of course dependent on the request (need for additional equipment to be delivered...) and it is essential that the DGs introduce their hosting requests as early as possible in the application life cycle.

To introduce a hosting request an electronic “THM form” (Transfer – Hosting – Modification) should be filled in using a technical template. Also, one additional template per domain should be provided.(ex: WebLogic, ColdFusion, webi, ORACLE).If help is needed when filling in a THM form, DG are invited to contact their corresponding Data Centre account manager for assistance either by phone or via the functional mailbox.

A technical project leader (or Change Coordinator in ITIL terminology) is also appointed within the Data Centre for coordinating the implementation of every request for change.

3.3.2. Change Management Process

Every hosting request creates a change in the Data Centre infrastructure and that change has to be managed in a controlled way. To achieve this, a change management process has been designed in alignment with ITIL recommendations including a Change Advisory Board with members of the Data Centre management team.

In this process, several checkpoints are foreseen to ensure that a request for change is handled systematically.

Concretely, 2 main checkpoints are defined by the CAB:

1. CAB approval for starting change implementation:

When a request for change is received, it is essential that the request from the client is clear and complete so that the impact on Data Centre infrastructure can be assessed accurately and the Data Centre can make a reliable capacity plan (machine workload, disk usage, staffing). Also, it is essential that the demand is aligned with the choices regarding architecture, standard products in the product catalogue and security policies. It is Data Centre policy that, in order to guarantee service levels, all components of Information Systems must be hosted in the Data Centre.

CAB approval will be given for implementing the request based on these criteria. Appropriate planning is required for implementing end-to-end monitoring and performing functional, performance (load/stress tests) and failover tests so that the implementation will be done “right first time”.

2. CAB approval for putting an application into production:

CAB approval will be given as soon as the tests mentioned above are performed successfully, the monitoring is in place and the configuration management database is updated.

To ensure that a systematic approach is adopted for signing off, a checklist will be used.

For more information regarding the change management process, contact the management team.

The changes initiated by DIGIT services (ex.: change of IP address) also follow the same phases. Only the procedures within a phase are adapted depending on the type of change (simple/multi domain, urgent...).

3.4. Incident Management for ISHS

What is Incident Management?

Information Systems Hosting Service incident management covers three areas:

- Incidents;
- Service requests;
- Complaints.

The following definitions apply:

An **incident** is any event which is not part of the standard operation of a service and which causes an adverse impact on the quality of service or the security of an Information System and/or DC components.

A **service request** is a request from a user for support, delivery, information, advice or documentation, not involving a failure in the IT infrastructure and not having an impact on the configuration items.

A **complaint** is a communication from the user about the quality of the service delivered by the ISHS.

There is one additional important term in incident management which is “**Configuration Item**”.

A **CI** is the item to which the incident refers. e.g. it can be an information system, a component of the IS like Oracle or similar. CIs are predefined in the Service Management Tool and incidents are linked to them when they are opened.

In SMT, there are two types of CIs:

- Initial CI: is entered by the originator of the incident;
- Failing CI: is entered by the Data Centre when the impacted CI is identified.

Incidents are opened by users via their local help desk or via the central help desk. They are processed primarily by the ISH Front Office with support from the ISH Back Office (see below). ISH Account Management has no role in resolving incidents.

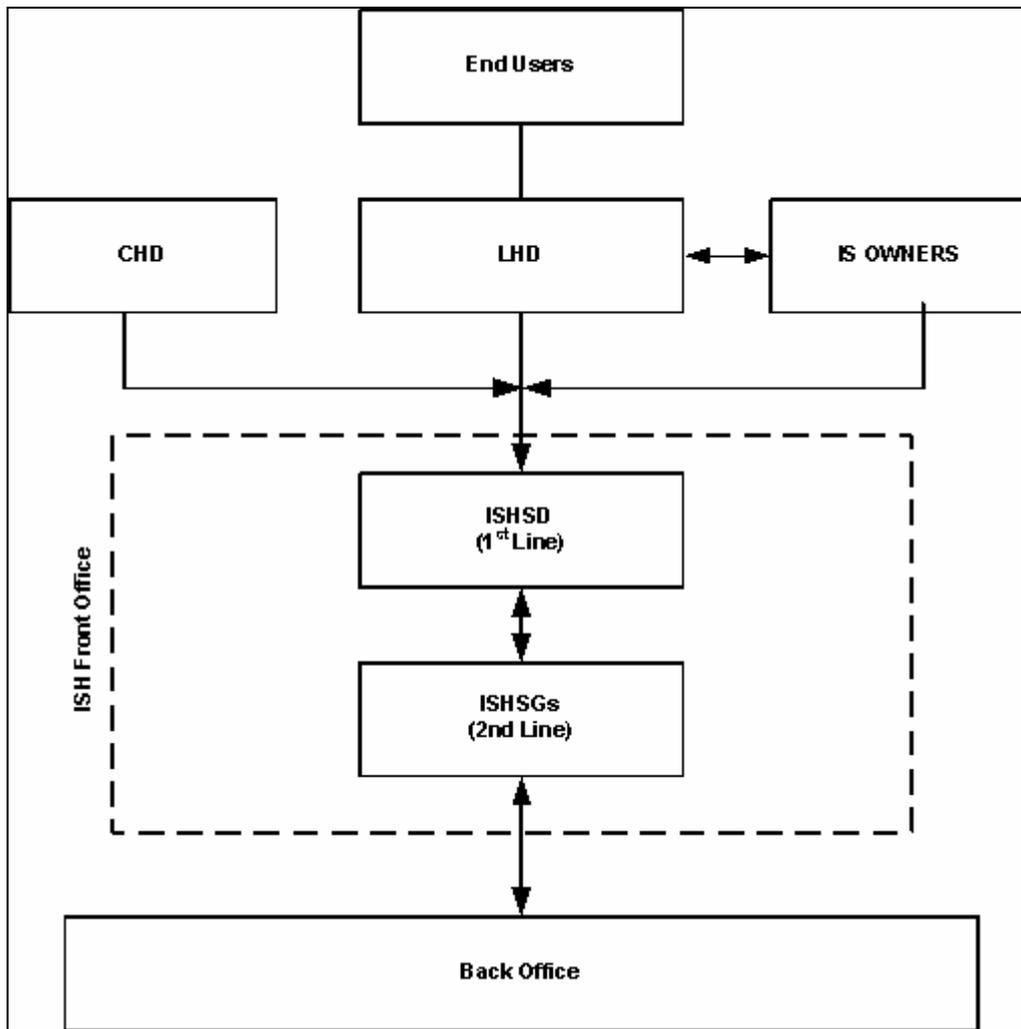


Figure 3 - Incident Management -

The organisation for handling incidents is:

- Front Office (ISHFO) – incident management
 - service desk (ISHSD) – 1st line
 - support groups (ISHSGs) – 2nd line
- Back Office (ISHBO) – planned activities and 3rd line support

The ISH Service Desk takes care of receiving incidents, assigning them to the ISH Support Groups and makes the necessary follow-up.

The ISHSGs solve the incidents while keeping contact, whenever required, with the ISH Back Office.

The detailed organisation of the ISH Front Office is shown above.

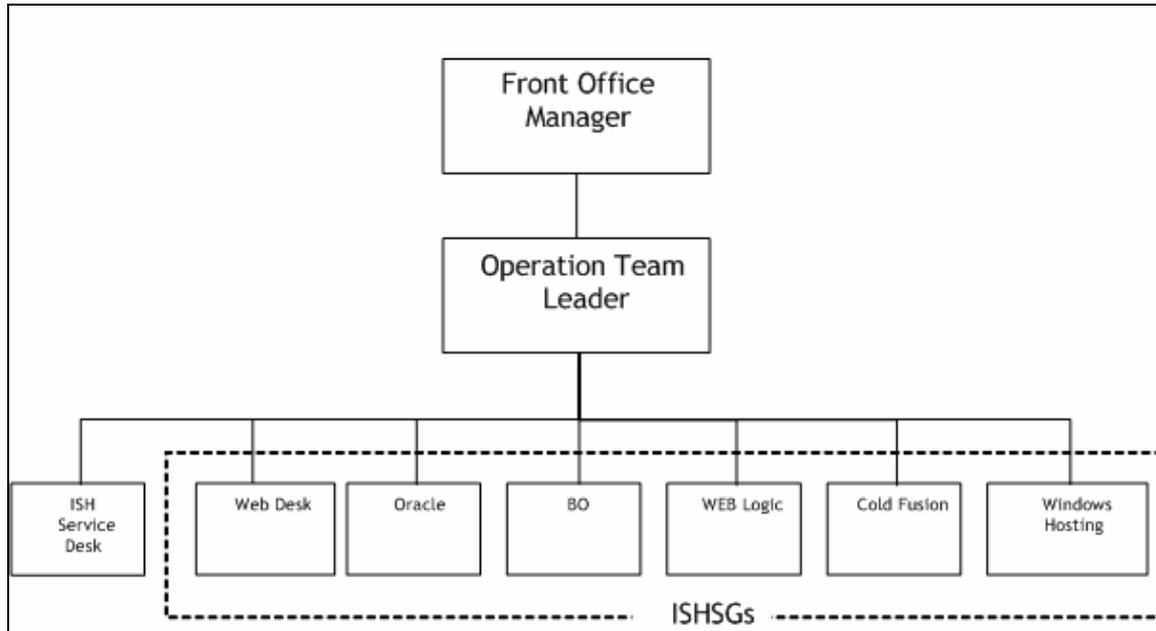


Figure 4 - Organisation of the ISHFO -

When is the ISH Service Desk available?

The ISHSD is available between 08:00AM and 07:00PM during normal working days.

Outside the period mentioned above, critical and urgent incidents regarding the Data Centre ISHS should be sent to the Central Helpdesk. Normal incidents should be reported on the next working day. When recorded incidents are categorised and prioritised.

Who can communicate incidents to the ISHS?

Only **reporting users** of the ISHS can communicate incidents, service requests and complaints to the ISHS.

The following are defined as reporting users of the ISHS:

- The hosted information system support group within the DG (information system owner), provided that the DG has access to SMT (see below);
- The DG Local Help Desk;
- The Central Help Desk.

It is important to note that the **end-users** of an information system **are not considered as reporting users** of the ISHS and, as such, they should contact the local help desk if they have an issue with an information system. It is the local help desk responsibility to follow the DG’s internal procedures for handling those issues.

HOW TO COMMUNICATE INCIDENTS TO ISHS

All the incidents for the ISHS must pass through the ISHS Service Desk (ISHSD).

The incidents can only be treated if communicated to the ISHSD through the EC standard incident management tool SMT, which means that each incident – **before it can be considered - must have a SMT assigned incident number.**

It also means that incidents arriving via current functional mailboxes or private mailboxes or via telephone will not be accepted.

There are two channels to communicate an incident to the ISHSD, depending upon the availability of SMT to the reporting DG. The following explains both cases.

DGs that have access to SMT

DGs that have access to SMT are able to create and assign incidents, service requests and complaints directly in SMT and to assign them to the ISH Service Desk.

In order to ensure a correct reception of the incident by the ISHSD, the subject line in the incident should be formatted as follows:

<application name> <operating system environment> <short meaningful description of the incident>.

If an application has several components, a clear identification of the failing component must be given.

For incidents with critical or urgent priority a special procedure must be followed. This procedure is described below.

DGs that do not have access to SMT

DGs that do not have access to SMT must pass through the Central Help Desk. They should communicate incidents, service requests and complaints to the CHD following the standard procedure. The CHD will then take care of entering the incident, service request or complaint into SMT and will assign it to the ISHSD.

For critical and urgent incidents, the CHD will apply the same procedure as described below.

Data quality

It is of vital importance that data registered for an incident is meaningful.

In particular, the following fields should be filled with meaningful information:

- Initial CI;
- Subject;
- Description.

The **subject** is used both in call dispatching/resolution and in statistics. In order to improve the resolution time and in order to produce meaningful statistics (the subject field is printed out as part of a report showing incidents by DG and by information system), the subject line should be compiled as follows:

<application name> <operating system environment> <product/technology> <short **meaningful** description of the incident>.

E.g. <application> <Solaris> <Oracle> <application does not respond>.

The **Description** field should contain a meaningful detailed description so that neither the ISHSD nor any support group has to call back the reporting user to understand the **nature** of the incident. It is very important, for an application having several components, to clearly identify the failing component.

When a critical or urgent incident is assigned to the ISHSD or to the CHD, the reporting user should follow up the assignment of the incident with a telephone call to, respectively, the ISHSD or the CHD to ensure that the right attention has been given to the incident.

HOW INCIDENTS ARE FOLLOWED UP

The ISHSD will follow up the progress of incidents and – whenever required – it will keep the reporting user updated. This is particularly true if – for any reason – an incident takes longer than expected to be resolved.

During its working hours (from Monday to Friday 8:00AM to 7:00PM), the ISHSD will follow up on incidents depending upon priorities as follows:

- **Critical:** the ISHSD will follow them up immediately with ISHS Support Groups. The targeted resolution time is of two hours after reception of the incident (not committed to in case of serious hardware failures or incidents that need in-depth analysis of the cause);
- **Urgent:** the ISHSD will follow them up within one hour with ISHS Support Groups. The targeted resolution time is of four hours (not committed to in case of serious hardware failures or incidents that need in-depth analysis of the cause);
- **Normal:** the ISHSD will follow them up with ISHS Support Groups within 4 working hours. The targeted resolution time is of eight working hours (not committed to in case of serious hardware failures or incidents that need in-depth analysis of the cause);
- **Low:** no follow up is done by the ISHSD;
- **Scheduled:** no follow up is done by the ISHSD.

The follow up information will be sent to the reporting user by e-mail.

Outside the working hours of the ISHSD, the Central Helpdesk should be called by the DGs instead of the ISHSD. The CHD is open 24/24 7/7.

At the same time – after the incident has been assigned by the ISHSD to a support group of the ISHS – the support group itself will have a direct contact with the reporting user if additional information is required.

Finally, the reporting user (please be aware of the definition of user given above) can contact the ISHSD by e-mail if incident status is required. In that case, the incident number must be specified.

Due to the potential additional heavy load upon the ISHSD, it is recommended to send those requests only in critical situations.

How to escalate an incident ?

If the DG has a perception that an incident is not correctly or timely managed by the ISHSD, the following escalation sequence may be activated:

- ISHS Front Office Team Leader:
- ISHS Front Office Manager:
- ISHS Service Manager.

3.5. Product Management

DIGIT maintains a product catalogue kept up-to-date with all the approved products that are supported. The product catalogue is updated regularly depending on the framework contracts signed by the European Commission, and product life cycles.

Products not in the product catalogue are not supported by the Data Centre.

Moreover, there can be a delay from when a product appears in “class B” or “class C” in the product catalogue and when it is actually available and supported in the Data Centre. This is dependent on the training of the technical people, the definition of the architecture and setting up of the new product within the Data Centre infrastructure.

In the case of “class B” products, competence centres are created in the Data Centre to support them.

In the case of “class C” products, the Data Centre hosts applications with these products but “ad hoc support” will be provided in agreement and close collaboration with the requesting client.

3.6. Data Protection

Any personal data included in the Contract will be processed in accordance with the requirements of Regulation (EC) 45/2001 on the protection of individuals with regard to the processing of personal data by the Community institutions and bodies and on the free movements of such data.

The data will only be processed for the purposes of the performance, management and follow up of the Contract by the Contracting authority (ies) without prejudice to a possible transmission to the bodies in charge of a monitoring or inspection task in conformity with Community law.

The Contractor may, upon request, obtain the communication of his personal data and rectify any inaccurate or incomplete personal data. Should the Contractor have any queries concerning the processing of his personal data, he shall address them to the Contracting authority (ies). As regards the processing of his personal data, the Contractor has a right of recourse at any time to the European Data Protection Supervisor.

The Contractor shall comply with Regulation (EC) 45/2001, notably article 23 thereof, and Council regulation (Euratom, EEC) N° 1588/90 of 11 June 1990 on the transmission of data subject to statistical confidentiality to the Statistical Office of the European Communities (OJ No L151, 15.6.1990, p. 1)".

As in general Regulation 45/2001 is applicable to all contracts, it is important to understand in a more detailed consideration, that there are 2 main issues a legal data protection clause has to cover in contracts:

- The processing of personal data by the institution in the contract documents, including any attachments;
- The processing of personal data under the responsibility of the institution during the execution of the contract;
- If in this context the sub-contractor is a Processor, the reference to Article 23 on "Processing of personal data on behalf of Controllers", pointing itself to articles 21 on "Confidentiality" and 22 on "Security", is important;
- If in this context the sub-contractor also has to act as a Controller subject to national data protection law the reference to Articles 7, 8 or 9 on "Transfer of personal data", is important.

The statutory staff and contractors working in the Data Centre have been duly informed on legal data protection provisions applicable to the Data Centre.

Contractors must respect the Data Protection Regulation by including a legal clause in their contract.

3.7. Security

Several levels of security measures are in place.

3.7.1. Physical access control

Security guards manage machine room access, access is only authorized by the Head of Unit.

3.7.2. UNIX production systems

In order to check the integrity of the UNIX environments, an automated routine is run continuously to verify each platform's UNIX configuration files and overwrites the files if a change is detected.

3.7.3. Secure environment

System administrators keep ROOT passwords; application owners have no access to the system with this role.

3.7.4. Password policy

Users should be prompted by information systems to use strong password in their logons.

Strong passwords have the following characteristics:

- Minimum length 8 alphanumeric characters;
- Mixture of upper/lower case and special characters;
- Maximum password age set (< 90 days).

3.7.5. Information system classification

The classification is based on Commission decision C (95)1510 on protection of information systems. Under the current regulations, none of the data contained in the information systems hosted in the Data Centre is classified "RESTREINT-UE" or higher. All systems process 'unclassified data' that do not require higher security measures than those currently implemented in the Data Centre.

Requests to host applications processing data classified "RESTREINT-UE" or "CONFIDENTIEL-UE" or that require non standard measures must be addressed in the first instance to the Head of Unit of the Data Centre.

3.8. Service Level Agreement

The Data Centre is currently working on the preparation of a service catalogue to structure its offer as well as on the implementation of monitoring to measure indicators (KPIs) and check them against committed service levels. The measurement of these KPIs will be the basis for defining service improvement actions.

Data Centre policy is to define with the client the service level that is required depending on the criticality of the information systems to be hosted, specifically regarding availability, monitoring, fail over, contingency and disaster recovery. For instance, it may not be necessary to foresee systematically 24/7 availability with contingency for each application. Through the negotiation process with ISHS account managers, SLAs will be defined and related conditions to be fulfilled by both parties (Data Centre and DGs) will be agreed so that service levels can be committed to by the Data Centre. For example, the Data Centre will give commitments regarding the availability of applications if changes are implemented on the production environments under its full control.

3.9. Preventive maintenance

3.9.1. No Break and Air conditioning maintenance

The Data Centre provides 24 hours/day, 7 days/week, the technical infrastructure (servers, disc space, software, tools for back up, monitoring and technical support) necessary for hosting the Information Systems of DGs.

This infrastructure needs permanent air conditioning and a no-break power supply on both sites. To guarantee operations the Data Centre organises annually several preventive maintenance interventions on the air conditioning and on the power supply.

Customers are informed in advance of the planning for these interventions and the possible risks for their applications.

Regarding the email to ensure the same availability for the contingency site and for the different sites in Brussels, one annual preventive maintenance operation per site is organised.

All maintenance dates are fixed each year in December for the following year.

An information note is sent in advance to all the Customers to explain the impact of the intervention.

3.9.2. Maintenance interruptions for Data Centre environments

In addition to these planned maintenance interventions required for testing contingency plans, installing patches and cleaning file systems, the Data Centre reserves the right to apply emergency interventions and system maintenance with minimum notice of :

- 5 (five) working days for planned interventions;
- As soon as possible during the day if an emergency intervention must take place after 07:00 PM;
- Without prior notice if necessary.

3.10. Configuration data

It is essential to keep systematically the information “up-to-date” whenever a configuration item is modified, added or removed within the Data Centre Infrastructure.

It is foreseen in the change management process to systematically update the configuration management database when a new change is implemented.

The benefit is that the impact of an incident or a change can be assessed within the Data Centre due to the links specified between the configuration items. Also the related documentation should be kept up-to-date.

3.11. Acceptance testing

Before authorizing a new application to be put into production, the Data Centre has to be sure that this new Information System will support the expected workload with acceptable performance and will not degrade the performance of other applications running on the same shared environment.

The acceptance tests required by the Data Centre are load tests, performance tests and failover tests if necessary. The load tests are executed by simulating a number of concurrent users. Additional performance tests are executed so that the resource consumption per process is measured.

A simplified procedure for testing corrective patches on existing applications will be defined by the Change Manager.

3.12. Reception Procedure - Production check list for new IS and major releases

The reception procedure consists of a series of items in the form of a production checklist. Each item in the list should be well documented as part of the implementation of new services and applications. The CAB will use this production check list to decide when a new information system is ready to be put into production (see chapter 3.3 Hosting Request and the Change Management Process).

The objective of the production checklist is to assure that every necessary element in relation to the administration of the hosting environment is well documented and available for all services before production is authorised.

It is intended that the input for several of these items is supplied by the client with the initial request for hosting (i.e. part of the “Mirella” form for new hosting requests).

Item 1. Architecture and description

The service/application architecture should be documented with a description of its purpose, clients and a simple flowchart illustrating the interaction between the various components.

Information concerning co-existence requirements (e.g. dedicated environment, co-existence only possible under certain conditions, etc.) is specified as an important part of this item.

Item 2. Predicted usage patterns

The predicted usage pattern (e.g. peak at the end of every quarter) is documented to allow appropriate planning of infrastructure interventions and evolution. The input will serve as a supplement to data obtained through the subsequent production monitoring.

Item 3. Installation procedures and routines

All installation procedures (from OS to application) are thoroughly documented and tested. Prerequisites for the installation and possible impact on other services (e.g. required server reboots) are included as part of the documentation.

Item 4. Software availability, support and licenses

The license coverage for all extra-OS software products are exhaustively documented, including a description of support channels (directly with supplier, etc.) and classification within the product management framework.

Item 5. Backup / Restore

Basic and application specific backup and restore routines are exhaustively tested and documented.

Item 6. Monitoring / Alerts

Exact specification of monitoring needs and methods are defined and documented in collaboration between the Data Centre and the client.

End-to-end monitoring is implemented for all application components.

Item 7. Reporting

With a view to establishing service specific SLAs (or in-line with existing SLAs) reporting mechanisms (e.g. response times, number of users, storage consumption) should be defined, documented and implemented in collaboration with the client.

Item 8. Responsibilities (including access rights)

With a view to establishing service specific SLAs (or in-line with existing SLAs) allocation of responsibilities between the Data Centre and the client is documented.

Administration of access rights to the service/application is thoroughly documented, e.g. who defines users, grants access to data, etc.

Item 9. Security

In addition to the applied base OS security and anti-virus, all application/service specific security requirements are analysed, documented and implemented, e.g. behaviour in relation to personal data, strengthened system lockdown, restricted administrator access, auditing requirements, etc.

Item 10. Communication

Based on the elaboration of a list of events requiring communication to/with the Client (e.g. service down, response times above preset threshold), the communication mechanism is documented (e.g. mail to Exchange distribution list/functional folder).

Item 11. Availability and contingency

Each new service or application is accompanied with a description of availability and contingency measures, if required:

Failover mechanisms are documented, including a detailed description of associated procedures and responsibilities.

A service specific contingency plan is elaborated.

Item 12. Operational procedures / contingency

Any service/application specific procedure is elaborated and documented, e.g.

- Clean up of temporary files;
- Procedure for move from test to production;
- Procedure (internal within the Data Centre, also between the Data Centre and the client) for activating the contingency if necessary.

Item13: Technical documentation

This should document the application as it has been implemented in the Data Centre (change of IP addresses of the servers, the flows...)

3.13. Publishing guidelines on Europa and IntraComm

For publishing data on the Commission's web sites, a set of rules has to be followed for the organization of the data.

These guidelines cover:

- Introduction;
- Data structure for web pages;
- Using ColdFusion;
- Weblogic Infrastructure;
- Uploading data;
- Updating the dissemination sites;

- Staging manager;
- CGI scripts and programs;
- Using the search engine;
- Publishing data bases on the Web;
- Working with the reverse proxy servers;
- Configuration related issues;
- Running "local" web servers on the IntraComm platform;
- Recommendations.

3.14. Monitoring

3.14.1. *Service Description*

The monitoring service is responsible for monitoring the Commission's central computing facilities in order to assure its availability on a 24/24 hours and 7/7 days basis. In this mission-critical context, the monitoring software verifies constantly the performance and availability of the servers hosting applications and information systems, as well as Europa, the European Union's website on the Internet.

This software can generate warnings to technical staff based on a very wide range of parameters. This allows anticipation and avoidance of most problems potentially affecting the servers' performance, thus preventing server crashes and the unavailability of applications and information systems.

The monitoring service is composed of 3 layers:

- **Information collection:** specific checks are performed at regular intervals, either locally, e.g. to verify the file system usage, or remotely, e.g. to verify the availability and response time.
- **Alert management:** for each parameter verified, it is possible to define thresholds that will trigger an e-mail or SMS message to the people concerned. This mechanism allows potential problems to be anticipated and to correct them before services will be impacted.
- **Reporting on collected information:** the objective is to give the Data Centre and the DGs a regular follow-up on the quality of service offered. It centralises all relevant parameters and provides consolidated reports. The analysis of the reports allows the verification and the follow-up of the quality indicators based on objective measurements. It is also possible to verify whether the level of service defined in the "Service Level Agreement" is being met provided that the corresponding indicators are well defined and measurable.

3.14.2. *What can be monitored?*

Availability and performance of the operating system and standard basic software.

Regular checks are performed at two levels:

- Various parameters of the operating system: CPU usage, file system usage, process availability, etc.
- Standard and basic software, this includes Oracle, web servers, etc.

These checks are mainly performed for internal use.

Verification of the availability and performance of information systems.

Verification of the availability of information systems is done through end-to-end checks to test whether the different components of the information system are working as expected.

The implementation of this kind of monitoring depends strongly on the application logic and usually implies a specific development that can be provided by the Data Centre based on monitoring specifications supplied and maintained by the DGs.

The objective is to check automatically the behaviour of an application and to be able to easily detect the failing component(s) so that appropriate action can be taken, if necessary.

Verification of the Quality of Service

Besides verifying the availability and performance of information systems, it is also important to have indicators on the quality of service offered. One of the key elements for verifying the quality of service consists in measuring the end-to-end response time of an application giving objective figures about the performance end-users are experiencing.

This is achieved by installing a monitoring probe in a DG's premises. The DG is responsible for supplying a PC and defining the generic transactions, which will be executed at regular intervals. The transactions simulate an end-user connection and should be as representative as possible. The Data Centre will install the necessary monitoring software. This procedure can be used in the framework of acceptance testing and it is systematically used to measure response times when the Data Centre signs a SLA with its clients.

Concerning information systems, which are accessed via Europa, it is possible to measure the availability and performance from the outside world. In other words, it is possible to measure the real end-to-end experience of someone on the World Wide Web who connects to services on Europa. This service is only implemented at the request of the system owner included in the initial hosting request presented to the CAB for decision (see 3.3 Hosting Request and the Change Management Process).

3.15. User Statistics – Data Centre Resources Usage

The resource consumption of the Data Centre's servers will be available by host machine, DG, information system and userid on a monthly basis and is currently under construction.

These statistics may be used for capacity planning.

3.16. User statistics – Web Consultation

3.16.1. Statistics on Europa and IntraComm

A number of different on line statistics (total number of hits, total number of visits, per DG, etc...) are at the disposal of users.

- The statistics related to web site consultation on Europa have been developed by DG PRESS.
- The statistics related to web site consultation on IntraComm have been developed by DG ADMIN

3.16.2. Statistics on other sites

A limited number of standard statistics on other web sites hosted in the Data Centre can be obtained by introducing a service request via the Central Helpdesk or by creating directly a ticket in the Service Management Tool.

4. INFRASTRUCTURE

4.1. Introduction

This chapter describes the infrastructure that the Data Centre currently provides. The intended audience are primarily the development teams.

After analysing the requirements based on the application architecture provided by the client showing the components and access paths, the Data Centre determines the servers and storage required to satisfy the expected performance and availability levels using the infrastructure described below.

4.2. Servers

Based on the contracts signed between the European Commission and its Suppliers, the Data Centre offers a range of UNIX and Windows/INTEL servers.

4.2.1. UNIX

The current range of UNIX servers that can be used for hosting information systems is:

Configuration A (entry level)

2 CPUs

6 GB

No local storage, only NAS storage.

Configuration B (mid range)

4 CPUs

16 GB RAM

No local storage, NAS or SAN storage.

Configuration C (high end)

8 CPUs

32 GB RAM

No local storage, NAS or SAN storage.

Configurations A, B and C are essentially used for those components, which can be scaled horizontally as is typically the case for servers used for the access, presentation and application layers.

Configuration D (highly scalable)

This server is a machine with a set of domains. Each domain has 1 or more basic building blocks of 4 CPUs and 32 GB RAM

Maximum scalability is 18 building blocks (72 CPUs, 576 GB RAM)

Configuration D is primarily used for those components that have to be scaled vertically as it is typically the case for the database layer.

Note: These configurations are given for illustration purposes and are not necessarily up-to-date as they are continuously evolving.

4.2.2. Windows/Intel

The current range of INTEL servers that can be used for hosting information systems is:

Configuration A (entry level)

From 1 x 2.4 Ghz to 2 x 3.0 Ghz processors

From 512 Mb to 6 Gb memory From 2 x 36Gb to 5 x 146 Gb Disks

Configuration B (mid range)

From 4 x 2.0 Ghz to 4 x 2.8 Ghz processors

From 1Gb to 16 Gb memory

From (8 x 73 Gb + 2 x 36 Gb) to 10 x 146 Gb Disks

Configuration C (high end)

8 x PIII Xeon 900MHz processors

From 4 Gb to 32 Gb memory

2 x 36 Gb disks

These 3 configurations are essentially used for horizontal scaling of components.

4.3. Time Servers

All Data Centre Servers clocks are synchronised with a time reference, the reference signal used in the Data Centre is the broadcast radio signal sent by the German atomic clock in Frankfurt.

In the Data Centre, there are 4 master time servers in the buildings of the Commission. These servers are the reference for the other 8 Time Servers accessible by clients.

4.4. Domain Name Servers

The Domain Name Servers provide the clients with a simple and stable way of addressing machines with a virtualized access to the servers. Only machine names have to be addressed, not the corresponding IP addresses.

The use of DNS has the advantage that the changes made within the Data Centre are transparent for the applications addressing machine names instead of fixed IP addresses. Also, it is a mechanism for implementing failover between machines (by mapping the DNS name to an alternative IP address).

In the Data Centre, the DNS are highly redundant (8 servers spread over different buildings) and have a high availability.

4.5. ORACLE Name Servers

The ORACLE Name Servers provide clients with a simple and stable way of addressing databases by providing a virtualized access to the servers whereby only ORACLE DB names have to be addressed, not the corresponding physical addresses.

The use of ONS has the advantage that the changes made within the Data Centre are transparent for the applications addressing ORACLE DB names instead of fixed links and addresses. If a database is moved from one server to another by the Data Centre, only the mapping in the ORACLE Name Server has to be updated and the change is therefore transparent for the users. ONS also provide a mechanism for implementing failover between DB servers.

In the Data Centre, the ONS are highly redundant and have a high availability.

4.6. DMZ

The Demilitarized Zone is a network subset isolated from the Commission network by firewalls. It is used to manage access to the Commissions network and services from the external world (Internet, TESTA, CCN...).

The contingency for the DMZ is implemented between Luxembourg and Bruxelles.

The current DMZ should be used for gateways between 2 domains and not for hosting applications.

4.7. Storage

4.7.1. Service description

The Data Centre offers storage capacity associated with backup and disaster recovery depending on the type of environment (production, non production)

4.7.2. Technology used:

Two storage architectures are currently implemented at the Data Centre:

- Network Attached Storage;
- Storage Area Network.

NAS

A NAS device is a combination of a server (filer) and a large amount of RAID storage. NAS devices are attached to the existing local area network and access to the file systems is made through file-level commands.

Two types of network file systems are used:

- Common Interface File System in the Windows world;
- Network File System in the UNIX world.

The current infrastructure is based on a NAS production cluster in one building, with disaster recovery and backup components in the other.

The main advantages of NAS are:

- Easy file sharing;
- File system localised in one device: simple set-up and maintenance.

Drawbacks are:

- Performance problems for large database applications;
- TCP/IP processing needs a lot of CPU;

- Asynchronous mirroring between 2 sites with a latency time of 5 minutes meaning that 5 minutes of transactions are lost in case of failover.

NAS is used for:

- Common Storage Service offering storage capacity to DGs for their needs in the area of office automation.
- Specific internal Data Centre needs:
- Static pages for web servers;
- Files related to COLDFUSION and WEBLOGIC;
- UNIX home for users to share among several servers;
- File systems which must be used on several servers, but which can be shared like standard system software (top, gcc, sas, ...), ensuring that the same copy is used everywhere.

SAN

A SAN installation is composed of a high-speed fibre channel network to connect a large storage subsystem to multiple servers. The file systems are located on the servers that run the application and in order to avoid two file systems from different servers unwittingly overwriting each other's data, each storage subsystem is partitioned so that a range of hard disk drives can be logically assigned to a specific server.

Currently two RAID levels are used: RAID-1 on EMC SAN, RAID-5 level on HDS SAN. RAID-1 is more expensive than RAID-5 and has equivalent read performance. However, its write performance is considerably better. RAID-5 gives acceptable performance for most applications, except for sustained load with write operations.

The main advantages of SAN are:

- Reliability because of the redundancies built into the storage subsystem;
- Scalability because of a common pool of spare disks;
- Performance and better disaster recovery;
- Synchronous mirroring between the 2 sites.

Drawbacks:

- High implementation costs;
- Lack of SAN standards;
- Interoperability in case of different vendors.

SAN is mainly used for:

- ORACLE and SQL data bases;
- Traditional file systems which are heavily used and which need high performance.

The choice of the technology used is made by the Data Centre based on the needs of the application. Thus it is important that the Data Centre receives all the input required to make the appropriate decision for meeting best the user expectations.

4.7.3. How to obtain the service

Storage should be requested by completing the appropriate form in MIRELLA for storage associated with new applications or additional storage for existing applications.

It is very important that the Data Centre receives realistic forecasts from DGs on storage usage for the forthcoming 2 years so that it can anticipate needs.

DGs should be pro-active by introducing requests early enough so that the Data Centre has the time to implement the required storage. Typically a request for storage capacities exceeding 1 terabyte should be received by DC 4 months before the required availability.

4.8. Contingency

The Data Centre may offer contingency with manual or automatic failover for production environments.

This service is implemented only if it is explicitly required by the client and clearly indicated in the application schema presented for CAB decision at the time when the hosting request is introduced.

Moreover, it is the responsibility of the information system owner to define setup and test the procedures to be applied in case of a disaster. The Data Centre will support the client in preparing these procedures.

Fail-over and contingency implementations generate considerable extra-costs and the client should examine whether those implementations are absolutely mandatory based on the criticality of their information systems.

If fail-over or contingency is required, then particular guidelines should be followed by the client when developing an application. Additional fail over tests must be undertaken before the application is put into production. For more details, please refer to the guidelines given for WEBLOGIC, COLDFUSION and ORACLE in Chapter 5. In particular, an application should never address directly server IP addresses or hostnames but rather via DNS with a dedicated name in the application.

Licences linked to servers should never be used.

4.9. Load Balancing

Load balancing may be required to ensure no downtime and/or for providing increased workload scalability.

Load balancing is not available at application level for the hosting platforms with the current technologies as actually implemented.

However, Load Balancing is available at telecommunication level and may be implemented.

If necessary, a request for load balancing will be subject to a feasibility study with regard to the offered service.

4.10. Backup Service for the local servers in the DGs

4.10.1. Service description

A backup/restore service is offered primarily for the data of the information systems that are hosted within the Data Centre. The main objective of this service is to restore the system, NOT to do "data archiving". The service is also mentioned specifically in the technology chapter below.

The Data Centre can offer a backup service for the DGs:

- With a standard retention period;
- With a backup cycle including a full backup and all incremental backups associated until the next full backup;
- For data from databases managed by the platforms given below (see § Technology used) and the data from the file systems.
- Starting at planned times agreed between the Data Centre and the DGs;
- Including the management of the cartridges stored within the Data Centre premises.

With an average backup performance of 11 GBytes/Hour and an average restore performance of 6 GBytes/Hour performance can fluctuate significantly depending on the number of the files and their size. Restore is the responsibility of the local system administrators in the DGs who execute restore for local servers.

4.10.2. Roles and Responsibilities between the DG and the Data Centre

This following table summaries the responsibilities for backup and restore operations:

Activity	DG	Data Centre
Electronic Request	Make request for backup via Mirella form	Analysis and CAB decision. Planning for implementation is agreed between the Data Centre and the DG. Contacts the user
Monitoring	The DGs get daily mails sent automatically by the backup server about the completion of the backup as soon as completed. If incremental backups have to be restarted, the DGs must ask for this explicitly to the Data Centre.	Monitors the completion of the backup and restarts the failed full backups. This is not systematically the case for incremental backups. Informs DGs about client software updates
Restore	DGs. are responsible for executing the restore. DGs are prompted to inform the Data Centre per e-mail about restores that will be launched BEFORE starting so that the Data Centre is aware and does the monitoring.	Monitor restore operation (tape activity)

The interlocutors of the Data Centre are normally local system administrators who will be in contact with the backup team of the Data Centre for installation, control of backup, change request, updates etc.

Every day, a notification of the backup status is sent to the DGs (at a pre defined address given by the DG at request time) so that DGs can insure daily follow-up of their activity.

4.10.3. Technology used

The software functions in "Client/Server" mode. Given the operating mode, a part of software has to be implemented on each client who wishes to use the backup service. The client part has always to be installed because it constitutes the communication base between the backup server and the client itself.

4.11. Environments

Environments are created and maintained for each information system based on its specific requirements.

The Data Centre offers typically the following standard environments:

- Development/test environments used for development, technical and acceptance tests;
- Production environments.

Based on requirements, additional environments may be offered (For example: training, performance, and conformance)

For validating the sizing before going into production, load/stress tests are executed on a separate environment within the Data Centre (not accessible for clients) with:

- A workload simulation tool to simulate user sessions;
- Other Data Centre internal performance measurement tools;
- A subset of production data;
- A copy of a stable application release.

5. TECHNOLOGICAL ENVIRONMENTS

5.1. Introduction

This chapter provides guidelines to be followed when developing new applications so that the integration of these applications within the Data Centre infrastructure will be optimally implemented. The intended audience are primarily the Customer's application development teams.

The main software technologies used for applications are described. All requests for these services should be introduced in MIRELLA using the appropriate forms and will be processed by the CAB. Account Managers may be contacted for further information.

5.2. Access Layer

5.2.1. Web access to Information Systems

5.2.1.1 Public access - Europa (http)

Europa is the web site dedicated to the Internet presence of the European Commission, the infrastructure is managed by the Data Centre.

5.2.1.2 Restricted Access

IntraComm is accessible only for the personnel employed by the European Institutions.

Webgate (HTTPS), the goal is to provide access to restricted sites in an encrypted mode between the webgate DMZ and the client.

5.2.2. Security conventions

5.2.2.1 Access from external companies

A security convention must be signed between the information system owner, the external company and the security directorate. The information on IP Addresses and ports is provided by the Data Centre to the DG within the framework of the change management process.

5.2.2.2 Access from other European institutions or Agencies

An ad hoc security convention has to be established between security directorate, the Data Centre, the information system owner and the remote institution.

5.2.2.3 Miscellaneous

A direct access, outside the EC network from a server located into the EC network.

5.3. Web Servers

The web platforms currently used are Iplanet V4.1 and Apache V1.3.x.

Aligned with DIGIT's open source policy, the Data Centre will progressively phase out Iplanet and replace it by Apache.

Consequently, **no** new web application dependent on Iplanet specific technologies will be accepted by the Data Centre.

5.4. Coldfusion

ColdFusion is a server-scripting environment for creating Internet applications. ColdFusion has had a successful history in the European Commission since its introduction in 1998.

The Data Centre provides access to this technology and makes available three types of environment: development, test and production.

By developing applications at the Data Centre, they will be compatible with the production environment.

5.4.1. Service description

The Data Centre offers a ColdFusion hosting environment for development, test and production use. The production hosting environment consists of three categories of dissemination site: Europa site, Intracomm site and local web site.

At present, two full versions of ColdFusion are available in the Data Centre: ColdFusion 5 and ColdFusion MX 6.1

Recommended version: ColdFusion MX 6.1

No new application using ColdFusion 5 will be accepted.

New applications should be developed using ColdFusion MX 6.1. ColdFusion MX 6.1 provides greater performance, new features and is stable.

ColdFusion 5 and MX 6.0 applications should be migrated to ColdFusion MX 6.1.

The Code Compatibility Analyzer must be used to find obsolete and not recommended features.

Architecture

Please refer to the schema representing the architecture of the Data Centre.

Failover is being implemented.

CFMX 6.1

ColdFusion MX 6.1 is installed in J2EE mode. This feature allows multiple server instances on a single machine to be defined, each running ColdFusion MX 6.1.

Application isolation is the biggest advantage of running multiple instances of ColdFusion MX.

Multiple ColdFusion MX 6.1 instances will be deployed on top of the JRun 4 application server.

The web server is currently Apache 1.3.28.

Reverse proxy

All access to the ColdFusion public sites is transparently channeled through reverse proxy servers. The reverse proxy servers map incoming HTTP requests to the appropriate web and ColdFusion servers.

Classification of hosted applications

Applications are classified as follows

Public sites

- Europa “Gateway to the European Union” <http://europa.eu.int>

The content is managed by DG PRESS (see IPG guide <http://europa.eu.int/comm/ipg>)

Restricted sites

- IntraComm “The Intranet of the European Commission”
<http://www.cc.cec>

The content is managed by DG ADMIN.

- Local sites. The content of these intranets is managed by the DG;

- Secure sites. The content of this extranet is also managed by the DG.

For more details, please refer to the chapter on the Access Layer.

Environments

Development

The Data Centre offers a shared development environment to the DGs. This environment is identical to the production environment.

It is recommended to develop on the Data Centre's servers to be sure that the application will work correctly on the production environment.

Load test

Each version of an application must pass load tests. A dedicated environment exists to perform such load and stress tests.

Test/Acceptance

The acceptance environment should be used to validate a version of the application before it is put into production.

Training/Production

The Data Centre provides ColdFusion production servers. The target is to dedicate one CF/MX instance to each application.

5.4.2. Conditions and constraints related to ColdFusion

5.4.2.1. ColdFusion MX6.1

Administration

The ColdFusion front office team installs and manages the ColdFusion MX 6.1 servers.

To install an application at the Data Centre, the DG must clearly specify their needs: datasources, mapping, scheduled tasks, verity collection.

Datasources

Only Oracle datasources will be created.

The datasource name must be {dg}_{application}_{d/t/p} where:

- {dg} is the DG's acronym
- {application} is the application name
- {d/t/p} is for development or test or production

The Oracle datasources can be configured with 3 different types of drivers:

- DataDirect
- Oracle Thin
- Oracle OCI

Oracle Name Server should be used.

If the application needs to use LOBs, the datasource must be created with the DataDirect drivers included in CFMX 6.1.

Application	JDBC drivers
Without LOBs	DataDirect or Oracle thin or Oracle OCI
With LOBs	Only DataDirect (Bug with Oracle thin)
With ONS	Oracle OCI
With LOBs and with ONS	No solution

Other constraints are specified in the security part of this chapter (see below.).

Oracle Datasource	
CF datasource name	{dg}_{application}_{d/t/p}
Driver	DataDirect or Oracle Thin or Oracle OCI
Maintain Connections	Checked
CLOB	On request
BLOB	On request
Long Text Buffer (chr)	64000
Blob Buffer(bytes)	64000
Allowed SQL	By default all the sql commands are allowed. On request, the ColdFusion team can disallow: select, insert, update, delete, create, drop, alter, grant, revoke, stored procedures.

Verity collections

A Verity collection is a group of information that can be indexed and searched as a set. The ColdFusion team creates the Verity collection and the application can manage it by itself.

This search engine is not linked to the Verity Search'97 Information Server which is used as the general search engine on the dissemination sites. The use of this central search engine is to be preferred.

The ColdFusion instance of the Verity Search'K2 search engine can be used within the boundaries of applications, in cases where the central search engine cannot give the desired results.

Please note that the central search engine does not include cfm pages in its indexes. ColdFusion pages are generated on the fly, and their contents, in principle, will change every time they are requested. It is therefore useless to include this unstable information in the search indexes.

Web Services

Web services provide remote application functionalities over the net. With a web service, requests to the remote application to perform an action may be made.

Extensions

Java Applets

ColdFusion has several Java applets that may be accessed using the CFFORM tag.

The ColdFusion team can register specific applets which users can add to CFFORM forms using the CFAPPLET tag.

CFX tags

CFX tags are custom tags written against the ColdFusion Application Programming Interface (API) to extend and enhance CFML.

The ColdFusion team must register specific customized CFX tags prior to using them in ColdFusion application pages.

Custom tags

Custom tags are allowed. They are implemented in users' own directories under site root: the ColdFusion team adds the appropriate path in the ColdFusion server instance.

Security

Security issues include architecture, data services, file access and sandboxes.

Architecture security

The first security is on the architecture level.

Each production instance of ColdFusion MX 6.1 has its own environment which means:

- It runs with a unique UNIX user;
- It runs in its own JVM;
- It runs only one application.

So, each server instance runs an independent application and problems encountered by one application have no effect on other applications.

Datasource security

The datasource security depends on the type of application. By default all SQL operations are allowed. On request, the following actions can be disabled: select, insert, update, delete, create, drop, alter, grant, revoke, stored procedures

File access

The complete ColdFusion application (cfm and cfc files) are stored under the /ec/prod/app/webroot directory.

According to the type of application, the application will be installed in one of the following directories:

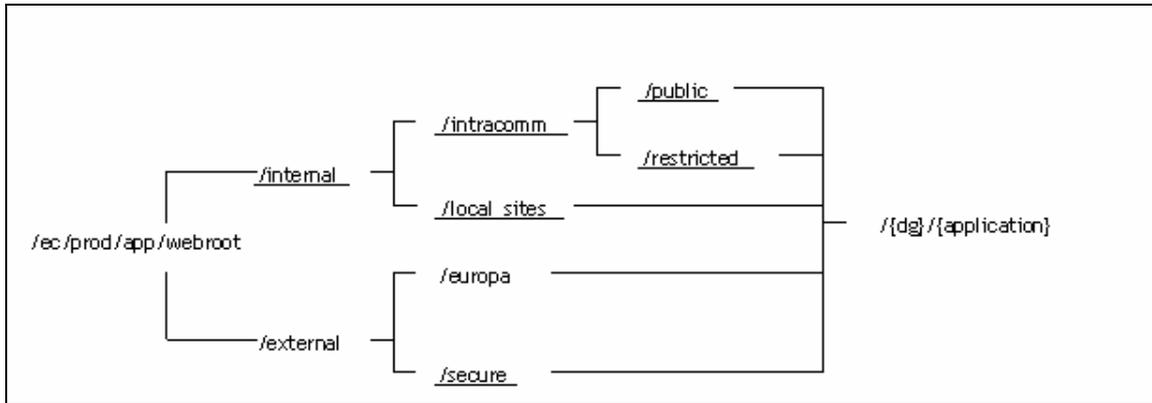


Figure 5 - ColdFusion Hierarchy -

A ColdFusion application cannot write under the `/ec/prod/app/webroot` directory.

But two directories are available to write temporary and permanent data.

Temporary data can be written in the work directory:

`/ec/app/prod/cf_app_doc/work/{dg}/{application}`

Permanent data can be written in the repository directory:

`/ec/app/prod/cf_app_doc/repository/{dg}/{application}`

Sandboxes

In ColdFusion MX Enterprise Edition, you can configure multiple security areas on a per-directory basis. These security areas are called sandboxes. A sandbox is a designated directory of a site to which security restrictions can be applied. Thus, sandbox security allows tags, functions, and resources (for example, files, directories, and data sources) to be specified which can be used by ColdFusion pages located in and beneath the designated directory.

- CF TAGS restrictions: see appendix 5
- CF FUNCTIONS restrictions:

All functions are allowed

- Files/Dirs restrictions:

On request, supplementary restrictions can be set on the directories or files.

Scheduled tasks

The Scheduling facility in ColdFusion MX Administrator allows the generation of static HTML pages to be scheduled. The scheduling facility is especially helpful for applications that do not require user interactions or customized output. ColdFusion developers often use this functionality to schedule daily sales reports, corporate directories, and statistical reports. Response time is fast because the output is an HTML page, not a database transaction.

ColdFusion MX allows pages to be scheduled for execution on a daily, weekly, or monthly basis. e.g. specify a time of day for execution; schedule a page to run only once, or on a specified date.

Debugging and logging

The debugging options are only enabled on the development and on the test server.

5.4.2.2. Coding Recommendations

Locking variables

Developers have to use the CFLOCK tag to ensure the integrity of shared data.

SQL commands must NOT be locked!

Please refer to the Macromedia website for more information at <http://www.macromedia.com/support/ColdFusion/ts/documents/tn18235.htm>.

Cookies

- CF permits the use of cookies. Nevertheless, for dissemination on EUROPA, the use of cookies is allowed only with certain restrictions:
- Cookies can only be used without explicit permission if they are limited to the current session;
- In the exceptional case where a cookie must be stored beyond the current session, explicit permission must be obtained, including an explanation of why it is necessary and the expiry period must not exceed one year. Furthermore the exact information which will be gathered must be listed and an assurance given that it will not be used for any purpose other than the one stated;
- If refused, the cookie must not simply try again indefinitely, nor must access to the site be refused.

XML

There are some basic XML editing capabilities available in ColdFusion but when using large files or more complex structures it's not a suitable product. An external parser is recommended for complex XML file generation and manipulation in a production environment.

Components

A ColdFusion Component (CFC) is an encapsulated ColdFusion application file that resides on the server. CFCs allow developers to make applications and functions available to a variety of clients including other ColdFusion developers, web browsers and web services. The CFCs can support multiple methods, making it easy to group similar functions into a single component.

5.4.2.3. Access to environments

Data

The application can be installed through FTP on the 3 environments development, test and production. The ColdFusion team configures the FTP access.

Administration console

As the ColdFusion team manages all the resources linked to applications users do not have access to the administration console

Log files

The application error file is automatically sent to the DGs

5.4.3. References

Useful links are:

ColdFusion MX 6.1 documentation:

<http://livedocs.macromedia.com/coldfusion/6.1/index.html>

ColdFusion technotes:

<http://www.macromedia.com/support/coldfusion/technotes.html>

5.5. WebLogic

5.5.1. Service Description

The Data Centre offers the following BEA WebLogic hosting environments on UNIX: development/test, production/training, and performance test. For these environments the Data Centre takes care of numerous aspects which will be developed in the following sections.

The Data Centre does the hosting and deployment of EAR, WAR and JAR archives.

The Data Centre does NOT create the archives, as it is the responsibility of the client to create them.

The Data Centre creates the domains from the “WebLogic form” information provided at hosting request time. The client will have only access to the development environment through a FTP server. The FTP server will allow J2EE applications to be deployed only on the development server (hot deploy). For the production, a “temp” directory is present on the development side so that the production is not impacted.

The Data Centre creates a domain per application. Generally, the Data Centre deploys all the applications on one server (the admin) to reduce the memory usage. Would split architecture be needed, the requirement should be sent via the Service Desk.

The environments are installed on a number of machines dedicated to WebLogic. However, WebLogic domains from several clients share the same machine: as a general rule, there is no machine dedicated to a single application.

WebLogic domain types are development / test, production / training and performance test. A machine is dedicated to one type of environment.

The choice of where a particular application is put depends on the information given in the request form.

The Data Centre is involved ideally in each phase of the development life cycle:

Development and Test phases

The Data Centre needs to be involved as soon as possible in the lifecycle of the applications to be hosted. This will allow the Data Centre to formalise and possibly, to adapt in an early stage the entire environment details, in order to guarantee the best integration with the production environment.

For this purpose, a development / test environment, similar to that of production can be made available.

To configure the environments, an installation guide is needed, including test cases to validate the installation.

Production phase

Before putting applications in production, the Data Centre needs to confirm that the performance tests have been executed successfully.

The Data Centre provides hosting for production applications in a failover configuration this means that if a machine goes down, the service can be restarted from another machine within the cluster.

In addition to the standard services related to operations, the Data Centre provides additional support and maintenance services to assist clients in case of errors.

Services to guarantee the operations of the hosted system include:

1. Monitoring

All production domains are monitored.

2. Tuning

The Data Centre uses a tool to collect information. This information is stored in a “CSV” file (available from the logs directory). One can easily interpret graphically these values with external tools (Excel for example).

Information is available concerning:

- Connection pool;
- Heap size;
- Thread activities;
- Servlet activities.

Additional information may also be collected if specified by the Customer.

5.5.2. Conditions and constraints related to WebLogic

The general architecture for web logic is that shown in chapter 2. Fail over mechanisms are currently operational at server level.

The Data Centre makes the latest version of the WebLogic server product available on all machines and this version is available to configure user domains. The same version of WebLogic is installed for all the domains associated with the same application (development and production).

The Data Centre manages all domains. All the configurations are made by the Data Centre following user requirements.

The FTP server allows applications to be deployed. These applications are in “EAR”, “WAR” or “JAR” format. The explode mode is not supported.

The applications should be deployed in the “applications” directory.

A subset of the directory is present inside a domain configuration:

- “Patches”: all the files present inside this directory should be “jar” files. This directory is used to patch the WebLogic API (mail API by example)
- “Lib”: all the files present inside this directory should be “jar” files. The other files are not loaded by the start script. All the elements of this directory are placed after the “weblogic.jar” in the classpath. This directory is used to install libraries for all your application as this one is present in the classpath of the server (server side not application side).
- “Classes”: this directory is used to store properties files used by application settings in relationship to the Data Centre environment. It is recommended to use this directory instead of storing the properties files in applications.

Default values are as follows:

Parameter	Value
JAVA_OPTIONS_AS	<i>-ms128m -mx256m</i> if other non standard values are requested they must be justified by stress test results.
LANG	<i>en_US</i>
START_MODE	DEV = false (auto deploy) PROD : true
Database Driver	Oracle Thin Driver For more informations, see this link

The Domain structure is as follows

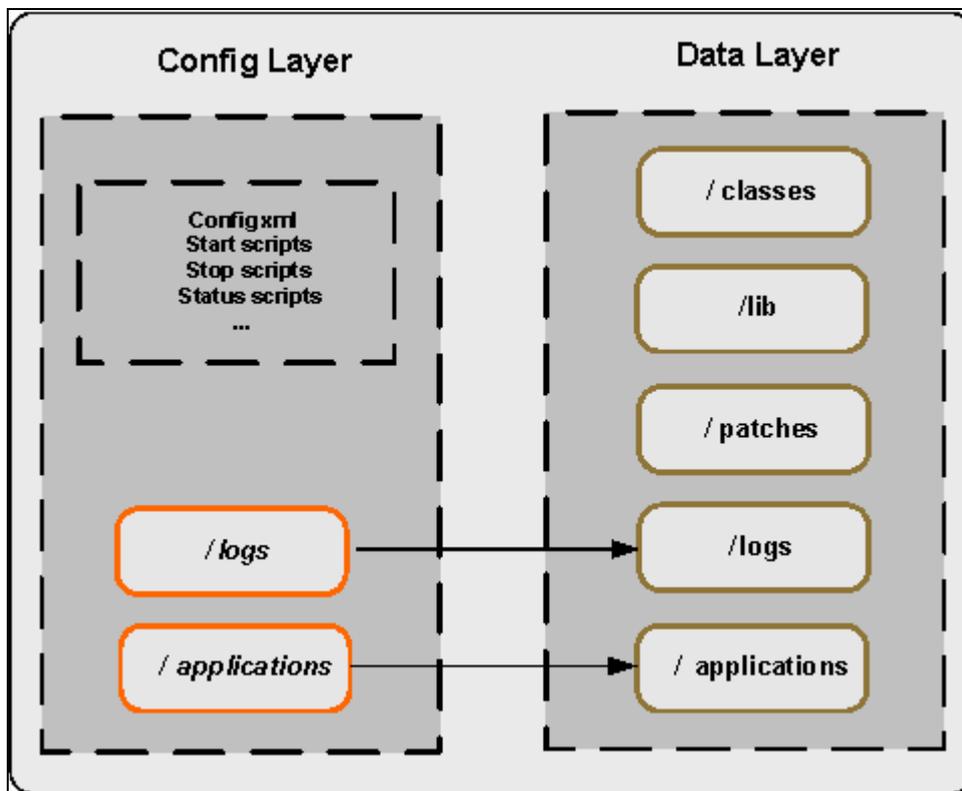


Figure 6 - WebLogic Domain -

5.5.3. References

Useful links are:

<http://www.bea.com>

<http://commerce.bea.com>

<http://support.bea.com/>

<http://edocs.bea.com/>

<http://dev2dev.bea.com>

5.6. Business Objects

5.6.1. Service Description

The general architecture for BO is that shown in chapter 2. Failover is being implemented at server level.

Webi technical support is provided at server level and not at application level. For instance, the Data Centre resolves only the incidents that occur on the server side.

The Data Centre provides the WebIntelligence module only on Unix Solaris platforms. Moreover, all the Webi functionalities are NOT available under UNIX for instance BO vba scripts which are typically for Windows or email diffusion for example.

Access to Webi is available from the Intranet or from Internet through a DMZ configuration.

The service offered is specific to each phase of the application life cycle.

Test phase

In this phase, the Data Centre provides a test instance totally in conformance with the production instance on a separated environment. The client has to test his application and all the BCA jobs to be sure that it will run in production. The Data Centre can tune specific parameters in order to offer the best performance.

Development phase

If the client wants to create or modify Webi reports with the Webi thin client, he has to do it on the test server. When the report is finalized and validated, it can be put in production via FTP.

In case of JSP development with the Webi SDK module, all development has to be done including the tests on the test server. When JSP pages are finalized and validated, they can be put in production via FTP.

Production phase

As soon as the test phase has passed successfully, the service can start in a production environment.

The Data Centre can create for clients an Oracle repository on a dedicated production database.

The Webi file systems are backed-up every day and the Data Centre is able to restore the system to any point in time within the preceding 30 days.

5.6.2. Conditions and constraints related to Business Objects

WebIntelligence Server

The current version is 2.7.x or 6.5.x (recommended).

When the Webi instance is installed, the application URL is sent to the user, including the repository main key if requested.

The Webi instance is installed with the default installation administration parameters. If a client wants to change it, he should send a request for change. No administration console access will be provided.

Only ORACLE databases can be accessed with this instance, an entry should be created in the Data Centre ONS for all the query databases.

If the client wants to connect the Webi instance to his own repository, the ORACLE user/password and the full connect string has to be sent.

The Webi instance can be configured with LDAP authentication. The repository has to be created with the LDAP aliases. At the login page the user will be prompted to introduce his internet login/password.

Broadcast Agent Scheduler Server

The current version is 5.5.x or 6.5.x.

The user may create an agent in his repository. When created it must be activated by the Data Centre Webi administrator. This activation has to be requested via the central helpdesk providing the BCA login/password and a general supervisor login/password. Based on this information the agent can be activated and made available to the user.

Usually the BCA is used to generate at schedule time some reports in PDF or HTML format. Depending on the need, an HTTP or FTP access to download generated files is offered.

WebIntelligence SDK

The current version is 2.7.x or 6.5.x.

This specific module is required for personalizing the infoview standard interface. The Data Centre provides the WIJSP application, which is an example of the infoview interface in JSP, on their instance. Clients are free to update the JSP application or to create a new one but no JSP support can be provided by the Data Centre.

The Webi SDK application can be configured with LDAP authentication. The repository has to be created with the LDAP aliases. At the login page the user will be prompted to introduce his LDAP login/password.

An FTP access is provided to upload JSP pages.

Business Objects Auditor

The current version is 5.1.x or 6.5.x.

This instance is shared between all clients.

An audit schema database is created on a dedicated database to receive all the Webi activities for each client. In order to create a personalized universe on this database, the ORACLE user/password is sent to the client.

To configure the Webi instance to log all activities in the auditor database, a general supervisor login/password has to be sent and a connection to this database has to be created in the repository.

For each application, the Data Centre creates a group in a dedicated repository and a local supervisor is created. This local supervisor is free to create other groups or users in this root group. To connect with the supervisor for this group, the auditor main key is sent.

The standard auditor reports and universe are linked by default to this group. If standards reports are not sufficient, it is possible to upload other reports or universe in this group.

Internet access

It is possible to access the Webi instance from the internet on a separated server isolated by a DMZ.

Files access with ProFtpd

An –internal only- FTP access limited to the European Commission's network, based on LDAP authentication may be provided to:

- Retrieve reports generated by the BCA or upload trigger files in order to use the BCA file watcher option;

- Upload JSP pages in order to update the Webi SDK application.

Http file access

An HTTP internal access can be provided for downloading reports generated by the BCA, over the European Commission's network.

The Data Centre may (on demand) apply security parameters on different directories in order to control the access. The user will be prompted for his LDAP login/password to access the directory. The application owner has to send to the Data Centre, for each directory, the list of the LDAP users or DGs who can access it.

5.6.3. References

Useful sites are:

<http://www.businessobjects.com/>

<http://www.techsupport.businessobjects.com/>

5.7. Oracle

5.7.1. Service description

The Data Centre offers an Oracle database hosting environment on UNIX for development, test and production use (with failover for production environments – see schema below). For this environment, the Data Centre takes care of numerous aspects, which will be developed in the following sections.

The Data Centre creates users and table spaces in these shared databases as requested by the client. Once they are created, the client has access to this new environment through the requested users using the usual Oracle client tools (either from the client's PC or from a UNIX environment that the Data Centre can also make available). The Data Centre considers that, from the moment that the environments have been given to the clients, they are under the responsibility of the client.

The Data Centre offers shared databases for use by several applications. Shared databases are dedicated to one client (e.g. a DG or a large system like for example Europa) and host all the data of the different applications for that client. The advantage of sharing databases is the reduction of the resource overhead (disk, memory, CPU).

The databases are installed on a number of machines dedicated to Oracle. However, databases from several clients share the same machines: as a general rule there is no machine dedicated to a single application.

Database types are development, test (staging, acceptance), production and training (separate databases are created for each type). Machines are dedicated to one type of database (e.g. a machine for "non production databases" and one for "production databases").

The choice of where a particular application is put depends on the information provided in the request forms.

The Data Centre also offers environments and support for the ORACLE related tools which have been approved.

The Data Centre does **not** take responsibility for:

- Application performance issues not due to the server;
- Application errors not due to the server;
- Application tuning (including database schema tuning);

- support on the client side;
- Creation or modification of application schema elements (e.g., tables, views, indexes, etc.);
- Data manipulation (modifying application data in the database);
- In general all interventions that don't require DBA privileges;
- Application security aspects.

To summarise, the Data Centre is responsible for the container (database) and the client is responsible for the content (the data and all application aspects).

The Data Centre is involved ideally in all phases of the development life cycle.

Development and Test phases

So as to supply a qualitative service, the Data Centre wishes to be involved as soon as possible in projects which will be hosted. This will allow the Data Centre to formalize and possibly, to adapt in an early stage the entire environment details, in order to guarantee the best integration with the production environment

For this purpose, a development and/or test environment, similar to that of production can be made available.

Production phase

Intensive contacts between development teams and Data Centre teams are necessary prior to putting a system in production. The following list must be verified before production is authorized:

- All the conditions and constraints for use respected (see chapter below);
- The database design is documented;
- Object creation scripts (with storage parameters) are supplied. An acceptable alternative is an export file;
- A representative from the client with technical knowledge about the installation is present or available;
- A detailed installation guide has been provided;
- All the (additional) products must be part of the official product list.

The Data Centre can provide production databases in a failover configuration if required.

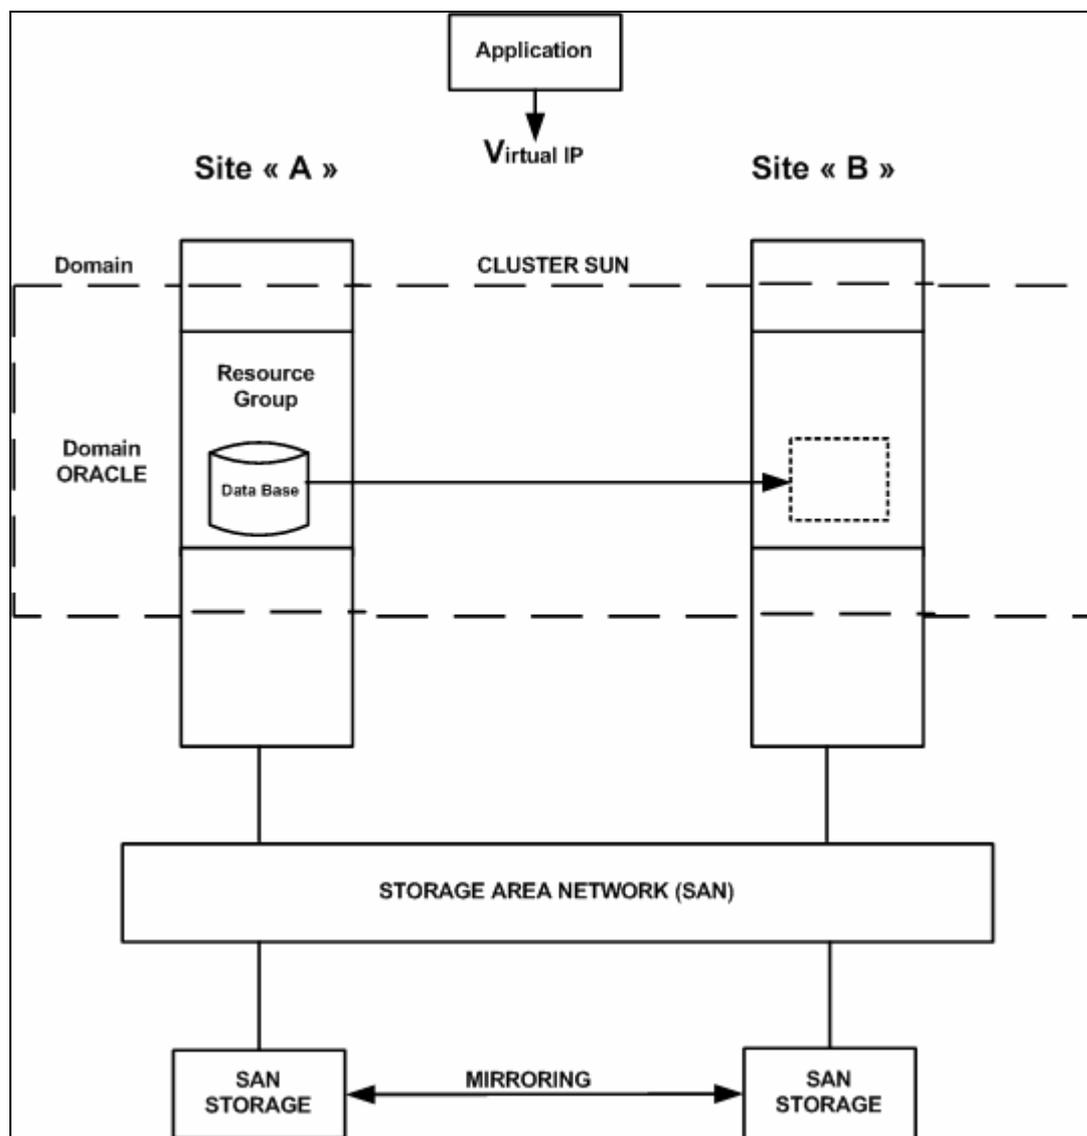


Figure 7 - Failover Architecture -

This means that services provided by a failing machine (hosting databases) can be transferred and restarted on a different machine. This mechanism is transparent to the users as the new machine adopts the network identity of the failed machine. The perceived downtime for the users is minimal.

In addition to the normal operations services, the Data Centre provides additional support and maintenance services to assist clients in case of malfunction or failure.

Services to guarantee the operation of the hosted system include:

Backup

The Data Centre takes care of daily backups. Backups are online and stored on cartridges. Full backups are executed once a week and incremental backups once or several times a day depending on the activity of a database.

All databases are in non-multiplexed archivelog mode. Archivelog files are backed up several times a day (this is part of the incremental backup).

As an additional but non-guaranteed backup measure, full exports are taken every day for databases that don't exceed a certain size.

Restore

At the request of the client, the Data Centre performs restores of data bases according to the information provided.

The execution time for a restore depends on the size and on the point of time to be restored to. The further away that point in time is from the last full backup, the more incremental backups will have to be applied and the longer the restore will take. Partial restores using the export files are generally faster but don't allow restores to exact points in time.

Monitoring

All production databases are monitored for availability and capacity problems. Preventive and corrective actions can be taken around the clock.

User administration

The Data Centre takes care of the usual work related to user administration: creation, deletion, password changes, privilege granting on request of the user (for instance in case of a lost password).

Data transfer between databases

The Data Centre performs data refreshes from one database to another (e.g. production to test) only if they require DBA privileges (e.g. copy of a whole database or copy of several users). The Data Centre does not perform transfers that can be performed with basic user privileges.

If these data refreshes are systematic and/or part of the application (e.g. transfer from staging to dissemination in the case of a web application) then they are to be carried out by the client.

Tuning

The Data Centre offers tuning of the infrastructure as part of the performance test process during the acceptance phase. However, no specific tuning is offered for applications except for solving problems incurred by the server.

Ideally, all phases of the development cycle are executed in the Data Centre environment. If this is not possible, then the application should periodically be deployed on a test database in the Data Centre in order to assess whether the application complies with the Data Centre environment.

Access to the database hosting service in the Data Centre can be considered at various stages, from the feasibility study to the production phase. Planning must be established as soon as possible. For major projects, the Data Centre must be involved as early as possible in the lifecycle, so as to take into account possible requirements for supplementary resources.

5.7.2. Conditions and constraints related to ORACLE

General technical constraints

All databases are created by the Data Centre. The SIDs (System Identifiers) of the databases follows a common non-negotiable naming convention

The choice of database used for a particular application depends on the information provided in the request forms: Depending on the requested application status (production, staging, development, testing, training/demo) the new environment is created in databases on production, test or development servers. The application type is also important for the decision of where to put the new application. If it is for the Europa website then the new environment is created in shared databases dedicated to Europa. For IntraComm, the environment is created in shared databases dedicated to IntraComm. Applications that are only of specific interest to a DG are put into shared databases dedicated to one client. For every corporate system, a corresponding dedicated database is created.

The Data Centre installs the Oracle software according to the Optimal Flexible Architecture proposed by Oracle. Several versions of Oracle may be installed in parallel on the same machine.

When Oracle releases new patches, releases and versions, the Data Centre aims to install them and align all the databases on the latest possible level. If necessary migrations are executed, in close collaboration with the DGs.

All databases must use the Cost Based Optimizer. The Data Centre analyses all databases on a daily basis for CBO to work correctly.

The allowed character sets for the databases are UTF8 (recommended) and WE8ISO8859P1.

All databases are accessible using SQL*Net through the Commission's Oracle names server. The use of the names server is the recommended way of connecting to the databases.

For each application and if requested by the client, the Data Centre supplies a UNIX account with access to all the usual Oracle tools (sqlplus, exp, imp ...).

All databases are configured for connection auditing. This means that access log files are kept which respect data protection regulations.

All databases on the same machine share one listener for 'SQL*Net'.

All databases have their control files and redo log files multiplexed.

The data files are not auto extensible.

The standard settings recommended above are summarized in the table below:

SETTINGS	VALUE
Optimizer mode	Cost Based Optimizer
Character set	UTF8: recommended for all new applications. WE8ISO8859P1: accepted
Auditing	Session
Control Files	Multiplexed
Redo Log Files	Multiplexed
Data Files	No Auto Extend

Data must be indexed where appropriate in order to provide acceptable performance.

Development and test databases that are inactive for 5 consecutive days are automatically shut down to preserve system resources. They will be restarted at the client's request.

When UNIX scripts are used to execute some automated operation these scripts must use the "Korn" shell language. These scripts must call the 'oraenv' script provided by Oracle to set the oracle environment (as opposed to hard coding the oracle environment variables). When such scripts connect to the database they must use sqlnet connections (as opposed to local connections) to provide for location independence. These scripts must have a log file so that errors can be traced. These scripts are not written by and are not the responsibility of the Data Centre.

If the application uses JDBC to connect to the database then it must use the OCI driver for JDBC instead of the thin driver. This is because the thin driver does not support the Oracle names server and thus such applications will no longer work when the database is moved to a different server. The Data Centre must be informed.

Change management

The application may not be tied to a specific version of the Oracle software (including the non-RDBMS software). It may require a minimal version but it must be able to run with higher releases within the same version (e.g. if an application runs with 8.1.6 it must also be able to run with any higher 8.1.x release).

Maintenance interruptions

In order to maintain the Oracle databases and the environment a certain amount of downtime is required and defined in a service level agreement between the Data Centre and the client.

Naming conventions

Any new application must be able to work together with other existing applications sharing the same database. Amongst other things, this means that the application must use a specific schema to own all the objects (no creation of objects in the sys or system schema is allowed). The use of public synonyms is only allowed when application specific prefixes are used. The same holds true for database roles. An overlap in the user community with other applications must be possible.

All segments must be created in application specific tablespaces (not in USERS, TOOLS or SYSTEM tablespaces).

A naming convention for all database objects is highly desirable but the Data Centre does not impose such a naming convention; the important point is that one is used. This should not be confused with the SID naming convention.

The Data Centre insists that the name of the main schema that owns all of the application objects be prefixed with 'APP_' in order to distinguish it from the other schemas.

Security

The complete application must work without DBA rights. No DBA rights (dba, sysdba or sysoper) **will be granted at any time**. All the individual privileges needed must be specified in a document. This allows the Data Centre to verify which privileges are needed and whether they can be granted.

The client can be granted exceptionally the right to create users but in that case he has to create them in such a way that the temporary and default tablespace do not point to the SYSTEM tablespace.

5.7.3. *References*

Useful sites are:

<http://www.oracle.com>: commercial site from oracle corp

<http://otn.oracle.com>: free technical support site

<http://metalink.oracle.com>: Official support site for customers

<http://docs.oracle.com>: All the documentation

5.8. UNIX/LINUX

5.8.1. *Supported Unix versions*

5.8.1.1. *Solaris*

The Data Centre uses a Reference Configuration for Solaris 9 that was built in cooperation with Sun Microsystems. Support is available for standalone servers, and clustered servers (based on SunCluster 3.1).

Operating System	Version	Status	Timeframe
Solaris	8	Phase out	Q4/2005
Solaris	9	Active	
Solaris	10	Planned	Q3/2006

5.8.1.2. *Linux*

The Data Centre uses the Redhat Enterprise Linux Reference Configuration as prepared by DIGIT A04 (STB). Currently support is limited to standalone servers, a clustering solution is under investigation.

Operating System	Version	Status	Timeframe
RedHat Enterprise Linux	3.x	Active	

5.8.1.3. *Others*

The Data Centre is currently phasing out its HP-UX servers.

Operating System	Version	Status	Timeframe
HP-UX	11.x	Phase out	Q4/2005

5.8.2. Supported Unix services

Every UNIX server is installed with the official operating environment including a standard list of packages.

For security reasons, the Data Centre limits the number of UNIX related services provided for the applications and application owners to the following list:

- Terminal emulation via ssh;
- File transfer to the machine from the EC network via ssh based protocols (scp/sftp)¹;
- Cron-jobs and at-jobs;
- Mail forwarding with sendmail.

Other UNIX related services may be provided on a case-by-case basis and must be requested as early as possible in the application lifecycle (see above).

5.8.3. NFS mounts

- The DC exclusively mounts NFS shares that come from the DC's NAS filers;
 - The DC does not mount NFS shares from machines not fully under his control;
 - The DC does not mount NFS shares between general purpose servers (cross-mounts).
- NFS is not supported in the DMZs or other separated environments.

¹ For legacy applications, inbound ftp on the standard port can be enabled on a case-by-case basis.

5.8.4. Cluster

If the application needs an OS Cluster to run, the Data Centre currently offers a Sun Cluster 3.1 environment based on Solaris 9.

Applications that are run in clustered environments must be aware of the limitations that a cluster places on them:

- (1) All parts of the application must be started by the cluster framework. The start scripts (shell scripts) must be provided to the DC for inspection and usage, and will not be writable by the application owners.
- (2) Stopping and starting the application must only be done via the cluster framework. Killing processes can trigger a failover of the resource group. Exceptions need to be arranged with the DC.
- (3) In the case of a multi-tiered application, which is using multiple resource groups, all tiers must be capable of automatically recovering from a switchover of one or more resource groups that the application is hosted on.
- (4) Applications are not guaranteed to always run on the same physical host. Always use the logical hosts when accessing the cluster resource.
- (5) Applications must not assume that other resource groups will always be on the same node as itself.
- (6) There is no failover for at-jobs. If the resource group is no longer running on the node when the time arrives to execute the at-job, an error message will be generated and sent to the user executing the at-job.
- (7) The application is responsible for registering and de-registering its cron-jobs during start/stop of the resource group.

DC has written two scripts to support applications to properly register and deregister cron-jobs. These scripts should be used in the start and stop scripts of the application. All changes to the crontab must be made via these scripts, 'crontab -e' is NOT supported in a clustered environment. For further information, please contact the DC.

In addition to this, the current cluster environment relies on local accounts, and does not synchronize accounts and passwords between machines. This is being addressed in the UNIX User Management Project.

5.9. Windows

All Windows servers are installed following the latest reference configuration.

Currently this is WINDOWS 2000 AdvancedServer with additional fixes.

5.10. SMTP relays for E-mail applications

5.10.1. Service Description

The Service offered by the E-Mail service is limited to relaying messages submitted to the INSEM3 infrastructure. The messages may come from the Internet, from TESTA or can have been submitted by applications hosted on the internal network or in the DMZ.

The target addresses of the relayed messages may correspond to an INSEM3 mailbox, an application hosted on the internal network or in the DMZ, an external recipient (be it reachable via Internet or via TESTA) or even, a fax number. Note that for security reasons, some of the combinations are not possible, e.g. messages coming from an external domain will not be relayed to an external e-mail address.

5.10.2. Conditions for use of SMTP relays

This section lists the reasons why a DG must use the SMTP architecture managed by the E-Mail service and what are the conditions imposed by the Data Centre for its use.

By default, the access to the SMTP architecture is forbidden for all machines not explicitly authorized

The following arguments can be presented as valid reasons in order to be authorized to access the SMTP infrastructure:

- An application has to send messages to INSEM3 mailboxes;
- An application has to send messages to an external correspondent;
- Alert messages need to be sent;
- An application hosted within the Commission network needs to be reachable from the outside, via Internet or TESTA.

Note that only servers will be authorized to submit messages to the INSEM3 SMTP relays, end-user workstation will not be authorized.

The applications allowed to use the SMTP service must not endanger the INSEM3 infrastructure.

5.10.3. Sending e-mail

This section describes how to submit e-mail from an application using the Internet Mail Technology, namely SMTP (Simple Mail Transfer Protocol).

Based on the underlying E-mail sub-system of the hosting server

The first possibility to send e-mail from an application is to pass the information to the e-mail sub-system of the hosting machine, e.g. the “sendmail” program when the application is running on a UNIX based platform.

The advantage of such solution is to put all the management aspects of the e-mail transaction in a dedicated product. Management aspects include configuration of the e-mail service, retry of delivery in case of problem, DNS resolution, etc. In this case, the application does not have to take care of all those aspects.

Built-in

Another way to interact with the e-mail system is to integrate the e-mail transaction aspects directly into the applications. In this case, the application is responsible for all the management aspects of the e-mail transaction, e.g. configuration, retry, queuing of e-mail, etc.

SMTP relay

The administrator of the application has to configure his system in order to relay the messages to the right SMTP relay. The infrastructure managed by the E-Mail service ensures a transparent load balancing from the applications point of view and a failover mechanism.

5.10.4. Incoming messages

In this case, the applications will be reachable using an addressing schema that has to be defined in agreement between all parties involved (i.e. the E-Mail service, the owner of the application, etc.).

5.10.5. Name resolution

The E-Mail service recommends using the DNS so as to be as independent as possible of any changes that may occur in the future. All applications hosted by the Data Centre have to use the DNS mechanism for name resolution. Applications that do not use DNS will not be accepted for hosting by the Data Centre.

Use of DNS gives flexibility and independence in relation to the changes that can occur at the level of the machines used to relay the SMTP messages. It means that the machines hosting the applications and the applications themselves do not have to be concerned about the IP addresses of the target machine.

5.10.6. Limitations of the Infrastructure

The application has to comply with the limitations imposed by the E-Mail service, e.g. max number of recipients per message, max size per message, etc.

5.10.7. Change management

Redundant environment

The E-Mail infrastructure is fully redundant, meaning that intervention on the machines used by the applications can be done transparently.

Maintenance interruptions

When possible, interventions are planned during weekends or holidays in order to minimise the impact on the end-users.

Information on interventions

As far as possible, planned interventions are announced well in advance so that local support teams are able to forward the announcement to the end-users and so that all concerned are informed.

5.11. Mailbox access for E-Mail applications

5.11.1. Service Description

This service allows applications to use a mailbox in order to send and receive messages.

In order to send or receive e-mail, some applications need to have access to a mailbox. It is not always possible for the administrators of the applications to host such mailboxes in their own environment and so, they need to define them elsewhere.

The Data Centre offers the possibility to define such mailboxes in the E-Mail environment. In addition to the standard mailbox functionalities, applications can also make use of the POP3/IMAP4 and SMTP protocols.

5.11.2. Conditions for use of the Mailbox service

The following arguments can be presented as valid reasons in order to be authorized to have a POP/IMAP mailbox:

- An application has to send messages;
- An application has to be able to receive messages;
- Alert messages needs to be sent;
- An application hosted within the internal network needs to be able to receive messages coming from the outside, via Internet or TESTA. POP and IMAP protocols can only be used within the internal networks. External messages will

be received through the E-Mail infrastructure using the SMTP protocol and then fetched from the mailbox by the application using the POP/IMAP protocol

The applications allowed to use the POP/IMAP service must not endanger the INSEM3 infrastructure. At any moment, the E-Mail service may decide to block such applications if the E-Mail infrastructure is in danger.

For security reasons, POP and IMAP connections cannot be initiated from outside the European Commission's network to access a mailbox.

Messages submitted by the applications have to comply with the **“ACCEPTABLE USE OF THE COMMISSION E-MAIL SYSTEM”** (Administrative notice N°88-2002).

All traffic that does not comply with this acceptable use of the E-Mail system will be blocked.

5.11.3. Sending of e-mail

This section describes how to submit e-mail from an application.

Sending e-mail using SMTP

Applications that want to send messages using SMTP protocols have to comply with the recommendation explained in chapter 5.10

Sending e-mail using MAPI functions

Another possible solution to send messages is to use the MAPI functions.

5.11.4. Incoming messages

POP3/IMAP4

When an internal application needs to receive e-mail, a specific functional mailbox may be defined. This mailbox may be accessed using the POP3 or the IMAP4 protocol. The access protocol may be implemented within the applications accessing the mailbox or by using a specific program, which is the easiest solution.

This service is only available for applications hosted within the internal network, POP3 or IMAP4 are not available from the outside

Internal sub-domain routing

In this case, the applications will be reachable using an addressing schema which is part of a sub-domain of the Commission domain.

5.11.5. Name resolution

The E-Mail service recommends using the DNS so as to be as independent as possible of any changes that may occur in the future. All applications hosted by the Data Centre have to use the DNS mechanism for name resolution.

Use of DNS gives flexibility and independence in relation to the changes that can occur at the level of the machines used to host the POP3/IMAP4 mailboxes. It means that the machines hosting the applications and the applications themselves do not have to be concerned about the IP addresses of the target machine.

5.11.6. Change management

Maintenance interruptions

When possible, interventions are planned for weekend or holidays in order to minimise the impact on the end-users.

Communication of interventions

As far as possible, planned interventions are announced well in advance so that the local support teams are able to forward the announcement to the end-users so that all concerned are informed.

5.12. ECAS & LDAP

5.12.1. ECAS (European Commission Authentication Service)

ECAS is the authentication system to be used for authentication of Commission users in web applications developed for the Commission.

ECAS supersedes the CED LDAP directory for most authentication purposes. The major exception to this is access to and from the Internet, where the proxy servers still authenticate using the CED LDAP.

ECAS is a common service that provides secure authentication facilities for users of Commission applications, ECAS is compliant with the Commission's security policy.

In addition to secure authentication, ECAS provides a single sign-on capability for users of its client applications.

The way to invoke the ECAS capabilities varies according to the client application platform.

The initial coverage extends to Weblogic, ColdFusion MX and Apache is planned.

5.12.2. CED LDAP Server (Lightweight Directory Access Protocol)

The CED LDAP server (Commission Enterprise Directory) provides applications with information about users such as name, email address or organisational assignment.

The CED LDAP authentication schema is NOT compliant with the Commission's security policy and therefore should NOT be used to authenticate users.

CED LDAP may be used for:

- "Off the shelf" products that cannot be adapted to use ECAS or NT authentication;
- Simple user identification (not authentication!);
- Small queries with a limited number of results.

5.13. Active Directory

The Active Directory Service (ADS) offers an environment permitting the Windows authentication of Commission users and computers from anywhere within the security boundary of the internal Commission network, facilitating the sharing of Windows resources across the Commission IT environment. However Active Directory is NOT the authentication system for Information Systems. (See chapter above).

The Data Centre provides a high-availability environment for the ADS servers with built-in redundancy and 24/7 monitoring. All ADS domain controllers are placed in physically secure environments to prevent unauthorised access to data on the servers.

In addition to the production environment, the ADS provides a permanent test environment in the form of a scaled down replica of the production environment

5.14. Corporate Web Content Management System

DOCUMENTUM is the strategic platform chosen for web content management. The Data Centre creates, configures and maintains test, training and production environments for WCM.

5.15. SAS

Several versions of the statistical package SAS are currently installed:

- v6.12;
- v8.2 (32 bits);
- v8.2 (64 bits) for web log analysis (contains Webhound);
- V9.1.3

The recommended version is V9.1.3 in conformance with the contract negotiated. Optional modules are provided only for specific requirements.

Possible Configurations:

- (1) In Client / Server via X-terminal session under UNIX;
- (2) With Integration Technologies;
- (3) Via browser;
- (4) Light Client Enterprise Guide on PC;
- (5) As fat client (all SAS on PC – mainly for developers).

Configurations 2, 3 and 4 require a specific installations on the server for the products IntrNet, Connect, Share, spawner, broker.cfg, etc ...

Monitoring

There is no specific monitoring for the infrastructure, only a monitoring of the file system containing the binary programs.

End-2-End monitoring must be agreed in a SLA between the DGs and the Data Centre for each information system.

5.16. FTP store

The guidelines related to the use of FTP STORE can be found on the Intracom.

The service FTPStore consists of two servers:

- Server for FTP via Internet;
- Server for FTP via TESTA.

The two FTPStore services allow temporary storage of files. They only support the FTP protocol (passive or active mode).

To use FTPStore, an account must exist on the server, access is granted after fill in the FtpStore User Registration Form.

The user must have an ftp client, in either text or graphic mode.

The confidentiality of the data is not guaranteed by the network and system which constitute FTPStore. To obtain confidentiality, the users must encrypt and decrypt the files before putting and after getting them, on the sending and receiving machines.