EUROPEAN CENTRAL BANK

EUROSYSTEM

8 March 2007

CONTRIBUTION TO THE T2S FEASIBILITY STUDY TECHNICAL FEASIBILITY

Introduction

This document is the contribution to the Feasibility Study of TARGET2-Securities (T2S) with regard to the technical aspects.

The technical feasibility of the T2S platform – and subsequently the technical cost estimates – was initially prepared assuming a completely greenfield approach (i.e. without leveraging synergies with any existing systems). In parallel, an estimation was made of the synergies resulting from the reuse of the components already adopted in the TARGET2 Single Shared Platform (SSP) (the “T2S on T2” concept). This process is summarised in the Economic Feasibility Study.

Synergies with TARGET2 are expected to result in a highly resilient infrastructure and in cost savings of roughly 20% when compared with a greenfield approach. The Governing Council therefore decided that the User Requirements will be prepared on the understanding that T2S will be developed and operated on the TARGET2 platform. The present document presents a possible way to implement T2S, and in particular T2S on the TARGET2 platform.

Feedback is welcome on this document, in particular on the validity of the assumptions made regarding the desired functionalities and the volume estimations. This feedback will facilitate the preparation of the User Requirements and will serve to adjust, if necessary, the elements that form the basis of the cost assessment of the platform.

Section 1 of this document describes the synergies that can be derived from the implementation of T2S on the TARGET2 platform. Section 2 presents the assumptions on the functional aspects and the workload of the system. Sections 3 and 4 describe the application and infrastructure high-level designs respectively.

This study concludes that there are currently no technical obstacles to building and running T2S. The case is reinforced by the substantial synergies that will result from developing and operating T2S on the TARGET2 platform.

1. Synergies

The T2S on T2 approach will take advantage of the synergies that can stem from reusing the TARGET2-SSP architecture, including:

- full integration of T2S into the architectural design of TARGET2 (two regions and four sites), which is the most advanced architecture from a business continuity perspective, and offers clear advantages from an organisational point of view as well;
- the use of the technical infrastructure of TARGET2, based on production and test and training environments, possibly scaled by increasing the already available components or adding new ones;
- the adoption of tools already established for TARGET2 with regard to the change management system, the trouble management system and technical monitoring;
- the use of the T2 internal network (3CBNet) for data exchange within the T2S system, adjusting the bandwidth appropriately;
- the flexibility of the segregation level between T2S and TARGET2, to ensure an optimal balance in terms of functional synergies and overall resilience;
- possible load-balancing between T2S and TARGET2, since peak hours differ in the two systems: payment activities occur mainly during the day, while most securities transactions are settled at night;
- enhancement of the archiving functionality that already exists on the TARGET2 platform for T2S purposes. This functionality will allow the long-term storage of T2S information.

The T2S on T2 concept also offers the following clear benefits from an application viewpoint:

- Reuse of SSP architectural features:

  The T2S application architecture will reuse some of the specific features of the SSP. The main reused architectural paradigms are inter-module communication techniques and protocols: different building blocks will communicate using a standardised technique, regardless of the technological layer, which can be mainframe messages or web-oriented tasks.

  A unique internal interface reduces the overall complexity of the application and increases reliability and maintainability. A solid and reliable communication and cooperation infrastructure will ensure that applications can be developed while remaining clearly distinguished from each other, without any hidden dependencies and unnecessary coupling that typically makes it necessary to maintain different components when a single component needs to be modified.

- Reuse of SSP application software:

  The reuse of some SSP components is foreseen within the T2S system, specifically to:

the accounting, where the structure (main and sub-accounts) and the liquidity management will be used, allowing the users to keep the cash management interfaces and internal processes as they are implemented for TARGET2;

the interface module, where the SSP *Information and Control Module* (ICM) will be adapted to T2S by adding new screens and designing new workflows. Reusing the same product will provide benefits for external users, ensuring a homogeneous presentation layer for TARGET2 and T2S, making it possible to integrate (at user workstation level) the full set of services provided by the SSP (i.e. TARGET2 and T2S);

the cooperative processing layer for external users, represented by the SSP *Ancillary System Interface* (ASI), which may be partially reused in order to provide a sophisticated XML-based application-to-application interaction tool. The significant benefits for users are expected as the ASI paradigm is well-known by central securities depositories (CSDs) acting as "ancillary systems" in SSP.

Finally, a state-of-the-art operations and support organisation will be put in place for TARGET2. In particular, this will allow operational experts to be located in two different geographic sites, which will operate and support both services with a high level of availability.

2. Assumptions

2.1 Functional assumptions (complexity-related)

The following assumptions apply to the application:

- The number of used message types in T2S is higher than in TARGET2 (a ratio of approximately 16:10);
- The process for security transactions is much more complex than the process for payments;
- MT5xx message types are more complex than currently used message types in TARGET2;
- The matching procedures depends on different instruction types (RvP, DvP, RFoP, DFoP, DVD) and underlying business cases;
- Complex parameterisation of matching rules is needed (possibly at CSD level), because the matching rules are not yet harmonised;
- The number of different instructions types (> 40) is high, and consequently the number of possible workflows is high as well;

- A high number of complex market rules (possibly at CSD level) will need to be covered at all stages of the lifecycle;
- Real-time and batch optimisation will require more algorithms than the seven foreseen in TARGET2;
- Complex optimisation algorithms are needed to bring into line simultaneously cash and security accounts, considering possible additional dependencies on linked instructions as well as time constraints with regard to finding the executable chain of instructions;
- Multi-currency settlement may be required;
- More items per function must be supported by ICM in T2S than in TARGET2;
- A huge variety of data are necessary to run the system, especially CSD participants, ISINs, accounts, rules, access rights, etc.;
- Different sources of static data have to be incorporated;
- Changes to static data have either to be activated in real time in all relevant modules during the business day, or might only be changed for the following business day.

2.2 Workload assumptions (system sizing)

The following volumetric assumptions have been taken into due consideration:

- The T2S average daily workload would amount to 2.1 million transactions (which corresponds to about four million settlement instructions).
- The system is capable of handling peak days with twice the above volumes (4 million transactions – 8 million settlement instructions) using, if necessary, capacity-on-demand features.
- 70% of transactions are processed during the night in batch mode – i.e. 1.47 million T2S transactions – with only 30% processed during the day in real-time mode – i.e. 0.63 million T2S transactions. Real-time processing, which is more demanding than batch processing, is the main criteria for capacity planning.
- The peak hour workload for real-time transactions is estimated as 30% of average daylight workload (189,000 transactions).[1] The worst-case scenario is considered, with the TARGET2 peak hour concurrent with the T2S peak hour (considering that the TARGET2 capacity planning is based on 380,000 payments per day, 500,000 payments on a peak day, 105,000 payments per

[1] The system will be able to process 1 million instructions per hour during the night in batch mode.

peak hour). The overall peak hour workload for real-time transactions on the SSP will be 294,000 transactions.

- The ratio between T2S transactions and messages is estimated at about 1:6.

3. High-level design of the application

The aim of this description of the T2S application architecture can be considered as an initial contribution to the architectural requirements which will drive the definition of the final application architecture.

3.1 Overall application architecture

The following schema is a high-level representation of the system from a functional point of view. It is based upon a set of main *functions* (or *modules*) that are conceived as independent, structured and “loosely coupled” application components. Communication among modules relies upon a standardised protocol.

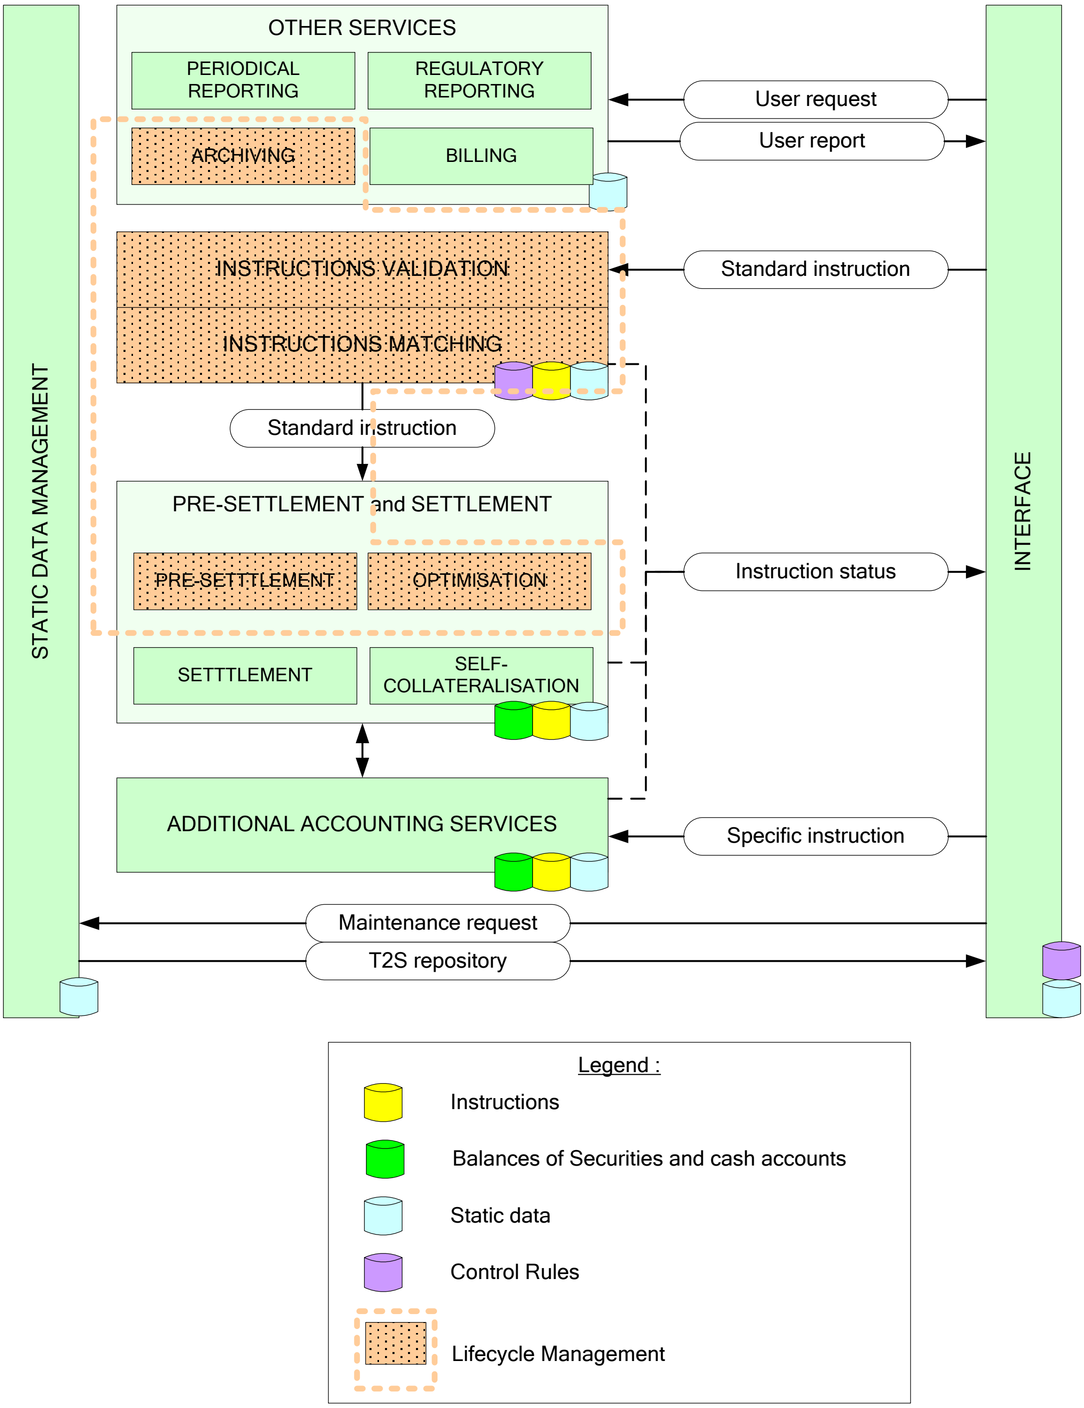

Instruction validation

This module is responsible for receiving and validating instructions against global and local (at CSD level) rules. It delivers each validated instruction to the other T2S module for further processing.

Instruction matching

This module matches instructions received and not already matched. Matching criteria can be global or local, and can foresee parameter-based matching tolerance. Partially matched instructions can be split.

Pre-settlement and settlement

Pre-settlement is responsible for collecting, preparing and prioritising instructions. It also provides a liquidity forecast for night-time processing.

Settlement controls optimisation algorithms, self-collateralisation, and the booking of securities and cash for daylight and night-time procedures. It also ensures the realignment of securities balances of investor CSDs with those of issuer CSDs.

Additional accounting services

These provide services for accounts such as opening, closing or modifying. Non-standard operations can also be handled (e.g. creating/destroying securities, conducting custody management operations, blocking/unblocking, pledging/unpledging and opening/closing loans).

Specific procedures are in place to reconcile with CSDs and provide end-of-day processing.

Other services

T2S provides other services such as archiving (also for legal purposes) and billing, with periodically produced user reports.

Static data management

This provides T2S modules and T2S users with an integrated and consistent set of common information. Static data management represents a single point of creation, modification and deletion of static data. Users can access static data accordingly with specific profiles or modes (i.e. *four-eye* vs. *two-eye* principles) established at local level.

Static data management module also provides versioning facilities and exceptional procedures for managing urgent changes.

Interfaces

The T2S interface supervises the whole traffic flow of information from and to users. It performs syntax, data format and authorisation checks.

Currently, data exchange via interfaces is foreseen using SWIFT FIN, FileAct and InterAct. Furthermore, SWIFT Browse can be used, as can application-to-application and user-to-application (A2A, U2A) approaches. The impact of dedicated links for CSDs on the interfaces will be analysed.

It is envisaged to reuse for the T2S application the same technologies adopted within TARGET2 (mainframe, adopting web-oriented and standard languages) in order to achieve the maximum level of synergies as well as availability. In any case, new technologies will undoubtedly emerge in the coming years, and the possibility of adopting them will be assessed in order to remain as up to date as possible.

The main typologies of the elementary components of the application will be:

- Real-time transactions, whereas response times are critical and must be taken under strict control. These transactions are expected to be widely parallel or, at least, to be "parallelisable", in order to ensure application scalability (see below);

- Batch processes, as workload-intensive components, where system efficiency has to be maximised (e.g. night-time processing). This kind of process will be coordinated by a single scheduling controller, ensuring centralised monitoring of activities.

Service Level Agreement (SLA). It is assumed that the SLA for T2S will be similar to that of TARGET2 in terms of confidentiality, availability and non-repudiation.

It is envisaged to implement the interfaces modules through the general reuse of the ICM. In addition, the network interconnection layer will make use of the same middleware infrastructure adopted in SSP, offering both A2A and U2A modes for all functions. The components will run on the *Websphere* application server in the mainframe environment. This provides, at the same time, advantages via the web architecture and the reliability and scalability of the mainframe environment. All functions are in any case available in both A2A and U2A mode.

3.2 Application scalability

The main architectural concerns deserving of special attention are linked to the scalability of the application.

These architectural issues arise from the need to achieve a high throughput level and, more importantly, to make the application to power up the system scaleable in case of unexpected peak occurrences, so as to deal with any unforeseen increase in workload.

To enable extensive parallel processing, the architectural design must foresee a strong concurrency control over the whole application. This will notably:

- reduce database holding time, splitting complex processing between several "light" tasks; elementary tasks should preferably be asynchronously linked to each other;
- design the database in order to maximise the isolation of critical pieces of information;
- predefine the data access paths in order to prevent database contentions;
- the application will adopt, whenever possible, an "uncommitted read-access" technique in order to read data which are not definitive (i.e. the task updating these data is still running when the data are fetched). This implies that the application should be designed to be resilient to this degree of data uncertainty, capable of making the needed consistency checks and taking into account the possibility of restarting processes that fail the checks.

To further increase the throughput, the *pipeline processing* technique could be used. This consists in splitting the application processes into elementary modules that are put into chains which can be scheduled in parallel. Considering the high volume of transactions to be processed, the technical design of T2S will rely, when appropriate, on this technology.

Database access is a critical issue for settlement processing. Appropriate measures must be considered in order to concentrate the load on the databases to the essential processes, and to avoid issues of contention.

4. Infrastructure high-level design

T2S aims at guaranteeing significantly high performance, resilience and processing capacity.

The main foreseen components of the T2S architecture are identical to these of the SSP architecture:

- A central processing system that is fully scalable (mainframe);
- A storage subsystem with synchronous and asynchronous mirroring functionality for efficient business continuity support;
- A secure wide-area network to connect CSDs' participants, market infrastructures, national central banks (NCBs) and CSDs (SWIFTNet);
- Optionally dedicated links with CSDs, mainly depending on traffic volumes[2];
- A dedicated network to connect the different processing sites (3CBNet);
- System software and management tools;
- Security systems (firewall, etc.).

One of the main characteristic of this architecture is that it is scalable both horizontally and vertically.

4.1 Environments

The T2S platform needs multiple independent processing environments to support technical integration, internal acceptance, customer acceptance and live operations. The number of environments fits the application development lifecycle.

4.2 Technical operation

The technical operation of T2S will be based on a high degree of automation in order to reduce human operator errors and to simplify the management of the infrastructure. In addition, T2S will have a good level of controllability and security (confidentiality, integrity, auditability, identification/authentication and access rights permission).

The mainframe-based operations will run alternately in the two different regions.

In each region two sites will be available: the primary site and the secondary or recovery site. The two regions will be connected by a dedicated network with adequate bandwidth that guarantees remote copy (asynchronous mirroring) between them.

[2] Dedicated connectivity for CSDs will be considered if needed based on the User Requirements.

4.3 Security

T2S will be fully compliant with the TARGET2 Security Requirements and Controls (T2SRC) as established by the Eurosystem. In particular, security will be managed at the system and network level. The T2S perimeter defence will be built using network firewall systems to filter communications and log network access.

4.4 External network connection

Dedicated connectivity for CSDs is not specified yet, as the technical solution may be radically different depending on the volumes that CSDs need to process, and the proportion of batches vs. individual messages. Implementing a single solution may prove rather inefficient and costly for smaller participants or may require cross-subsidising between CSDs.

Accordingly, there should be the possibility to implement an option adapted to each CSD's needs. A CSD with very high volumes may need a high-capacity dedicated line-based solution with full backup and extreme availability for batch transfers and individual transfers, whereas smaller CSDs may have lower capacity needs and might use a different solution, such as SWIFT.

However, in the case of a dedicated connection with a CSD, it will be necessary to provide and develop application and transport services such as SWIFTNet.

From a T2S system perspective, it is necessary to offer the capacity to receive messages from several sources and to transmit them into the processing queues. It is understood that T2S will send and receive messages in one single format, regardless of the connection type. It will be the responsibility of the transmitting parties to format their messages accordingly.

From the perspective of installing and supporting such a hybrid solution, there cannot be open-ended flexibility. It is therefore likely that the Eurosystem will define a "catalogue" of connection solutions and select the appropriate one on a case-by-case basis. In addition, the Eurosystem will retain responsibility for specifying the provider and the equipment to be used, where necessary. This might also cover message pricing negotiations with the providers.

Having a dedicated connection with a CSD will require application and transport services to be provided. This is all in place in the TARGET2 infrastructure for SWIFT.

4.5 Internal network connection

T2S will adopt the same business continuity solution as the TARGET2-SSP (see paragraph 4.7), based on two regions and four sites (two sites for each region). The internal connection between the two sites in each region is based on a fully scalable fibre-optic channel (dark fibre with a multiplexer device – DWDM). These are used for remote control, synchronous remote copy of data, and remote access.

The current network bandwidth for the SSP is:

Region 1 SITE A - Region 2 SITE D   1 Gbit/sec direct point to point link
Region 1 SITE B – Region 2 SITE C   1 Gbit/sec direct point to point link

Region 1 SITE A – Region 1 SITE B   2 Gbit/sec
Region 2 SITE C – Region 2 SITE D   2 Gbit/sec

Connection between T2S and TARGET2 is based on internal channels.

4.6 Technical monitoring

The main aim of technical monitoring is to detect problems well in advance before a service interruption or degradation occurs, with the objective of solving the problem immediately through either automatic or human interventions.
The main objectives of technical monitoring are:

- controlling the availability of physical and logical critical resources (status);
- managing and reacting in case of specific events such as error messages, hardware faults and system or application software failures, etc.;
- measuring system throughput and application response time to verify in real time the service level; and
- measuring resource utilisation, and reacting if some thresholds are exceeded.

Automation improves availability and reduces operational costs by managing system operations as automatically as possible. It also reduces the amount of tasks that are normally carried out by the operators, thereby avoiding human errors (i.e. system activation/deactivation, etc.). Moreover, alert conditions due to abnormal situations (restart failed processes) are automatically dealt with.

Depending on the seriousness of the problem (a warning, trivial error, serious error, etc.) an operator intervention (telephone call, mobile call) may be required; information based on automatic event analysis is provided to the operator to help determine the problem.

4.7 Business continuity

Business continuity is a crucial aspect of the T2S design, especially with respect to its infrastructure, as it will determine the ability to perform tasks under abnormal circumstances such as a trivial failure (i.e. disk failures), a serious failure (i.e. a node/system failures), a site failure (i.e. the unavailability of one or a group of buildings), and regional failure (i.e. the consequences of September 11-style disaster events, which could result in unavailability over a wide area). As a result, the infrastructure of the T2S will be based on cutting-edge technology for both business continuity and contingency measures.

Deploying T2S on the SSP architecture allows these objectives to be met by reusing the TARGET2 architecture design (two regions and four sites), which is the most advanced architecture from a business continuity perspective. The T2S infrastructure will be based on cutting edge technology for both business continuity and contingency measures and service level, and will guarantee:

- high throughput (to process an abnormal peak workload or a peak workload due to a temporary stop);
- high-level performance;
- automation of operations;
- manageability (i.e. no ambiguity in operations, reduced complexity);
- controllability and audibility (confidentiality, integrity, identification/authentication and access rights permission).

Staff must be available under all circumstances.

With a view to ensuring business continuity, T2S should be protected from the effects of short-term security failures on the one hand and major failures and disasters on the other.

After September 11, particular importance has been attached to the following business continuity objectives:

- Rapid recovery and timely resumption of operations following: (i) wide-scale regional disruption; or (ii) the loss of staff or inaccessibility in at least one major operating location; or (iii) the unavailability of a number of sites that may even be located in different areas (a rolling disaster, simultaneous terrorist attack, etc.).
- A high level of confidence so that critical internal and external continuity arrangements are effective and compatible.

The business continuity model proposed for T2S should be able to face regional disasters. The recovery of activities and processes during a wide-scale regional disruption calls for the establishment of out-of-region arrangements for operations, plus the related personnel and data needed to support such an activity. The objective of establishing out-of-region arrangements is to minimise the risk that primary and back-up sites and their respective labour pools could be impaired by an individual large-scale regional disruption. Out-of-region sites should not depend on the same labour pool or infrastructure components used by the primary region, and should not be affected by a wide-scale evacuation or the inaccessibility of the region’s population. As a consequence, it could be necessary to operate from the other region for quite a long time, and therefore a high level of system availability should be guaranteed there too.

The proposed model is based on having two regions and four sites (**two sites for each region**).

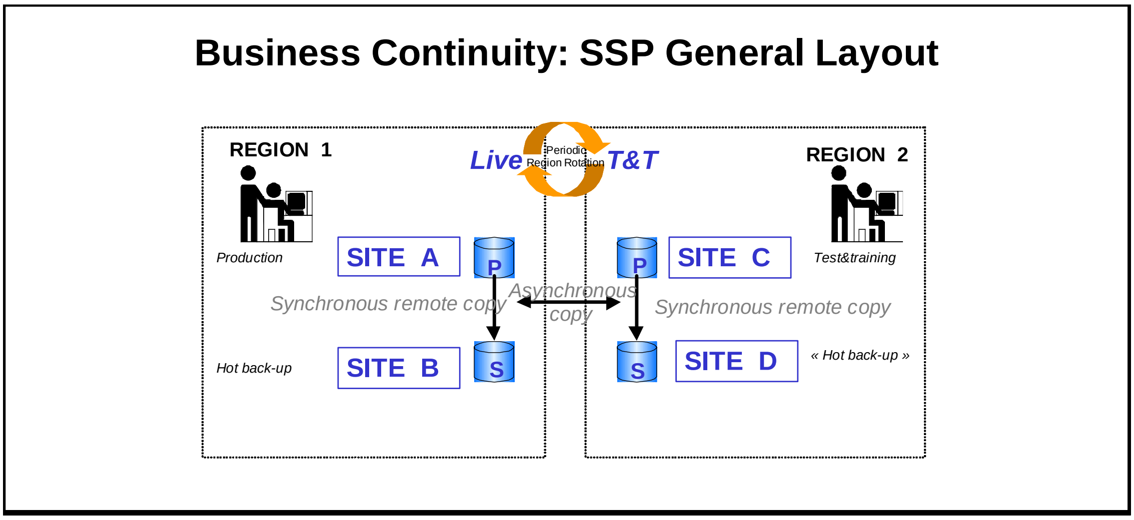

Integration into the TARGET2 infrastructure will ensure:

- Intra-region recovery: in each region the two sites are located at a distance of some kilometres and connected by means of a fibre-optic channel. The two sites are fully equivalent, and the same technological resources are installed in each. Recovery within a region is assured by synchronous remote copy activated on the entire T2S environment between the two sites in the same region. During normal operations, the T2S live environment would run in one site, while the other would remain in "hot standby" mode.
- Inter-region recovery: recovery from one region to the other is only possible by asynchronous remote copy, activated on the entire T2S environment.

The T2S Business Continuity Model will take full advantage of the organisation of the support already in place for TARGET2. The availability of skilled personnel with a suitable technical and operational background is mandatory. In order to deal with regional disaster recovery, there need to be adequately trained staff at both sites to meet recovery objectives for a period that could last for quite a long time. Periodical live environment exchanges between the two regions are necessary to maintain skilled personnel in both regions (inter-regional rotation). This is a very effective way of guaranteeing the alignment between the infrastructures of both regions and of maintaining staff in each region who are skilled and prepared for such a recovery situation. It is also worth mentioning that the organisational and procedural arrangements of such a rotation process are already in place for the TARGET2 SSP and could easily be used for T2S as well.

The workload would be distributed between the two regions. Periodical swaps between the two regions ("rotation") would ensure that the technical and operational staff in each region remain fully trained. The system and the application software will be kept updated in the two regions (via asynchronous remote copy), so that after the rotation the system can restart with the same customisation (i.e. naming conventions, security policies, management rules, etc.). T2S offers a single interface to its users, i.e. they are unable to tell in which region a certain module is running. Moreover, rotation is completely invisible to NCBs, users and market infrastructures, i.e. no configuration changes in customer systems are envisaged.

Address: Kaiserstrasse 29, 60311 Frankfurt am Main, Germany
Postal address: Postfach 16 03 19, 60066 Frankfurt am Main, Germany
Telephone: +49 69 1344 0
Fax: +49 69 1344 6000
Website: http://www.ecb.int