



Council of the  
European Union

Brussels, 23 July 2014

12131/14

**CSCI 23**  
**CSC 175**

**NOTE**

---

From : General Secretariat of the Council  
To : Delegations  
Subject: Information Assurance Security Guidelines on Data Separation

---

Delegations will find attached the Information Assurance Security Guidelines on Data Separation as approved by the Council Security Committee on 18 July 2014.

This page intentionally left blank

**IA Security Guidelines on Data Separation**  
**IASG 5-07**

## TABLE OF CONTENTS

I.	PURPOSE AND SCOPE .....	5
II.	BASIC PRINCIPLE .....	8
III.	SEPARATION OF DATA .....	9
	III.1 Data at Rest .....	9
	III.2 Data in Use - Memory .....	12
IV.	VIRTUALISATION AND CONSOLIDATION .....	12
	ANNEX - SEPARATION OF DATA ON MOBILE DEVICES .....	14
	TERMS AND DEFINITIONS .....	15

## **I. PURPOSE AND SCOPE**

1. These guidelines, agreed by the Council Security Committee in accordance with Article 6(2) of the Council Security Rules (hereinafter 'CSR'), are designed to support implementation of the CSR.
2. These guidelines describe minimum standards to be observed for the purpose of data separation methods in communication and information systems for protecting EU classified information (EUCI) in terms of confidentiality, integrity, availability and, where appropriate, authenticity and non-repudiation, especially in consolidated systems where several CIS share common components. This document covers mainly the use of technical data separation methods inside common hardware used by one or different CIS or between different hardware components inside a contiguous secure area.
3. The Council and the General Secretariat of the Council (GSC) will apply these security guidelines in their structures and communication and information systems (CIS).
4. Member States should use security guidelines as a benchmark when EUCI is handled in national structures, including in national CIS.
5. EU agencies and bodies established under Title V, Chapter 2, of the TEU, Europol and Eurojust should use these security guidelines as a reference for implementing security rules in their own structures.
6. These guidelines should be implemented by the teams responsible for designing and supporting CIS, taking into account the volume and nature of the information being handled by the system in question and the estimated protection needs for the information and the CIS.

7. While access control guidelines<sup>1</sup> cover the requirements for providing access to data or services of CIS<sup>2</sup> to authorised clients while preventing unauthorised access, in this document data separation means the use of technical solutions to separate data, to support and enable access control and reduce the risk of unauthorised or unintended data flows which could compromise the security of CIS and /or of the information processed.
8. In Appendix A, the CSR define security modes of operation of CIS such as dedicated mode, system high mode, compartmented mode and multilevel mode. The applicability of these guidelines to each is as follows.
- (a) These guidelines are chiefly intended for use in a set of CIS or CIS operated in "system high" mode, which requires need to know separation even though all clients are cleared<sup>3</sup> to the highest classification level of the information processed. Where multiple CIS are involved, it is assumed that these operate at the same classification level and that the users have the necessary clearances to access all the information on all the CIS, even if they do not necessarily have the need-to-know.
  - (b) They are also applicable to compartmented mode where rules are set to allow clients to access only those parts of a CIS or set of CIS for which they are formally authorised and have a need to access for reading or writing. Where multiple CIS are involved, it is assumed that these operate at the same classification level and that the users have the necessary clearances to access all the information on all the CIS, even if they do not necessarily have the need-to-know.
  - (c) These guidelines do not usually need to be applied where data separation is not required, e.g. when a set of CIS or CIS is operated in "dedicated mode" where all clients have a common clearance level and a common need to know. When it is considered necessary to provide methods of data separation even in systems operated in 'dedicated mode', these guidelines must be used. Where multiple CIS are involved, it is

---

<sup>1</sup> IASG 5-04, doc. 17547/13

<sup>2</sup> The term CIS covers the set of hardware and software components which runs applications (customised layered products or custom code) providing services to clients. CIS are housed in core and remote hardware, interconnected using network devices and cabling to components such as workstations and/or to other CIS. CIS components are located inside facilities - administrative or secured areas as required by the classification level of the EUCI handled.

<sup>3</sup> In case the client is a CIS or a service thereof, it must be accredited to the required classification level.

assumed that these operate at the same classification level and that the users have the necessary clearances to access all the information on all the CIS, even if they do not necessarily have the need-to-know for information on all the CIS.

- (d) The subject of multilevel security is not intended to be covered by these guidelines as at present data separation solutions of sufficient capability and maturity for this purpose have not been approved for processing of classified information. Separate guidelines may be produced when such approval becomes available.

9. These guidelines are intended to be used where there is a need for:

- (a) separation of data inside consolidated systems;
- (b) separation of data in transit inside secured areas by need to know;
- (c) separation of data based on security requirements and criteria other than confidentiality;
- (d) separation of data at rest (which is not in transit between CIS or components thereof);  
and
- (e) separation of data using means other than approved encryption products.

10. While an organisation would want to use separate hardware and cabling for CIS handling different kinds of information it processes, based not only by classification level but also by the "least privilege" and "need to know" principle, such complete segregation may not be justified by the security requirements of the CIS. This document provides guidelines on data separation in consolidated<sup>4</sup> CIS hardware such as virtual machine (VM) parks and computer clusters, and on 'fixed' or removable storage devices and media respectively.

11. These guidelines apply to CIS or a set of consolidated CIS sharing common components, all located inside an administrative or secured area<sup>5</sup>. Such a set of consolidated CIS is treated as a set of separate CIS and interconnections between them, located inside a secured data centre.

---

<sup>4</sup> In this context the term consolidation means using a common hardware set-up to serve multiple CIS.

<sup>5</sup> If the information handled is not classified higher than R-UE/EU-R, the hardware can be placed in an administrative area.

## II. BASIC PRINCIPLE

12. Data separation is performed using hardware, firmware, software<sup>6</sup> or a combination thereof, which throughout this document are referred to as a "data separation component" (DSC). This DSC must be trusted and approved to the highest level of classification of the data being processed by the set of CIS or CIS<sup>7</sup>. This ensures that the DSC can handle all the information it processes and is trusted by all participating CIS or components of CIS.
13. Data handled by CIS or a particular service of one CIS is made wholly or partly accessible to clients by the DSC, depending on the need to know and clearance level of the client. This is the logical equivalent to the processes of a communication centre for physical forms of classified documents, where the staff of the centre act as DSC. They are cleared to the highest level of classification of the documents handled, are trusted to handle all documents and tasked to separate them according to classification level and need to know before distributing them by secure means to users with the required personal security clearance level and need to know.
14. Irrespective of the method chosen to achieve data separation, risk analysis must be performed to assess the likelihood of cross-contamination or leakage of data, its impact and the likelihood that such compromise is promptly discovered thus enabling the incident to be easily contained. Based on this analysis the chosen solution must be implemented and accredited, residual risk being ultimately accepted by the Business Owner(s) of the set of CIS or CIS in question.
15. Where data separation is between a set of CIS or on consolidated components shared by multiple CIS, then these guidelines must be read in conjunction with the Security Policy on Network Defence (IASP 4), the Security Policy on Interconnection (IASP 3) and the Security Guidelines that support these two policy documents.

---

<sup>6</sup> e.g. RADIUS (Remote Authentication Dial In User Service), TACACS (Terminal Access Controller Access Control System), EAP (and variants) Extensible Authentication Protocol, LDAP Lightweight Directory Access Protocol (e.g. Active Directory AD) in combination with access control lists, firewall rules etc.)

<sup>7</sup> See however paragraph 37



### III. SEPARATION OF DATA

#### III.1 Data at Rest

16. Unless otherwise specified, this section deals with separation of data at rest handled within consolidated CIS inside a secured area.
17. Data at rest is typically contained in databases or files stored on fixed or removable storage media both at the transmitting and receiving end of components of a CIS or between CIS. While the external boundaries of CIS are often easy to determine and separation there covered by the Information Assurance Security Policy on Interconnection (IASP 3), the Information Assurance Guidelines on Boundary Protection Services (IASG 3-02) the Security Policy on Network Defence (IASP 4) and by their associated Security Guidelines, internal boundaries<sup>8</sup> in consolidated or virtualised systems are not easy to identify and data separation approaches at such interfaces need to be defined.
18. In any case, it is strongly recommended to encrypt any sensitive information on removable storage media or mobile devices with storage capability which could leave the secure physical facility or environment in which the data is processed or generated. If encryption products which are not approved are used for EUCI on removable devices or media being carried outside secured areas, the media or device must be treated as if it were a document of the highest classification level of the contained information , that is as if the information were not encrypted at all.

#### *Data at Rest of Different Classification Level*

19. The risk assessment process should determine whether logical access to stored information inside a set of CIS or CIS needs or does not need to be separated by mandatory access control measures or cryptographic products approved to the corresponding level of classification.
20. Placing data at rest of different levels of classification in different storage devices should be the goal, even when these "devices" are logical 'volumes' rather than different physical disks or disk arrays. Such devices can also be encrypted "containers" or "virtual disks", access to which is controlled by key files, passcodes/passwords etc. Access to such storage devices

---

<sup>8</sup> e.g. inside a server, cluster or virtual machine farm

must be controlled by access control measures which allow their mounting only in CIS accredited for handling information of the appropriate classification level or higher and where clients have a need to know for the stored information.

21. It is often desired not to store different "sanitised" versions of the same set of information but to regulate access to the information based on authorisation and need to know. Such differentiated access to stored information is analogous to the internal transmission of portions of classified documents to users with different levels of authorisation and need to know.
22. Guards and gateways between CIS must be customised to allow clients with the appropriate need to know and authorisation to see all relevant parts of electronic documents while masking or blocking the view of entire documents or portions thereof for which the client has either insufficient need to know or insufficient authorisation.
23. By analogy to the considerations about transmission of EUCI detailed above, where risk management allows, a DSC which is a "guard" or "gateway" but neither a diode nor a cryptographic product could be used in some cases for sets of CIS inside common hardware and/or contiguous secured areas to provide role-based access control (RBAC) or discretionary access control (DAC) to stored data on storage device hardware<sup>9</sup> common to members of sets of CIS.
24. When deployed on 'internal' interfaces of clustered or virtualised environments, this implies that the data flows out of the consolidated environment from one CIS to the DSC and thereafter returns to another CIS. Internal filtering, where the DSC resides inside the same consolidated environment as the CIS between which it is filtering, may not be used due to the danger of the corruption and malfunction of the DSC if the management interface of such pooled resources is compromised.
25. Depending on the risk resulting from potential malicious or inadvertent human error (tags being incorrectly added), when used as DSC such guards or gateways should be supplemented by other data separation methods such as suitably customised "data leak prevention" solutions.

---

<sup>9</sup> e.g. arrays of tape devices, optical disks, fixed disks or solid state disks

## Different Sets of Data of the Same Classification Level

26. Even inside a secure device or secured area, data separation should preferably be provided using methods<sup>10</sup> different from the built-in access control functionality of operating systems and networks<sup>11</sup>. The reason is that in many modern operating systems, there is no technical barrier to prevent privileged<sup>12</sup> users from performing tasks for which they have no authorisation.
27. Network segmentation using switches (VLANs) and routers (subnets) should in any case be used extensively to group components and CIS with identical security requirements and segregate them from others with different needs. In combination with security and access control mechanisms of firewalls<sup>13</sup>, this can provide data separation inside CIS or sets of CIS.
28. While it is possible to manually create network segments by configuring the individual devices, such segmentation is better set-up and managed via a network management interface. This management interface and its supporting hardware must be configured and maintained, subject to the restrictions referred to in paragraph 36.
29. Measures need to be put in place and monitored to ensure that only authorised and trained persons can access the management system. All changes to the management interface itself, the management ports on the managed devices, and the configuration set on the devices must be recorded as not all such 'management interfaces' provide the required level of security capability and maturity. Special attention must also be paid to access to "mirror ports" on network switches since at such ports, the entire traffic handled by different ports and VLANs of the switch can be inspected and intercepted.
30. It is also important to retain a record of all configurations over time of the hardware details and of the software versions and settings used for data separation, rather than keeping only a record of changes and of the original set-up.

---

<sup>10</sup> e.g. VPN (Virtual Private Network), PKI (Public Key Infrastructure) using tokens carrying digital certificates for identification.

<sup>11</sup> e.g. access control lists, firewall rules

<sup>12</sup> Users which can modify system settings and/or affect the work of other users are considered privileged.

<sup>13</sup> or similar functionality provided by combined 'switch/router/firewall' gateways

31. After obtaining network access, a client should be forced to authenticate to access each of the various services and/or sets of data which the client is authorised to access.

### **III.2 Data in Use - Memory**

32. For data held in memory, separation and access control are provided by the memory management functionality of operating system(s) and/or by the hypervisor of virtual environments - the assignment of memory to individual processes or virtual machines and prevention of access to such sections of memory by other processes or machines. Depending on the security requirements the CIS or set of CIS is expected to meet, evaluation and testing for the effectivity of such built-in mechanisms must be carried out prior to adopting such solutions.
33. Though physical memory (RAM, etc.) is volatile and information in virtual memory such as swap and page files is frequently overwritten, information can remain in virtual memory long after it is needed. The use of a common pagefile or swap space by different virtual devices must be prohibited by the operating system or hypervisor or both, either by default or as a result of configuration settings. Further, page and swap files must as a rule be regularly wiped clean<sup>14</sup> at times of low usage of the CIS or during planned maintenance outages. Hibernation files must also be prohibited as they are a memory dump of the system at the time of hibernation, so that passwords and sensitive data can be stored in them. Should there be any doubt about the reliability of the separation of data via swap and page files, whenever this is technically possible they should be removed and sufficient physical memory made available to cover the needs of processing.

## **IV. VIRTUALISATION AND CONSOLIDATION**

34. Having a set of different CIS and/or services provided by a CIS hosted on one or more physical or virtual machines governed by the same management system is a feasible and cost-effective method to implement need to know separation for multiple CIS, especially if they have identical or very similar user populations, security requirements, etc. Sharing of resources such as storage, virtual memory, network interfaces, etc. without mechanisms for

---

<sup>14</sup> Overwritten with known characters such as 'hex null' (zeroisation) or random data.

data separation between virtual machines must as a rule be prohibited. Pooled services such as "cloud" or "software as a service" solutions must be subject to careful scrutiny and constant security evaluation prior to EUCI being handled in such environments.

35. With the current increasing use of virtualisation of servers and services, each and every device in a virtualised or consolidated environment must implement the set of technical security measures - the 'security policy' - for protecting the data which it handles, independently of an overarching 'security policy' which governs the entire array and the interactions between the various components.
36. Data separation methods as well as boundary protection and monitoring products such as intrusion detection sensors, content filters, malware filters, etc. deployed to control or monitor data flows for the purposes of network defence may themselves also be consolidated, but the same VM park or cluster may not host both the data separation and/or monitoring systems as well as the CIS separated by them, as also mentioned in paragraph 24, 28 and 37.
37. In all consolidated environments, care must be taken to protect whatever management system (e.g. the hypervisor manager of the VM park) is used to configure the merged systems as it is a single point of failure and a prime target for security attacks and incidents. It is strongly recommended to classify the management system at the same level but to separate it using physical and/or logical separation from the managed systems themselves, restricting access to the management system to a select few of trained experts of high integrity, e.g. via a management network.
38. As in standalone CIS, provision must be made for inspecting and analysing data flows between CIS components and even more so between different CIS for the purpose of network defence. It is therefore not allowed to build consolidated systems where all the data flows within and between CIS are internal and where there is neither a method to intercept and block unauthorised data flows, nor to detect inadvertent or malicious input which could disrupt the correct functioning of the CIS or compromise the security of the information handled. In practice this implies that the processed data must exit the consolidated hardware which performs the "business" functions of the various CIS to different hardware which performs network defence and boundary protection external to the (consolidated) CIS. The network defence and boundary protection hardware may itself be consolidated as long as this is performed in hardware different from the one housing the CIS being protected, as mentioned in paragraph 36.

## ANNEX - SEPARATION OF DATA ON MOBILE DEVICES

39. This section deals only with active devices - devices which can store and process information such as laptops, smartphones, tablets, etc.
40. Besides the general requirement to protect any and all data stored on removable storage media by suitable solutions which must be accredited for the purpose, special considerations come into play when an active mobile device is used to connect to classified systems in order to input or extract data from them. It is vital to ensure that the mobile device is set up in a way<sup>15</sup> which prevents malware taking over the device at all times.
41. As described in paragraph 20, data at rest needs to be separated on mobile devices which house or need to access CIS with different requirements for any domain of security<sup>16</sup>. This should be achieved by setting up different virtual machines on the common hardware, each virtual machine having a different data storage device and/or using different encryption methods to separate the data of one CIS from the others and to protect each collection of data from unauthorised access.
42. The use of different built-in or removable storage media, each having a different security profile and access control restrictions, is also a possible solution for such purposes.
43. When the separation is to be achieved using a hypervisor as DSC between different virtual machines, Type 1<sup>17</sup> ("native" "host-less" or "bare metal") hypervisors should be used, rather than a Type 2 hypervisor running on a host operating system.
44. The method of handling, carrying and using such devices must be carefully evaluated, defined and approved as a prerequisite for the CIS to be eligible for accreditation.

---

<sup>15</sup> e.g. using read only boot media, disabling unneeded network interfaces, removing unneeded services, using multiple network defence products.

<sup>16</sup> confidentiality, integrity, availability, etc.

<sup>17</sup> Type 1 hypervisors run on the firmware of a computer and do not require a host operating system.

## TERMS AND DEFINITIONS

Client	<p>A client can be a physical person, another CIS or another "service" inside a CIS.</p> <p>A client accesses services offered by a CIS to input, process or retrieve data.</p> <p>A client is sometimes termed the "subject" e.g. in the context of public key infrastructure (PKI).</p>
Components	hardware, software or cabling used to build the CIS.
Computer cluster	<p>A computer cluster consists of a set of loosely connected computers called nodes.</p> <p>The nodes share a high speed internal connection (coaxial cable connector or high speed SCSI [small computer standard interface] are the most common) so that in many respects the set of nodes can be viewed as a single system. They usually run the same operating system and various nodes in the cluster can take over processing and network loads if another node fails.</p>
Consolidated systems	<p>Consolidated systems are usually either computer clusters or virtual machine parks (farms) Consolidated systems share storage media and can be reconfigured to shift processing and transmission loads from one machine in the consolidated system to another.</p>
Data	The structured representation of information (in electronic communication and computer systems).

Diodes	Diodes ensure the unidirectional flow of data packets between networks. They can thus be used to block the flow of information from the higher to the lower classified level. They intercept the normal process of establishing connections in network protocols such as TCP/IP, which require both sides of the connection to exchange data packets for "handshaking" required to set up, maintain and tear down a network connection.
Facilities	Buildings and rooms used to house components of the CIS.
Gateway	Technology (software and hardware) that transforms content, protocol or security information from one format to another to enable interoperability, at a boundary between networks with different security policies.
Guard	Technology (software and hardware) used to control transfer of information at a boundary between networks of different security levels.
Hypervisor	<p>The software which manages the resources of virtual (guest) machines in a virtual machine park.</p> <p>A type 1 or bare metal hypervisor runs directly on the computer hardware. A type 2 hypervisor software is a layered product running inside a host operating system which in turn is running on the hardware common to both the host and the guest machines. The guest machines can have operating systems different from that of the host and different from that of other guest machines.</p>
Layered product	A highly configurable software package which provides one or more services.



Mobile devices "Smart" devices which can be temporarily attached locally to a CIS or remotely over a protected interconnection. The distinction between smart removable storage media and Smart mobile devices is becoming difficult to determine with increasing consolidation of functionality in small mobile devices with huge fixed and/or removable storage capacity (up to 64Gb at time of writing).

Record A record is a piece of information which the business owner has decided is of paramount importance to the work of the entity, e.g. correspondence with other organisations on policy, Curriculum Vitae of leading management figures, legal commitments, contracts, etc.

A record can be of any form - hardcopy, electronic e.g. email or instant messages, web sites, reports resulting from database queries, etc.

Removable storage devices Removable storage devices and media can be magnetic (tape, floppy), optical (CD, DVD) or solid state (USB/SSD) but can also be active devices which have internal storage - smart phones, digital cameras, mp3/4 players etc., when these are used primarily for transport of data.

Service In this document the term service refers to a particular functionality provided to authorised and authenticated clients by all or part of a CIS - e.g. e-mail, instant messaging, database, web, file transfer, enterprise resource planning, document management, authentication and authorisation (LDAP, active directory etc. ).

A service is sometimes termed a "resource" or an "object" e.g. in the context of public key infrastructure (PKI).

The term as used here does not cover built-in memory-resident modules or executables which form part of the operating system or layered product itself (e.g. the Event Log service of Microsoft Windows).

Virtual machine park A virtual (computer) machine park consists of several different "guest" machines (computers or network devices) running under a common "hypervisor" and sharing common hardware.

The virtual machines can run the same or different operating systems.

VLAN A Virtual Local Area Network consists of a group of active devices with a common set of requirements. In network terminology, it is called a broadcast domain. VLANs have the same attributes as physical LAN and can be created by assigned physical ports on a network switch to the VLAN.

Access control lists, defined on the operating system of the switch, or in an external DSC such as a router or firewall, limit the data flows to and from such a VLAN to other VLANs or networks.

Mirror ports Port mirroring in a network switch means that a copy of the network packets seen on one switch port (or an entire VLAN) is sent to a network monitoring connection or to another switch port (the mirror port). Such mirroring is typically used for intrusion detection systems, eavesdropping (passive probes) and debugging problems on the network. Mirror ports are often known as SPAN ports (Switched Port ANalyzer - CISCO), other vendors having different names for them.

