Measurement Frameworks and Metrics for Resilient Networks and Services: Technical report

Discussion Draft

enisa
European Network and Information Security Agency

About ENISA

The European Network and Information Security Agency (ENISA) is an EU agency created to advance the functioning of the internal market. ENISA is a centre of expertise for the European Member States and European institutions in network and information security, giving advice and recommendations and acting as a switchboard of information for good practices. Moreover, the agency facilitates contacts between the European institutions, the Member States and private business and industry actors.

Acknowledgments:

ENISA would like to express its gratitude to the stakeholders that provided input to the survey and the Workshop on Metrics. Our appreciation is also extended to the ENISA experts and Steering Committee members who contributed throughout this activity.

Contact details

For more information about this document, please contact:

Dr. Panagiotis Trimintzios

Panagiotis.Trimintzios@enisa.europa.eu

Internet: http://www.enisa.europa.eu/act/res

Table of Contents

List of figures

BLANK PAGE

Executive summary

Reliable communications networks are becoming increasingly important to our society today. The European Commission has acknowledged this and EU Commission's recent Communication on CIIP[1] recognises the importance of the area and confirms ENISA's role in the field.

ENISA, fully recognizing this need, devised a Multi-annual Thematic Program (MTP) with the ultimate objective to collectively evaluate and improve the resiliency of public communication Network and Services in Europe. As a part of that program, a study was done with a group of ENISA stakeholders on resilience measurements [1] – it became apparent that **lack of a standardised framework or good metrics** was considered to be one of the main challenges experienced by the respondents. Resilience was not considered to be a **well-defined term** and depending on the context, it encompassed several interpretations and viewpoints. Additionally, there was consensus on the fact that **information sharing** and sources of **consolidated information on resilience metrics** were not readily available. These challenges were recognised as serious obstacles towards the adoption of resilience metrics.

Addressing these concerns, this report represents an attempt to create a single technical source of information on resilience metrics, the taxonomies and the open issues. It puts together work that has been done in the areas of security, dependability and specific taxonomy research under the single umbrella of resilience. It is intended to become a source of information for the community interested on resilience and measurements, but also the cause to initiate more in depth works on the subject.

The first section of this document includes the **definitions of a number of important terms** such as 'resilience' and 'metric'. We analyse each part of the resilience definition and explain its impact on metrics and measurements.

In section 3 we **overview the different initiatives, works and frameworks related to resilience metrics and measurements**. We look at regulations, research efforts and a number of most of the related taxonomies available in the literature.

The report then continues with section 4 by presenting a two-dimensional approach **to categorising resilience metrics.** This section is a first attempt to bring together different taxonomies in single unified model. The model includes an incident- and a domain-/discipline-based dimension. Finally, this section briefly explains the **current open issues** when trying to apply these metrics on a larger scale.

[1] *http://ec.europa.eu/information_society/policy/nis/strategy/activities/ciip/index_en.htm*

Subsequently, a **number of metrics are identified and presented** in a detailed and consistent way. The aim of the section is to respond to the request of an overview of good set of baseline resilience metrics. While the section in no way claims to be exhaustive, it should provide experts with a starting set of metrics.

Introduction

ENISA devised a Multi-annual Thematic Program (MTP) with the ultimate objective to collectively evaluate and improve the resiliency of public communication Network and Services in Europe. In order to achieve the desired resilience of the involved networks, measurements are needed and expected.

There are several specific and commonly recognised needs and drivers for adopting resilience (and security) metrics and frameworks, like for instance:

- The need to show and provide **assurance and evidence** on the level of resilience and/or security achieved;
- The need of a metrics system for **validating the conformance with regulations**, policies and business requirements;
- The **practical need to analyse** in an effective and efficient manner the increasing number and complexity of technical logs;
- The **identification of trends** in the different communications networks, such as the level of attacks, common failure causes, etc.

As a part of that program, a study was done with a group of ENISA stakeholders on resilience measurements [1]. As was anticipated by ENISA, it became apparent that the stakeholders are concerned with the many challenges in the measurement of resilience.

In an attempt to alleviate these concerns, this report represents a first discussion draft on resilience metrics, in an attempt to provide a holistic view on resilience. It is the result of thorough study and puts together work that has been done in the areas of security, dependability and specific taxonomy research under the single umbrella of resilience. It is intended for technical experts on resilience and measurements thereof.

The main challenges identified were [1]:

- Resilience was not considered to be a **well-defined term** and depending on the context, it encompassed several interpretations and viewpoints.
  - Section 0 presents a definition of resilience to provide a common platform for future discussion on resilience.
- **A lack of a standardised framework or good metrics**. Organisations have their own specific approaches and means of measuring resilience, if they actually have any at all. There was acknowledgement that meaningful examples of metrics categories can contribute to a systematic and comprehensive practical approach when metrics need to be considered.

  - Section 0 provides definitions for and indicates the differences between a metric, a measurement and an indicator. It describes different aspects of measurement and tries to create awareness of different taxonomies described in the research literature.
  - Section 0 presents the anatomy of selected taxonomies and the rationale behind these taxonomies. Two example taxonomies are presented in more detail.
- Major hurdles in the **identification and implementation of adequate metrics** or measurement frameworks, either because the metrics do not exist or because the organisations are unaware of their existence. In general, maturity of current practices is low. The main advice towards measuring resilience in [1] is to base the resilience and security metrics on existing business requirements and to start out with a small set of metrics which gradually expands.
  - Section 0 presents a number of metrics in a hands-on approach. They are presented in a consistent template, specifying the definition of each metric, its usefulness to the measurement of resilience and the method of measurements method.
- Additionally, there was consensus on the fact that **information sharing** and sources of **consolidated information on resilience metrics** were not readily available.
  - Section 0 outlines a number of regulations regarding resilience as well as relevant research projects which have been formulated in the technical literature.

Therefore, the **aims of this report** are to provide **practical guidance** towards the implementation and usage of resilience metrics by:

- Clarifying the **key concepts regarding the resilience** of networks and services, as well as regarding metrics and measurement frameworks in this context;
- Presenting **reference practices** for stakeholders to measure the effectiveness of efforts related to the resilience of communications networks and services, based on the analyses techniques, methods and metrics frameworks currently existing and used by stakeholders;
- Highlighting **applicable regulations and key ongoing research projects** in the context of network service resilience and metrics.

This report is the result of a thorough study and we believe it provides an overview of the different areas of resilience measurements. It should be considered as a first draft for discussion - we would like to encourage the readers to share their comments and suggestions on the report with us.

Definition of network service resilience

When considering network service resilience, it is essential to clarify the key concepts relating the topic.

First of all, the term 'network service' is defined based on 2 out of the 3 security layers as defined in ITU.T X.805 and bears the security plane separation of that recommendation in mind [36]:

A **network service** consists of the infrastructure building blocks of a network (such as individual routers, switches, servers, Ethernet links, etc.) and the services provided to the end-users built on those infrastructure building blocks (such as Frame Relay, IP, Wi-Fi, VoIP, QoS, Location services, etc.). A network service relies on three different types of activities that occur on a network (depicted in Figure 1):

- **End-user:** Access and use of the network by the customers for various purposes (basic connectivity, VPN, VoIP, etc.);
- **Control/signalling:** Activities that enable efficient functioning of the network (such as routing, machine-to-machine communications; ..);
- **Management:** The management and provisioning of network elements, services and applications.

As noted in the definition of a network service, web browsing, e-mail etc. are not considered to be network services but are network-based applications built on top of network services.

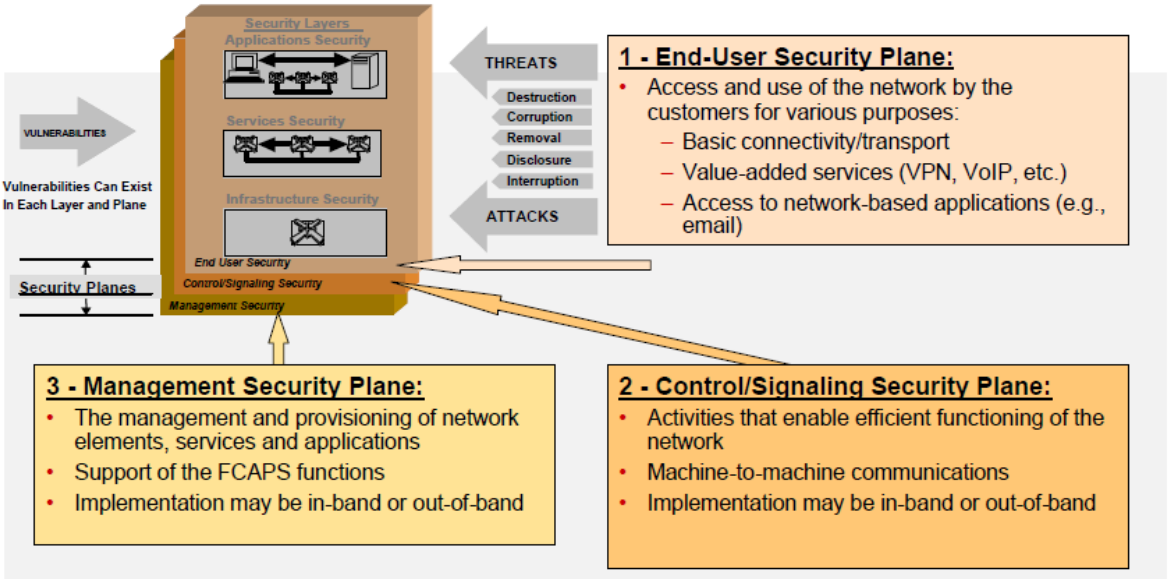

Figure 1: ITU.T X.805 definition of network service (adopted from [36])

A widely used and generally accepted definition of the term 'resilience' in the context of networks and services is the following [35] [38]:

Resilience is the ability of the network to provide and maintain an acceptable level of service in the face of various faults and challenges to normal operation.

This definition of the network service resilience can be decomposed into a number of more tangible elements for additional clarification. More specifically a number of key questions and considerations should be taken into account when looking at resilience:

- Network services should be prepared against faults and challenges by implementing resilience provisions in order to provide and maintain **an acceptable level of service**. Subsequently the following should be considered:
  - What is an acceptable level of service?
  - Which provisions ensure the ability to provide and maintain a level of service?

- **Faults and challenges** (disturbances) will have an impact on the network which needs to be measured in order to verify operational state and service degradation of the network services. Subsequently the following should be considered:
  - What faults and challenges are networks and services facing?
  - What is considered to be normal operations?
  - How the risk and impact of faults and challenges to the network be measured?
  - How will risk be managed for the network service?

These questions are addressed in the following subsections and will serve as a basis to further define a reference measurement framework for network service resilience.

Acceptable level of service

The aim of resilient networks and services is to provide an acceptable level of service (and be able to maintain that level of service) when faults are occurring in the network or the level of service is being put at risk by challenges (for example: the incoming network traffic exceeds the traffic rate the service can handle). Therefore it is fundamental to specify the acceptable or desired level of service and align any measurement practices with such definition.

In the domain of telecommunications and networking, acceptable service levels are typically defined in a Service Level Specification (SLS), often as part of a Service Level Agreement (SLA) between the network service provider and customer. The SLA describes the service levels that are considered to be **acceptable to the customer**. What is considered to be acceptable can also be determined by regulatory requirements and standards set out for the operators (some of these regulatory requirements and standards implicitly target societal acceptance of the network service level).

Network service level agreements and specifications are commonly defined in terms of **quantitative service parameters** such as service availability, throughput (bandwidth), latency (average round trip time), packet loss, jitter (packet delay variation), etc. These availability and service quality elements express whether the network service is actually delivered and can be measured as a function of

time. However, currently no universally accepted taxonomy exists to consistently express levels of different network services at a granular level by means of such properties.

It should be noted that acceptable service level definitions can also be refined based on further classification of the **service disruption impact**. More specifically, the significance of the service impact can be quantified using a number of impact metrics such as the extent of the network impacted in terms of users, services or network portions or in terms of recovery times (these metrics will be reviewed in section 0).

The specification of the level of a network service typically consists of defined minimum thresholds for all relevant, quantitative properties of that service.

By monitoring the current service and comparing the measured properties to the defined service level thresholds, one can assess whether the level of service has been met.

Based on defined target and minimum service level thresholds, resilience can be defined as a function of service level in the face of faults and challenges as illustrated in Figure 1 below. The property of service operating above or at the target service level can be defined as an **acceptable level of service**. When the network service operates between the target and the minimum service levels, the performance represents an **impaired level of service**. Finally the performance of a service operating below the minimum service level would be at an **unacceptable level of service**.

Network service levels depend on the operational parameters of the network that supports the service. More specifically, the faults and challenges the network faces have an impact on the service level perceived by the service consumer. Therefore, the objective of pursuing network service resilience is to lower the impact of operational network parameter degradation on the network service parameters that will express the final service level delivered.

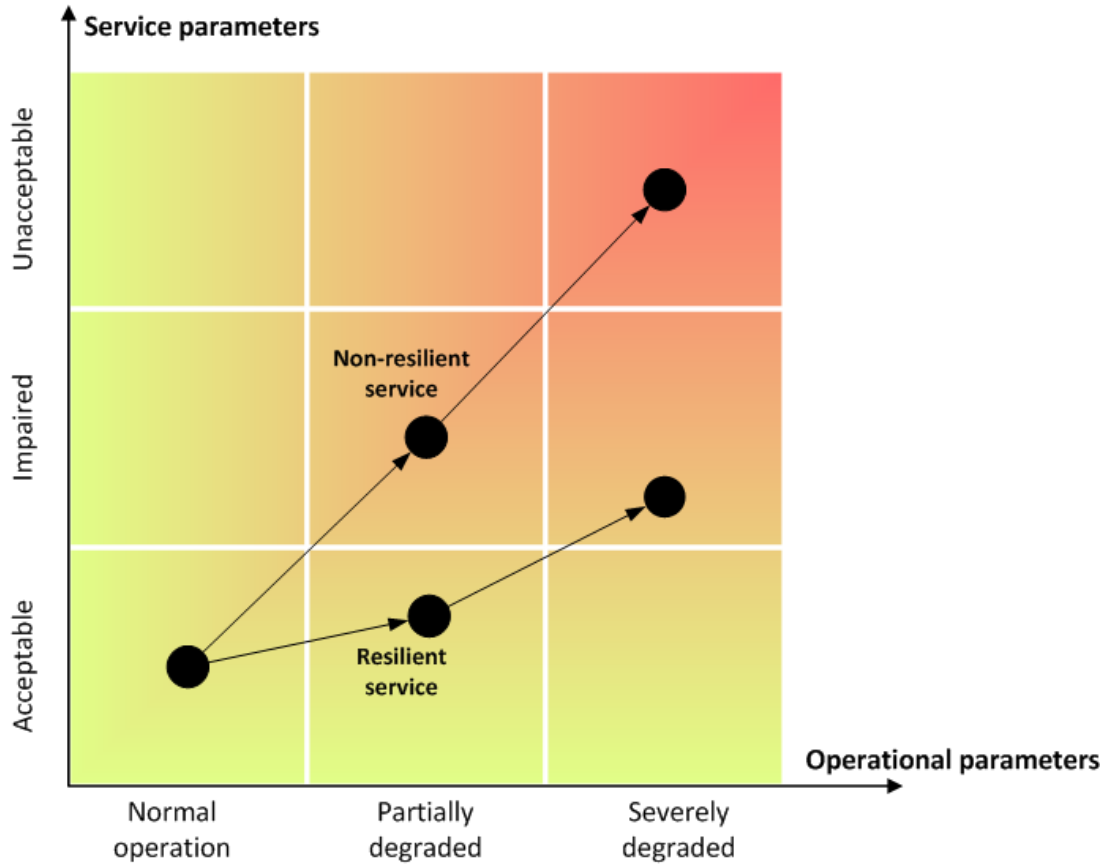

Figure 2: Resilient vs. non-resilient service

In fact Figure 2 is a visual representation of the definition of resilience as given in section 0. Facing the same operational parameter degradation (these are the faults and challenges the network is facing), the service level of a resilient service will have less degradation compared to a non-resilient service, where the non-resilient service will reach the impaired service level with less degradation in operational parameters compared to the resilient service in face of challenges and disturbances to normal operation. Reversing the previous sentence, one could also say that to remain with the acceptable level of service range, a resilient service can tolerate more operational degradation than the non-resilient service.

The service level agreement (SLA) with customers or other parties really determine the level of network service resilience which will be built into the network (the SLA defines, among others, network service parameters, such as a maximum guaranteed delay or a minimum guaranteed bandwidth). The SLA also requires proper **business continuity management** to ensure that the network service is delivered to the service consumer according to the SLA parameters, even when facing network faults.

As resilience-enabling measures bear a cost to the organisation, organisations should always base the resilience requirements on the set of service level agreements that must be met (and are, in many cases, contractually agreed upon). Taking into account that no network is infallible, the SLA can be seen as a driver for business continuity requirements which in turn will drive the resilience

requirements. In section 0 we propose a number of metrics to quantify the resilience of a network service. Example thresholds that could define an acceptable level of service are also provided.

Provide and maintain

In the definition of resilience in section 1.1 above, the term '**provide'** represents the delivery of the network service on an acceptable level (as specified in the service level agreement) given normal operational parameters.

The term **'maintain'** represents ensuring that the network will provide service in the normal condition for a maximum fraction of operating time, particularly in the light of faults and challenges. It refers to the goal of delivering an acceptable or highest possible network service level, by taking measures to prevent challenges, minimizing their possible service impact, and rapidly restoring the network service level in case it was degraded.

Faults and challenges

Faults and challenges to normal operation will also be referred to as disturbances or risks to the communication networks and services.

While faults represent errors and/or failure in the different subsystems that support a network or a service, different categories of challenges can also threaten the service level of a specific service.

Examples of challenges are (based on [38]):

- Unintentional misconfiguration or operational mistakes: Non-malicious errors made by humans, e.g. in the configuration of network components;
- Large scale disasters (natural and human-caused)
- Malicious attacks from intelligent adversaries, e.g. denial of service attacks or hacking of network systems;
- Hardware destruction: Destruction of physical components due hardware failure;
- Events or situations where a large surge of legitimate traffic is observed;
- Failure of a service provider: Outages to other services affecting the network, e.g. a loss of the connections to the Internet due the WAN (Wide Area Network) provider;

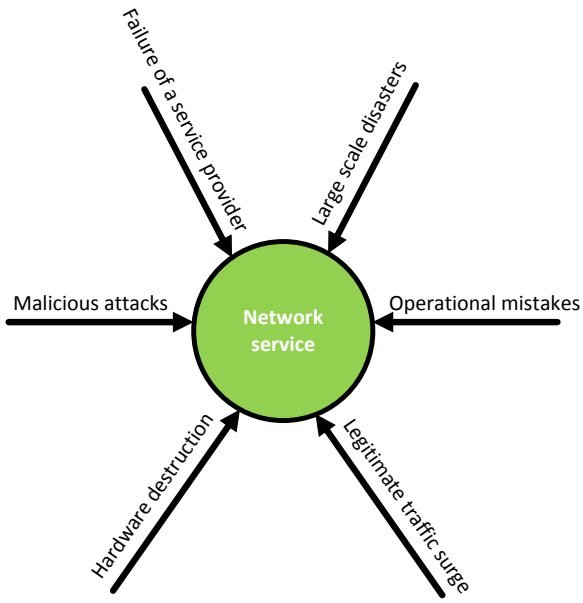

Figure 3: Network challenges

Managing the risk of faults and challenges

The impact of the loss of critical services on its stakeholders is the same regardless of the cause of the disruption. It is of vital importance for the organisation to take appropriate risk management measures and implement controls to ensure that their networks and network services are resilient enough to ensure optimal provision of services in the face of faults and challenges.

An organisation should manage **the risk of the network services** by:

- Determining the risk management objectives:
  - **Reduced failure probabilities** – the reduced likelihood of faults and challenges to a critical infrastructure, systems and components;
  - **Reduced consequences from failures** – in terms of service disruption, damage (including financial damage) and negative economic and social impacts;
  - **Reduced time to recovery (TTR**) – the time required to restore service to a normal level.
- Determining an appropriate **risk methodology**;
  - The level of risk an organisation is willing to accept depends on the **risk appetite** of the organisation.
- **Identifying the different risks** in providing network services;
- If possible, **quantifying** the risks (using for example the metrics of section 0);
- Identifying **likelihood** of those different risks and the possible **impact** on the network service if a disruption occurs;

Choose **appropriate controls or mitigations** (or accept the risk). In order to manage the risk associated with disturbances, it is fundamental to define how to measure the impact of disturbances in the service offered to the organisation, the stakeholders and the users. It must be noted that not all risks are measurable: they should be qualitatively assessed instead of measured.

The metrics chosen by an organisation to quantify the impact of disturbances on the organisation, stakeholders and users will depend mainly on what metric is most relevant to the organisation's goals, priorities, and business. Combined with risk likelihood figures, they can be used to quantify the risks of these disturbances.

Impact measures can also be incorporated in service level specifications in order to provide more fine-grained control over the specification. For example, an acceptable level of service for consumer Internet connections could contain the following thresholds:

- 95% of all the service users have an availability of 99,99% measured on a yearly basis. Assuming a population of 1 million service users, this implies that the service provider must provide 950.000 users with Internet access that can go below the acceptable service level for 0,87 hours yearly for each user.
- 99% of all the service users have an availability of 99,99% measured on a yearly basis. Assuming a population of 1 million service users, this implies that the service provider must

provide 990.000 users with Internet access that can go below the acceptable service level for 0,87 hours yearly for each user.

The example clearly shows that availability guarantees are not the only factor the service provider will need to take into account: In order to provide service guarantees to 99% of all service consumers (in this example: 40.000 extra users), the provider will need to foresee extra measures in the network to avoid failures.

Whether or not these fine-grained specifications are an attainable and measurable quantity depends largely on the type of service. For example, while a cellular provider can approximate the number of service users by the region that is out of service, it is very difficult for a housing provider of an externally facing web application to foresee how many users may or may not experience degraded service in case of an outage.

Measurement frameworks and metrics - discussion

It is a widely accepted management principle that 'an activity that cannot be measured, cannot be managed', and this also applies to network service resilience. Metrics can be a very effective tool for network / security / resilience managers and engineers to assess and manage the effectiveness of their resilience policies and controls.

The importance of this principle is also supported by the reliance on measurements in the typical continuous improvement cycles that can be found in international security and management standards and good practices, such as for example ISO/IEC 27001:2005 ISMS, Six Sigma, and various quality management system approaches. Measuring the effectiveness of resilience policies and controls put in place by organisations is a challenging issue as the discipline is still in the early stages of development. Many organisations do not use the concept of resilience, although their policies, procedures and controls refer to information security, availability and similar concepts.

An activity cannot be managed, if it cannot be measured.

Metrics and measures

It is essential to clearly define certain concepts of metrology. In popular literature on the subject, there are some contradictions in the various terminologies used. For example, some reports do not make a distinction between measures and metrics; others may talk about qualitative metrics. Throughout this report, several key definitions, concepts and relationships have been used from the IATAC document on Measuring Cyber Security and Information Assurance [11]:

- A **measurement** is the act or the process of measuring, where the value of a quantitative variable in comparison to a (standard) unit of measurement is determined.
  A measure is a variable to which a value is assigned as a result of the measurement.
  According to the Webster dictionary; a measure represents the dimensions, capacity, or amount of something ascertained by measuring. For example, seven seconds (7 sec) is a value of duration.

- A **metric** is a system of related measuring enabling quantification of some characteristic of a system, component or process. A metric is composed of two or more measures. For example, the number of information security incidents per day is a security metric. The metric represents incident rate, which is related to the "security" attribute of a system, in function of time. The composing measures are 3 incidents and 1 day.

  Note that according to this strict definition, metrics are always quantifiable, qualitative measures cannot be used as metrics. However in various information technology metrics research initiatives, qualitative measures are also used as metrics, due to the fact that certain aspects of information technology are not easily quantified.

- An **indicator** differs from a metric in the sense that the value of an indicator is calculated and not measured. It is also not time-dependent. An indicator is a quantified property that does not require measurement: it is a calculated property. An example is the redundancy of a network path between 2 nodes: It is calculated from the topology and does not need to be measured nor is it time-dependent (unless the topology is changed of course).

At the moment of writing, several research projects on the subject of network service resilience exist. However, a standard taxonomy or framework for resilience measurement has not yet been globally defined and accepted. Every initiative uses its own framework and categorization of resilience concepts, controls, etc. A full measurement framework with regards to network and service resilience is non-existent to date. Academic projects such as AMBER or ResumeNet (refer to section 0 and 0) are performing research on the subject but have not yet publicly delivered practical metrics or an integral measurement framework.

Several information security and information technology performance management measurement frameworks define initial guidance regarding resilience and security performance metrics. In annex of this document, we list the relevant information security performance management measurement frameworks that have served as inspiration to our approach for the development of a measurement framework for resilient networks and services.

During the development and implementation of a measurement program, several key principles need to be considered in order to define a measurement framework that is based on good metrics:

- **Technical characteristics**: Good metrics meet the following characteristics in order to allow for accurate and useable comparison of different measurements:
  - **Quantifiable**: Metrics are per definition based on quantitative measurements; as a result any metric should be quantifiable. The measurement method will define the quantity and unit of measurement associated to a metric;
  - **Repeatable**: The measurement methods for the metric must be reproducible. This means they must yield the same result when repeated by a second assessment;

  - **Comparable**: The measurement values of a metric must be a linearly ordered set, such that any 2 measure values can be compared to each other and a conclusive comparison result be reached.
- **Business characteristics**: In addition to technical characteristics, good metrics must also possess certain non-technical characteristics in order to be used and to be useful to the business objective under measurement.
  - **Easily obtainable**: Metrics must facilitate easily obtainable measures. If not, the complexity of measurement will outweigh the advantage and purpose of measurement by utilizing too many resources or not being timely;
  - **Relevant**: The metrics must provide useful information for tracking performance, managing resources and directing the strategy towards the business objective. They must be relevant to the mission and provide added value. If not, metrics will only be perceived as overhead and will not be used.
  - **Continuous improvement**: Metrics should be used to monitor and improve resilience on a continuous basis. While metrics can provide benchmarks against target values for one moment in time, they should be also used to track the resilience improvement process.

Metrics taxonomies overview

In publicly available information sources (research reports, white papers, standards etc) several types of information technology metrics classifications have already been proposed and suggested. However most of these metrics taxonomies are specific to general information security or a certain area of information technology services, for example software development. At the time of writing, no publicly available taxonomies have been defined for communication networks and services resilience.

The most common forms of information security metric classification are according to:

- **Information security objectives:**
  - Metrics are grouped according to security control objectives such as the objectives defined by the ISO/IEC 27001:2005 standard:
    - Security Policy
    - Organizing Information Security
    - Communications and Operations Management
    - Information Security Incident Management
    - Etc.
  - Metrics can also be grouped according to their business functions as defined in the CIS Security Metrics report [5]:
    - Incident Management
    - Vulnerability Management
    - Patch Management
    - Application Security
    - Configuration Management

- Financial Metrics

- **Organisational aspects:** organisational, operational and technical metrics (e.g. WISSSR Measures structure [13], I3P Taxonomy of Security Metrics for Process Control Systems [34]);

- **Network properties:** A theoretical poster on quantifying metrics for resilient and survivable networks **[34]** provides the following metrics taxonomy:
  - Density (Number of nodes, Area of spread, Distribution pattern, Topology change rate
  - Mobility (Velocity of the node, mobility model, predictability)
  - Channel (Capacity distribution, propagation model, bit error rate, error rate model)
  - Node resources (Electrical power, computing power, memory, TX/RX power, location awareness)
  - Network traffic (Distribution, packet size, source/sink placement, Quality of Service)
  - Derived properties (Degree of connectivity, propagation delay, queuing delay, node willingness)

- **Measurement aspects or measure type**: In certain cases metrics are also grouped according to an aspect of measurement or a type of measure. For example, NIST Special Publication 800-55 Revision 1 "Performance Measurement Guide for Information Security" **[6]** defines 3 types of metrics:
  - **Implementation**: used to demonstrate progress in implementing information security programs, specific security controls, and associated policies and procedures;
  - **Effectiveness/efficiency**: used to monitor if program-level processes and system-level security controls are implemented correctly, operating as intended, and meeting the desired outcome;
  - **Impact**: used to articulate the impact of information security on an organisation's mission.

While the taxonomies listed above provide a broad overview, many of these are specifically focused on the security aspects of systems and not as such on specific metrics for measuring resilience in networks. The IATAC document on Measuring Cyber Security and Information Assurance [11] provides an excellent overview on the cited taxonomies. A summary of the most relevant taxonomies has been consolidated in section 0 of this document.

In section 0, we propose a metrics taxonomy for networks and services resilience that is based on the objectives of network service resilience and focuses on practical applicability instead of academic interests. The taxonomy quoted in the CIS document is used as our primary reference [5].

Aspects of measurement

Many standards exist on the different generic properties that define the measure. They can describe the nature of the metrics, the difference between quantative and qualitative measures, the

difference between intrinsic and relative measures, etc. In the case of NIST Special Publication 800-55 Revision 1, these properties are used for a classification of metrics.

The following are aspects of measurement that have been discussed mainly in the information security context. Most can also be applied to the resilience context.

- NISTIR 7564 – Directions in Security Metrics Research [7] discusses the following aspects of measurement:
  - Correctness & effectiveness
  - Leading, coincident & lagging indicators
  - Qualitative & quantitative properties
  - Measurement of the large & the small

- In NIST SP 800-55 Rev.1 – Performance Measurement Guide for Information Security **[6]**, metrics are categorized / typed according to the another aspect:
  - Implementation
  - Effectiveness/efficiency
  - Impact metrics

- ISO / IEC 27004:2009 Information Security Management – Measurement also makes a distinction between objective and subjective measurement methods.

- Furthermore additional potential classifications are listed on a NIST website on metrics and measures[2]:
  - Static & dynamic metrics
  - Objective & subjective metrics
  - Intrinsic & relative metrics

Because these properties have been already comprehensibly described in the literature, we will refer to the selected literature instead of providing them explicitly.

[2] *http://samate.nist.gov/index.php/Metrics_and_Measures.html*

Overview of related works

Applicable regulation

Directive 2009/140/EC of the European Parliament and of the Council

Directive 2009/140/EC of the European Parliament and of the Council, chapter IIIa 'Security and integrity of networks and services', article 13a 'Security and integrity', list item number 3, states 'Member States shall ensure that undertakings providing public communications networks or publicly available electronic communications services notify the competent national regulatory authority of a breach of security or loss of integrity that has had a significant impact on the operation of networks or services.'

The relevance of the EC directive paragraph with regards to network service resilience metrics resides in the concept of 'significant impact on the operation of networks or services'. What constitutes a significant impact on the operation of networks or services? The term significant impact is a qualitative expression of the effect an event can have on the operation of networks or services. The interpretation of 'significant' is subject to interpretation and discussion.

However network service resilience metrics can provide a solution here. 'Significant impact on the operation of networks or services' is a concept that can be defined using qualitative or quantitative resilience metrics for networks and services and can provide guidance to Member States in the implementation of Directive 2009/140/EC.

Federal Information Security Management Act (FISMA) and Government Performance Results Act (GPRA)

We also identified a number of laws, regulations, and policies in the US that include compliance verification requirements that mandate the use of measurement for verifying compliance or, at a minimum, suggest the use of or imply a preference for measurement as the best approach to verification of compliance. Key examples of this are relating to the Federal Information Security Management Act (FISMA) and the Government Performance Results Act (GPRA).

FISMA provides a comprehensive framework for securing federal government IT resources by defining key federal government and agency roles and responsibilities, and by requiring agencies to integrate information security into their capital planning and enterprise architecture processes. FISMA requires that agencies conduct annual information security reviews of all programs and systems, and report the results of those reviews. Annual FISMA guidance is published that includes specific performance measures to be reported as a part of annual and quarterly reporting. A requirement was included regarding three performance metrics that agencies need to use to measure the effectiveness or efficiency of security policies and procedures based in NIST SP 800-55 Rev1.

The Government Performance Results Act (GPRA) in the US does not explicitly mandate security planning, measurement, or reporting. However, it is desired that federal government agencies tie all their activities to their strategic and performance planning processes. NIST SP 800-55 Rev. 1 suggests that agencies tie their information security goals and objectives to the overall agency goals and objectives, and that agencies use information security measures to track accomplishment of their information security goals and objectives.

Key network service resilience research projects

Several research initiatives are ongoing around security and resilience metrics. It is clear that there is currently no clarity or alignment regarding exact definitions, taxonomies, applicability, purpose, etc. We highlight below the key research projects and initiatives relevant to the study.

AMBER project[3]

AMBER was a project finished in 2010 aimed to coordinate the study of resilience measuring and benchmarking in computer systems and components, fostering European research in order to address the big challenges on resilience assessment posed by current and forthcoming computer systems and computer-based infrastructures. The AMBER project developed a research agenda for resilience assessment as input for the EU Seventh Framework Programme (FP7) of the Information and Communication Technologies (ICT) research activity.

AMBER brought together leading research teams on assessment, measurement and benchmarking of resilience in computer systems in order to coordinate the effort of defining metrics and benchmarks for comparative evaluation of the resilience of computer systems and components. The consortium included seven partners (universities of Coimbra, Budapest, City, Chalmers, Florence and Newcastle and the company ResilTech) from five EU countries and relies on a large and representative Advisory Board that constitutes the necessary link between the coordination action and the influential parties in industry and government, thus ensuring that the views of major stake-holders are being taken into account by the AMBER Consortium.

The project had several work packages for which the key deliverables included:

- A web portal made of two distinct parts: an intranet accessible only by AMBER partners (using authentication) and an extranet part accessible by the community in general;
- A Data repository to analyse and share field data on computer failures and resilience evaluation experiment results (the goal is to build an infrastructure that integrates data from different sources in such way that enables comparison and cross-exploitation in a meaningful manner);
- State of the art report on resilience assessment methods;
- A research roadmap on assessing, measuring, and benchmarking resilience based on the identification of gaps and research opportunities.

[3] *http://www.amber-project.eu*

In the conclusion of the State of the art report around "Resilience assessment", the AMBER project concludes that research challenges include both pushing the boundary of the problems that can be addressed by quantitative techniques, and finding clearer indicators for these boundaries. They also identified the need for sound guidance on the advantages of measurement "as far as they go" while avoiding a potential collapse into unrealistic, theoretical decision-making.

ResumeNet project[4]

The EU-funded ResumeNet project is currently investigating a framework and mechanisms for resilience in a future Internet. At the centre of the project is a straightforward strategy for building resilient networked systems, called D^2R^2 + DR – Defend, Detect, Remediate, Recover, Diagnose and Refine.

The strategy is as following: initially, one must install appropriate *defensive* measures, e.g., configure firewalls and use appropriate redundancy and diversity of services to ward off anticipated challenges. In many cases, there will be unforeseen events (or those that are too expensive to build defensively for) that will breach defensive measures and cause a degradation of service. Such challenges should be *detected* in real-time and the network dynamically adapted to *remediate* them. This implies an underlying monitoring system. Most likely, there will be a cost associated with remedying a challenge (e.g., sub-optimal paths are used to route around a malicious node); a *recovery* stage in the strategy reflects that mitigation mechanisms should be disengaged when a challenge has abated. It is assumed that the system is not perfect; therefore, the aim is also to *diagnose* shortcomings and *refine* the networked system.

The research work on resilience in the context of ResumeNet is structured around three main directions:

- **Framework:** The framework aspects of resilience are investigated in the first work package (WP1). This is where the main ingredients of a systematic approach towards embedding resilience in future networking are investigated. The identification/classification of challenges, the quantitative assessment of their impact, the search for resilience metrics, as well as the role of policies and cross-layer techniques lie at the core of studies in this WP.
- **Mechanisms:** The realization of the resilience framework studied in WP1 raises certain requirements for the network. The more generic and sometimes network-agnostic structures need to be actually supported and implemented in the real network; (some) network nodes need to be equipped with certain functionality, information about the state of the network has to be collected, shared, and made available to decision-making entities, policies need to be enforced, remediation mechanisms need to be coordinated and synchronized. How this can be optimally accommodated in the network- and service-layer infrastructure is the subject of WP2 and WP3. In the same time, these two WPs

[4] *http://www.resumenet.eu*

accommodate tasks tailoring the more generic solutions devised in these two WPs to the particular study cases treated in more detail in the experimentation work (WP4).

- **Experimentation:** Assessing the efficiency of the framework and the mechanisms supporting it, in particular when this has to be done quantitatively rather than qualitatively is challenging. An exhaustive investigation of the full space of challenges and countermeasures is not possible; many of those anyway are specific to a particular networking context. Therefore, in the experimentation phase of the project (WP4), four study cases have been selected for assessing to what extent the framework studied in WP1 and the mechanisms devised in WP2 and WP3 could drive tailored solutions for improving their resilience. These four study cases were deliberately chosen to address promising networking scenarios, whose wide-spread deployment is to high extent impeded by the lack of resilience: wireless mesh and delay tolerant networks, peer-to-peer voice conferencing and service provision over heterogeneous smart environments.

The second deliverable of work package 1, D1.2a 'Defining metrics for resilient networking', is the most relevant with regards to Measurement Frameworks and Metrics for Resilient Networks and Services. However the deliverable has not been published yet, publication has been postponed to an undisclosed date.

ResiliNets[5]

Society increasingly relies on computer networks in general and the Internet in particular. Therefore, the consequences to disruption of the network are increasingly severe, and threaten the lives of individuals, the financial health of business, and the economic stability and security of nations and the world. Resilience and survivability are therefore regarded as critical to the future of our network infrastructure. The ResiliNets initiative aims to understand and progress the state of resilience and survivability in computer networks, including the Global Internet, PSTN, SCADA networks, mobile ad-hoc networks, and sensor networks.

The ResiliNets initiative is a collaboration between the University of Kansas (US) and Lancaster University (UK), and aims to understand and progress the state of resilience and survivability in computer networks, including the Global Internet, PSTN, SCADA networks, mobile ad-hoc networks, and sensor networks. The initiative provides a wiki that is designed to facilitate collaboration and provide the content for the ResiliNets portals.

Other European projects

We noted several key European projects regarding resilience in the context of the EU Framework Programmes on Information and Communication Technologies (ICT) research activity including:

[5] *https://wiki.ittc.ku.edu/resilinets/Main_Page*

1.1.1.1 DESEREC - Dependability and Security by Enhanced Reconfigurability[6]

The main interest of the proposed DESEREC approach is to improve the dependability by the combination of three technologies: modelling & simulation, incident detection, and response. The work in DESEREC is highly model-driven, including a vulnerabilities and fault model, thus providing a methodology that allows for the assessment of both security and dependability.

1.1.1.2 HIDENETS - Highly Dependable IP-based Networks and Services[7]

HIDENETS provides end-to-end mobility-aware resilience solutions addressing both accidental and malicious faults, where the user perception of trustworthiness is a key issue. Scenario-based analysis and validation of these solutions is performed via analytic/simulation models, and via an experimental proof-of-concept prototype.

1.1.1.3 RESIST - Resilience for Survivability in IST[8]

In the context of ReSIST Network of Excellence, an important role is played by the challenges dependability assessment faces. In particular, existing technologies are evaluated taking into account the scaling challenges of large and evolving modern-day systems.

1.1.1.4 CRUTIAL - Critical Utility Infrastructure Resilience[9]

In CRUTIAL, assessment was studied in the context of interdependent critical infrastructures in general, and electric power system infrastructures, in particular. The project applies model based assessment, using discrete-event simulation to deal with difficult to analyse practical fault models including dependencies.

1.1.1.5 MASTER – Managing Assurance Security and Trust for Services[10]

MASTER will provide methodologies and infrastructure that facilitate monitoring, enforcement, and auditing of security compliance, especially where highly dynamic service oriented architectures are used to support business process enactment in single, multi-domain, and iterated contexts. MASTER focuses on the regulatory requirements related to IT support of application of security policies to business processes in organisations. The project includes a 'State of the art on security and continuity management tools' deliverable aimed at providing an overview of the existing security and continuity management tools that can inform and inspire the design of a control cockpit, but also have a sort of report on an inspirational benchmarking.

Summary of taxonomies referenced in IATAC

This annex summarizes the taxonomies referenced in IATAC:

---

[6] *http://www.deserec.eu*
[7] *http://www.hidenets.aau.dk*
[8] *http://www.resist-noe.eu,*
[9] *http://crutial.cesiricerca.it*
[10] *http://www.master-fp7.eu*

WISSSR taxonomy

The WISSSR taxonomy structures the metrics around certain aspects of information security. The subject matter addressed in the WISSSR Workshop fell into two main categories:

- Organisational security
- Technical Target of Assessment

NIST types of measures

NIST SP 800-55 Rev. 1 provides an informal taxonomy in Section 3.3, 'Types of Measures.' The publication identifies three categories of measures:

- **Implementation measures**: Used to demonstrate the organisation's progress in implementing information security programs, specific security controls, security of system-level areas, and policies and procedures associated with any of these;
- **Effectiveness/Efficiency measures**: Used to determine whether program-level processes and system-level security controls have been implemented correctly, operate as intended, and achieve their intended (desired) outcomes. Effectiveness/efficiency measures reflect two aspects of the results of security control implementation: the robustness of the result itself (i.e., its effectiveness) and the timeliness of the result (i.e., its efficiency);
- **Impact measures**: Articulate the impact (i.e., business or mission impact) of information security on the organisation's ability to accomplish its mission.

I3P taxonomy of security metrics

The purpose of the Institute for Information Infrastructure Protection (13P) taxonomy is to categorize measurement of security of process control systems, e.g., SCADA systems.

The developers of this taxonomy used as a starting point three implied IA metrics taxonomies:

- Categorization of CS/IA measurement subject matter at the WISSSR;
- Control objectives in ISO/IEC 17799 'Information technology – Security techniques – Code of practice for information security management';
- Categories of technologies in American National Standards Institute (ANSI)/International Society of Automation (ISA)-TR99.00.01-2004 'Security Technologies for Manufacturing and Control Systems'.

They identified 5 categories of metrics:

- Security controls in ISO/IEC 17799
- Security controls in (ISA)-TR99.00.01-2004
- Organisational metrics
- Operational metrics
- Technical metrics

Department of Public Safety and Emergency Preparedness Canada

This taxonomy was defined for the Department of Public Safety and Emergency Preparedness to measure results of network assessments. Its measures fall into three categories, with the same three sub-categories within each category:

- Security metrics
  - Technical
  - Organisational
  - Operational
- Quality of Service metrics
  - Technical
  - Organisational
  - Operational
- Availability metrics
  - Technical
  - Organisational
  - Operational

VTT Technical Research Centre of Finland Security Metrics Taxonomy

VTT Technical Research Centre proposed a taxonomy, intended to 'bridge the gaps between business management, information security management, and information and communication technology product security measurement practices.'

The proposed taxonomy is divided in 3 metric groups:

- Business level security;
- Security metrics for organisation's Information Security Management (ISM);
- Security, dependability and trust metrics for products, systems and services.

Daniel Geer's Balanced Scorecard-based taxonomy

In his tutorial Measuring Security, Daniel Geer suggests a taxonomy based on the four corners of a balanced scorecard:

- Financial vs. Security;
- Internal Business Process vs. Security;
- Learning and Growth vs. Security;
- Customer vs. Security.

Towards a unified taxonomy of resilience metrics

In this section we present a two-dimensional approach **to categorising resilience metrics.** This is a first attempt to bring together different taxonomies in a single unified model. The model includes an incident- and a domain-/discipline-based dimension.

The two-dimensional classification is a flexible model. On one hand takes the incident-based view of classifying resilience metrics before an 'event' happens that is preparing for resilience and delivering the intended service, and after the 'event', while trying to respond and recover to normal operation.

On the other hand the model recognises the multi-disciplinary and multi-domain nature of resilience, covering for example areas from disciplines, called domains thereafter, such as security, dependability, performability etc.

Up to now the taxonomies that were proposed in the literature (see previous section for some examples) are mainly referred to the domain dimension of this classification. The domains included and the levels of details of domain-based classifications differ from proposal to proposal.

The two-dimensional model though is independent of the actual domains included in the resilience domain dimension. Therefore we do present two example domain-based approaches. These should be seen only as a proposal for possible candidate domains to be included in the general taxonomy.

The two-dimensional taxonomy

The goal of a common taxonomy is to logically structure the different metrics in groups, in order to emphasize the common properties. This enables a better understanding of the generic resilience properties and provides a common language and understanding of the issues behind the underlying concepts.

During the study that ENISA undertook on the subject in 2010 on the different resilience metrics, it was found that up to now the taxonomies that were proposed in the literature (see previous section for some examples) mainly referred a discipline-/domain-based grouping of the various metrics. The domains included and the levels of details of these domain-based classifications differ from proposal to proposal.

During the same study it became apparent that there was another dimension in the classification of resilience metrics which is related to the temporal view of an event/incident. During the discussion with experts even though this approach was not heavily represent in the literature had wider acceptance mainly because it relates directly with the definition of resilience.

In an effort to come with a unified approach we present a two-dimensional classification model for resilience metrics is a flexible model.

On one hand the model takes the incident-based view of classifying resilience metrics before an incident happens that is preparing for resilience and delivering the intended service, and after the incident, while trying to respond and recover to normal operation. On the other hand the model recognises the multi-disciplinary and multi-domain nature of resilience, covering for example areas from disciplines, called domains thereafter, such as security, dependability, performability etc.

The two-dimensional model though is independent of the actual domains included in the resilience domain dimension. Therefore we do present two example domain-based approaches. These should be seen only as a proposal for possible candidate domains to be included in the general taxonomy.

The one dimension of the classification model is based on the principle that resilience metrics can be categorized according to a **temporal dimension related to the incident**. In our classification system we will call this the **incident-based** dimension. This way of looking into grouping the resilience metrics is explained in more detail in section 0 but in summary, it is possible to express resilience over the 3 different time phases with respect to challenges and faults (events) that threaten the normal level of service:

- **Preparation** phase: Resilience provisions are implemented in order to prepare the network/service for coping with faults and challenges. Metrics in this dimension measure how well systems and services are prepared to cope with challenges/faults.
  A high preparedness metric indicates a reduced failure probability, i.e. reduced likelihood of damage & failures to critical infrastructure, systems and components.

- **Service Delivery** phase: The network/service is operational and detects occurrences of faults and challenges. Metrics in this dimension measure the difference in service level before, during and after the fault or challenge.
  A low metric (a high difference in service level) indicates that the consequences of a fault or challenge on the network are reduced.

- **Recovery** phase: When the network/service is no longer at an acceptable level of service, recovery is initiated to restore normal operations. Metrics in this dimension revolving about how fast a service/network can recover from faults/challenges.
  A low metric indicates reduced time to recovery (the time required to restore a network or a service to the normal level of functionality).

The second dimension of the proposed taxonomy is based on the **parts of different disciplines,** called **domains** thereafter, which collectively constitute the notion of resilience. A metrics domain is a group of metrics which are measuring different aspects of the same resilience property. This approach is explained further in section 0. One can define domains at various levels of detail and abstraction. An example of a possible high level domain is ‘security’ while an example of a finer abstraction domain is ‘Patch management’ where all metrics belong to that measure which systems are regularly patched, what the average time is to patch a system, etc.

Both dimensions of this taxonomy model could be considered as classifications. This document does not provide an exhaustive list of possibilities, for all possibilities of domain-based classifications.

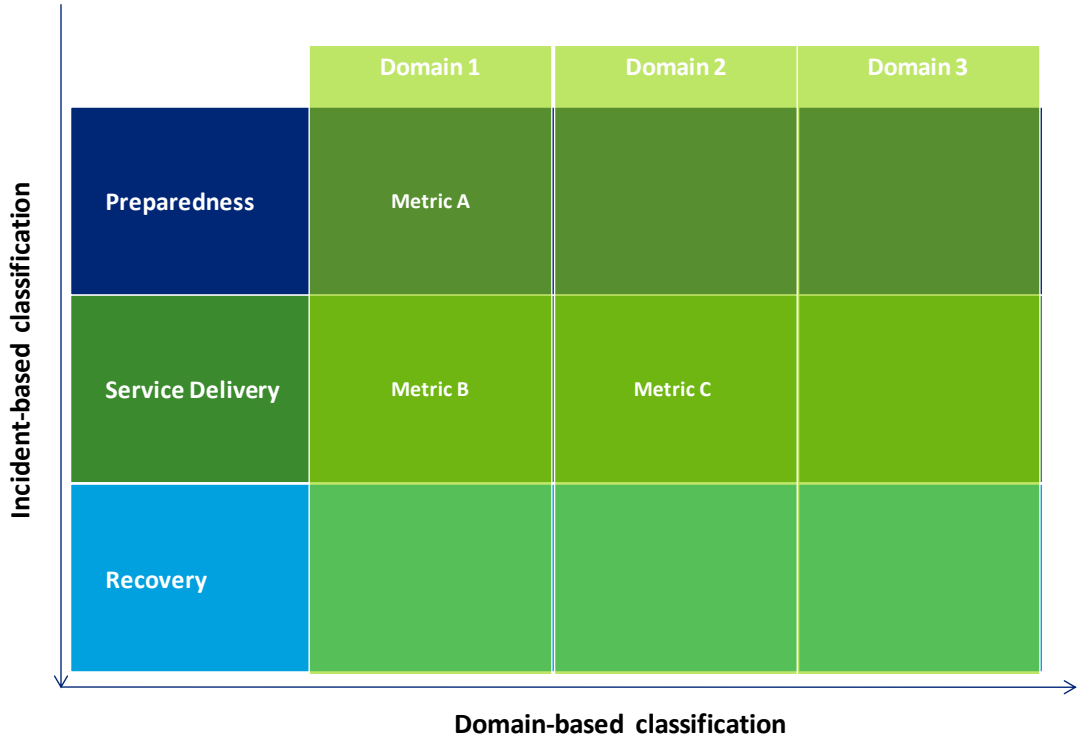

Figure 4: The two-dimensional taxonomy

Figure 4 demonstrates this principle:

- Metric A is a metric belonging to the incident-based dimension 'preparedness' and to 'Domain 1' in the domain dimension;
- Metric B belongs to the same domain as Metric A, but this metric is measured during the service delivery phase: it thus belongs to a different incident-based dimension compared to Metric A;
- Metric C is also belongs to the service delivery incident-based dimension but measures another resilience property compared to metric B.

The following sections illustrate different possible options, based on the two-dimensional taxonomy as shown above.

The incident-based classification

The classification in this paragraph concentrates on the idea of provision and maintainability of an acceptable level of service. It comes directly from the definition of resilience in section 0.

The service level can be compromised when an event such as a security incident, a system failure or a human error occurs. If nothing happens the service delivery remains stable. Our approach of identifying resilience metrics is event-based.

By dividing time into phases according to the state of the network/system when an incident occurs, we conclude to the following high level taxonomy.

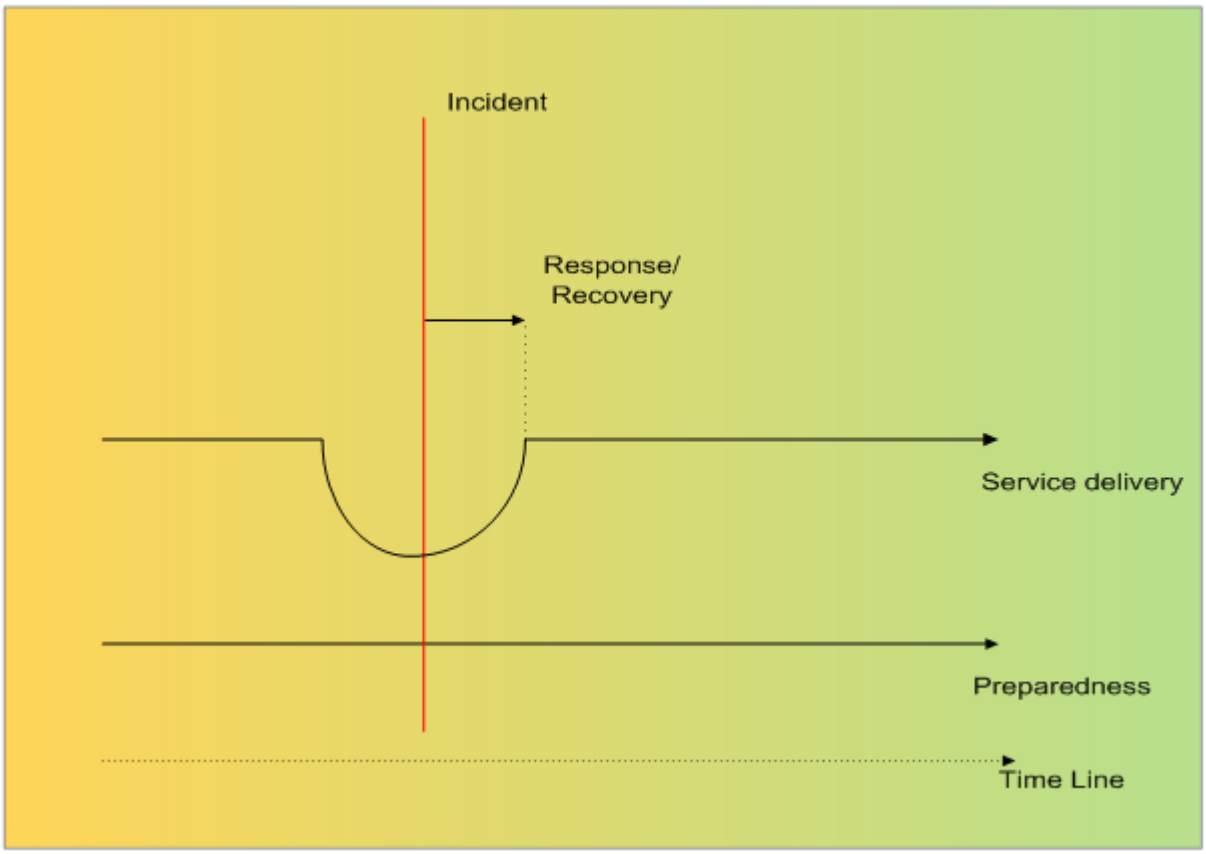

Figure 5: Time- and incident-based service delivery separation

The life time of the network when an incident occurs, is divided in three phases; the *preparedness*, the *service delivery* and the *response-recovery* phase.

During the **preparedness** phase the state of the system is stable. Preparedness includes all the actions and measures taken so as to prevent an incident from happening, or diminish impact to the minimum level. Preparedness measures are the umbrella that covers end to end the system and is fully operating even during the incident time. In effect, the preparedness level doesn't differ even when an incident happens. Normal operation of the system is parallel to preparedness.

The level of the **service** the network delivers is stable (in a level higher than the minimum level of service) and decreases to the minimum level of service when the incident occurs. After the incident the system tries to recover to the previous state; during this period, from the occurrence of the incident until the system recovers fully to its previous level of service delivery, the **response/recovery** phase emerges simultaneously. Without the occurrence of a breach of security which causes the system to fail, the response phase does not occur. Recovery includes all the mechanisms used to eliminate the impact and bring back the system to its service level.

All the metrics proposed derive from the analysis of the incident occurrence conceptual scheme. An incident occurs. The impact of this incident is measured by how much the level of service decreases, how steep is the curve. In the same direction when the recovery- response phase starts, the steepness of the curve defines the efficiency of the system to recover. But how can we define an incident?

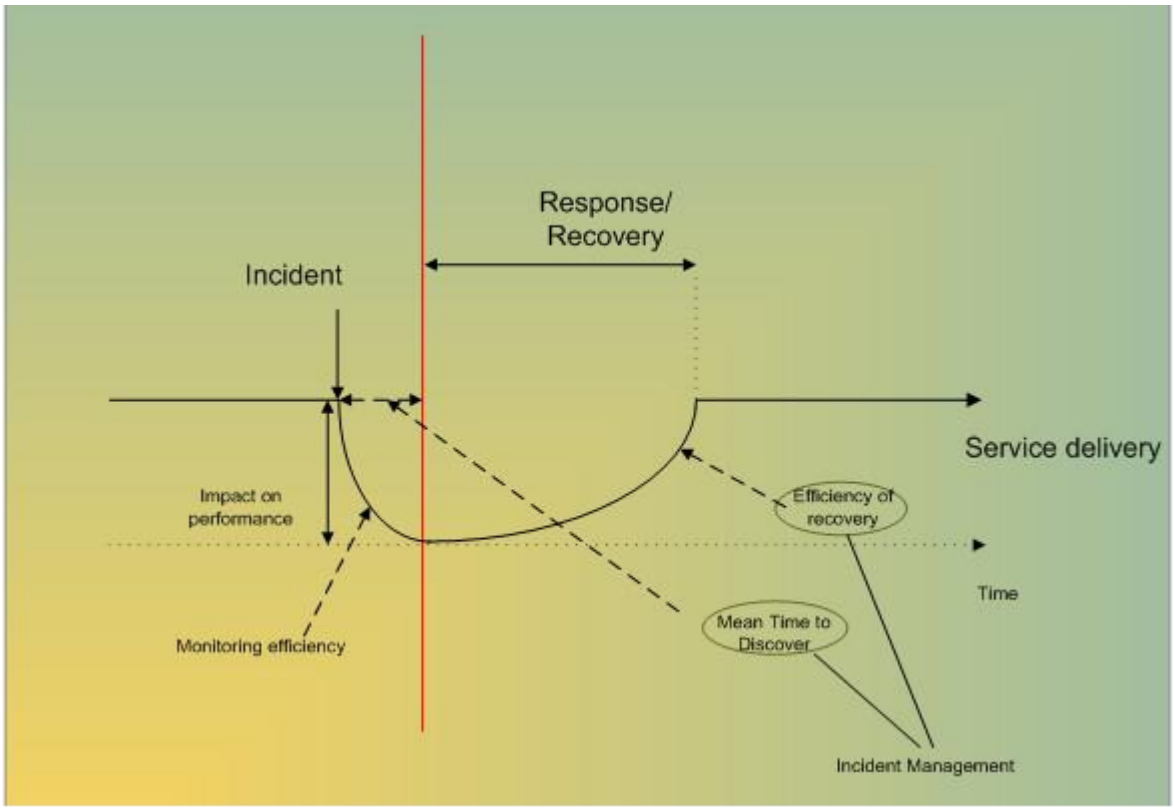

Figure 6: Example metrics in various time periods

In a period of time if a network is monitored, many events occur which bypass the peripheral security mechanisms and can degrade the level of service. Security events are the actions that can cause an incident. The incident to be measured has specific characteristics: a) it has an impact (light or severe impact), b) it has duration and c) it can be discovered and eliminated.

An incident can be categorised as light or severe according to the decrease (the steepness) of the service (the curve). If the service level 'falls' lower than a predetermined 'minimum security level' then automatically it can be characterised as severe as the impact can be grave for the system's performance. The efficiency of the systems' protective mechanisms can be measured by the percentage numbers of incidents / number of events (incidents have impact, events don't) and by the category if the incidents (light or severe). The minimum security level used for classifying incidents (which is objective and is set by each company according to its priorities) must in line with the service level requirements in place. If the minimum service level used is too strict or too loose, the incident report will not reflect the reality experienced by the network service users.

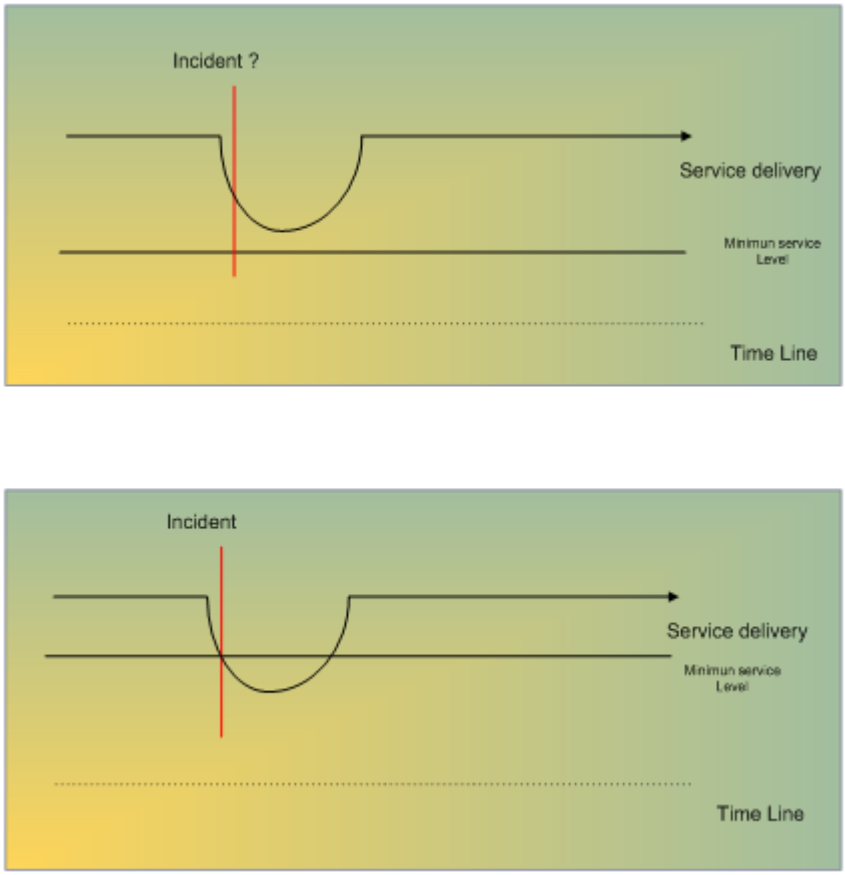

Figure 7: Minimum service level

The key elements in the proposed taxonomy are the Preparedness, Service Delivery and Response – Recovery Phase. To apply the resilience metrics, each phase includes domains by which the metrics (indicators) are concluded. Having a closer look, a two-dimensional taxonomy is created which can be incorporated to any system. The flexibility of this scheme is that domains can be included and metrics added, according to the field of each corporation.

The domain-based classification

The second dimension of the proposed taxonomy is based on **the parts of different disciplines,** called **domains** thereafter, which collectively constitute the notion of resilience. A metrics domain is a group of metrics which are measuring different aspects of the same resilience property. This model recognises the multi-disciplinary and multi-domain nature of resilience, covering for example areas from disciplines, called domains thereafter, such as security, dependability, performability etc.

One can define domains at various levels of detail and abstraction. An example of a possible high level domain is 'security' while an example of a finer abstraction domain is 'Patch management' where all metrics belong to that measure which systems are regularly patched, what the average time is to patch a system, etc.

The taxonomies that were proposed in the literature are mainly referred to the domain–based classification. The domains included and the levels of details of domain-based classifications differ from proposal to proposal.

In the two-dimensional model the domain-based classification is recognised as one important dimension. It is though independent of the actual domains included. Therefore we do present two example domain-based classification approaches. These should be seen only as a proposal for possible candidate domain-based classification to be included in the general taxonomy, as consensus on the domains included still needs to be achieved.

In the first example of domain-based classification below we present the classification model (including the domains that have been identified) as defined by the ResiliNets research initiative. The second example classification proposes a simplified model of the ResiliNets classification.

Domain-based classification example 1: The ResiliNets classification

The network services resilience framework presented is defined by the ResiliNets[11] initiative collaboration between The University of Kansas, US and Lancaster University, UK).

The research initiative provides a decomposition of the resilience concept into a number of disciplines broadly classified into two categories: Challenge Tolerance and Trustworthiness. This is graphically illustrated in the figure below:

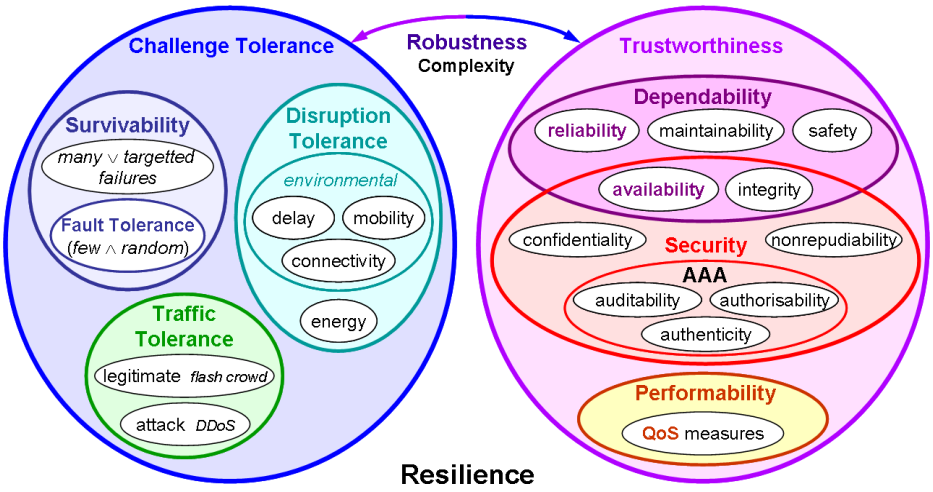

Figure 8: ResiliNets metrics taxonomy

The ResiliNets initiative identifies following dimensions of resilience:

The **'Challenge Tolerance'** dimension considers the ability of the network to withstand faults and challenges:

[11] *ResiliNets initiative - https://wiki.ittc.ku.edu/resilinets/Main_Page*

- **Survivability** is the capability of a system to fulfil its mission, in a timely manner, in the presence of threats such as targeted attacks or large-scale natural disasters resulting in many failures, in addition to the few random failures covered by fault tolerance. **Fault Tolerance** is the ability of a system to tolerate faults such that service failures do not result. Fault tolerance generally covers random single or at most a few faults, and is thus a subset of survivability, as well as of resilience. Survivability is thus a superset of fault tolerance but a subset of resilience.
- **Disruption Tolerance** is the ability of a system to tolerate disruptions in connectivity among its components. Disruption tolerance is a superset of tolerance of the environmental challenges: weak and episodic channel connectivity, mobility, delay tolerance, as well as tolerance of power and energy constraints.
- **Traffic Tolerance** (also referred to as elasticity) is the ability of a system to tolerate unpredictable offered load without a significant drop in carried load (including congestion collapse), as well as to isolate the effects from cross traffic, other flows, and other nodes. The traffic can either be unexpected but legitimate such as from a flash crowd, or malicious such as a Distributed Denial of Service attack.

The **'Trustworthiness'** dimension brings a number of additional quantifiable properties of resilience:

- **Dependability** is the property of a system or network such that reliance can justifiably be placed on the service it delivers. It generally includes the measures of availability (ability to use a system or service) and reliability (continuous operation of a system or service), as well as integrity, maintainability, and safety.
- **Security** is the property of a system or network of being protected from unauthorised access or change, subject to policy. Security properties include AAA (auditability, authentication and accountability), confidentiality, and non-repudiation. Security shares with dependability the properties of availability and integrity.
- **Performability** is the property of a system such that it delivers performance required by the service specification, as described by QoS (quality of service) measures.

Based on this decomposition, metrics of resilience could be classified in 2 main types:

- **Fault and challenge tolerance** metrics reflect the resilience provisions of the network and indicate the preparedness of the network services to provide and maintain an acceptable level of service in face of disturbances.
  These metrics indicate if certain defensive provisions have been taken in order to maintain an acceptable level of service when facing challenges and how well these challenges can be tolerated. An example of a challenge tolerance metric is the number of correctly patched workstations in an organisation: it indicates, before hackers try to attack these workstations occurs, how many workstations are protected and immune to certain types of attacks.

- **Trustworthiness** metrics measure network service resilience in terms of the operational state, with regards to resilience, of a network or service after (or during) the occurrence of disturbances to normal operation.

These are quantifiable measures that characterize the quantifiable properties of the main resilience objectives: dependability, security and performability.
An example of such a metric is the availability of a service: High availability of a service indicates that it is able to maintain the needed level of service, when experiencing challenges.

The ResiliNets initiative also defines a strategy for implementing resilience, referred to as D^2R^2+DR:

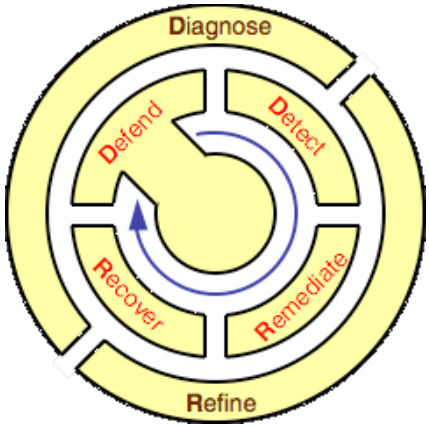

Figure 9: ResiliNets D^2R^2+DR strategy

The strategy is defined as a control loop, consisting of following processes:

- **Defend** against challenges and threats to normal operation;
- **Detect** when an adverse event or condition has occurred;
- **Remediate** the effects of the adverse event or condition to minimise the impact;
- **Recover** to original and normal operations.

While these processes are running, the overall effectiveness of the control loop should be **diagnosed** and **refined.**

Domain-based classification example 2: A simplified domain-based classification approach

In this example domain-based classification the different domains are based on a simplified version of the ResiliNets classification for network service resilience metrics.

Quoting from the ResiliNets initiative [43], 'Resilience subsumes a number of disciplines, many of which are tightly interrelated but have developed separately, sometimes with inconsistent language. We broadly classify these disciplines into two major categories: those that are related to the tolerance of challenges and faults, and trustworthiness that considers aspects that can be measured'.

Based on this one can limit the domains to the aspects of trustworthiness that can be measured, i.e. dependability, security and performability as illustrated also in Figure 10 below.

The ResiliNets initiative provides the following definitions [42]:

- **Dependability** is the property of a system such that reliance can justifiably be placed on the service it delivers. It generally includes the measures of availability (ability to use a system or service) and reliability (continuous operation of a system or service), as well as integrity, maintainability, and safety.
- **Security** is the property of a system and measures taken such that it protects itself from unauthorised access or change, subject to policy. Security properties include AAA (auditability, authorisability and authenticity), confidentiality, and non-repudiability. Security shares with dependability the properties of availability and integrity.
- **Performability** is the property of a system such that it delivers performance required by the service specification, as described by QoS (quality of service) measures.

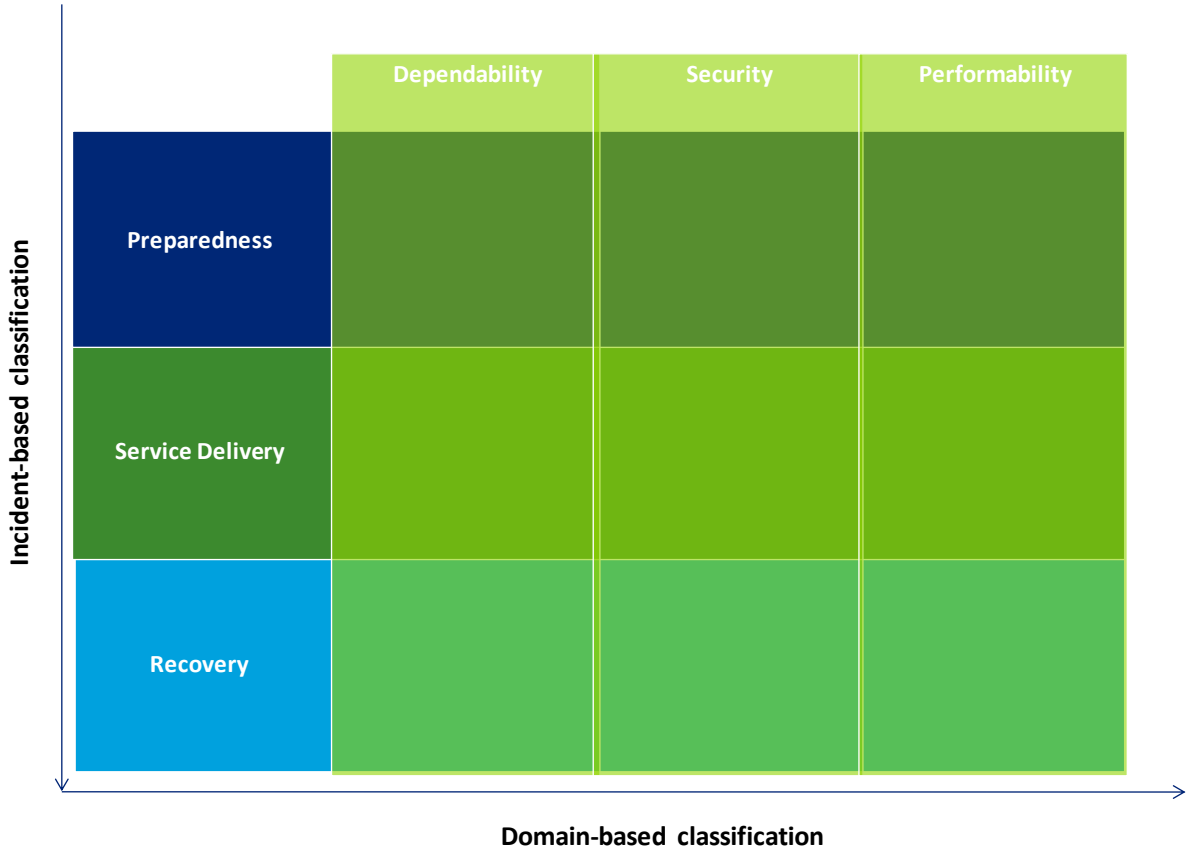

Figure 10: A simplified approach to resilience metrics

The two examples presented above just highlight the different options that one can have in a domain-based classification. The domains used are mainly a matter of definition of the different disciplines and a subject to further work in order to reach to consensus.

Open issues – composition and aggregation

In the previous sections, the two-level taxonomy used represented the metrics that could be measured and used within a single corporation or at a level where single and unified measurements are possible. It is clear that this one is not enough when one wants to have the resilience status at

different levels of abstraction. One example is when we want to understand the *resilience status* beyond the level of a single corporation or entity, for example on sector-wide basis, on national basis or even on a pan-European level.

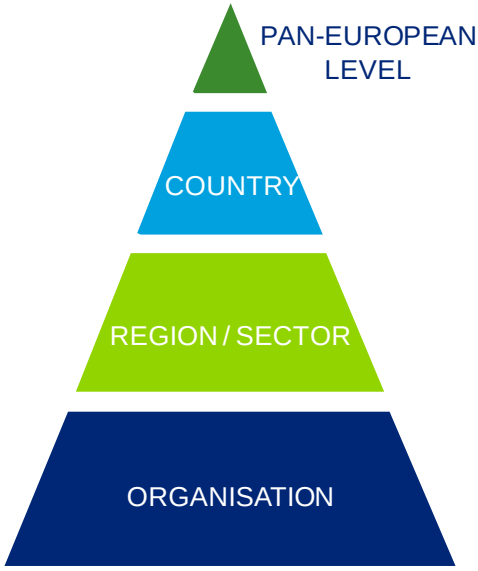

Figure 11: The use of metrics in different levels of detail

Following the proposed taxonomy and the individual baseline metrics that are presented in the following section, each corporation will use these metrics to indicate the resilience of the system. It is still an open issue how this can be used at different level of abstraction, as in the hierarchical view represented in Figure 11.

In order to assess the resilience status at higher levels of abstraction the use of individual metrics is not enough. Aggregation and composition of metrics will be required to achieve this.

With these two techniques we may be able to formalise the definition of different metrics, using other metrics. This area needs certainly to be further investigated and studied in the field of resilience metrics.

This topic is directly tied to the European directive 2009/140/EC article 13 (cf. section 0) which obliges the different Member States to 'ensure that undertakings providing public communications networks or publicly available electronic communications services notify the competent national regulatory authority (NRA) of a breach of security or loss of integrity that has had a significant impact on the operation of networks or services'. The competent NRA in turn needs to inform the NRAs of the other Member States.

It is exactly this quantification of breaches of security or losses of integrity that is enabled by using a common set of resilience metrics between the different organisations and NRA's. The development of aggregation/composition techniques which will express metrics from the level of an organisation or sector into a country-wide or pan-European level will enable a common understanding of the

different breaches and integrity losses across all stakeholders in Europe, using comparable metrics that can express the communications services' resilience level.

Baseline resilience metrics

This chapter will describe a number of quantifiable measures for network service resilience.

Section 0 will detail a number of measures to quantify the impact of network service degradation. Impact is a critical part of the resilience equation, describing the effect of outages and measuring its impact is therefore very important.

In section 0, some important baseline resilience metrics are explained and structured according to the incident-based dimension (cfr. section 0).

These metrics apply to ICT systems, which are defined as “systems related to Information and Communications Technology required for an organisation to run its business”.

While section 0 describes the different preparedness metrics identified in this document, a number of theoretical counterparts to the preparedness metrics were identified – in this report, they are called preparedness indicators or design-based resilience metrics. They are described in section 0.

The following template will be used as a generic template to describe the resilience metrics in detail. It includes the most important aspects of metrics that should be outlined by a measurement framework for resilient networks and services.

Stakeholders implementing the framework may choose to customize the template in accordance with their preferences and needs, by using a subset of the given fields or by adding more fields.

Metric name	Standard or assigned name, used to reference the metric.
Source	Indication of the literature from which the metric was adopted.
Description	Description of the metric, explaining the concept / attribute under measurement and the measures from which the metric is derived.
Objective	Description of the resilience measurement goal. What value does it bring to measure the metric? What conclusions could be derived from the metric? What purpose does the metric serve?
Measurement method	Description of the base measures and units of measurement, and the formula to calculate the numeric metric value of the metric; The formula consists of a mathematical function of 2 or more measures. The measurement method for these measures needs to be accurately described as well, in order to assure repeatability and comparability of metrics.
Frequency	Number of times per period that the data will be collected in order to measure the metric. The frequency will be dependent on several factors, including the rate of change in the measure attribute, compliance & reporting requirements, business specifics, etc.
Target values	Threshold for an acceptable value of the metric. The target value can be part of, for example, a service level agreement or a performance goal in a Capability Maturity Model.
Reporting format	Description of an example reporting format to visually or verbally best characterize the metric.

Additional information that could be collected for metrics as part of a measurement framework would be for example (refer to [6] [9]):

- Type (according to a certain or multiple classifications);
- Implementation evidence;
- Control or process under measurement;
- Object and attribute under measurement;
- Derived measures;
- Analytical model;
- Decision criteria;
- Stakeholders.

Impact metrics

Section 0 already introduced the importance of impact quantification for incidents which degrade network service. This section indicates a number of metrics which measure the impact of disturbances in the network service offered. While these impact metrics do not provide insight in the level of resilience of the network service, they do indicate the graveness of the network service incident.

The most common impact metrics are:

- Number of users affected by the network service disruption;
- Number of network elements affected by the network service disruption;
- Geographical area impact by the network service disruption;
- Financial impact of network service disruption (financial liabilities such as contractual fines);
    - Monetary cost;
    - Loss of market share;
    - Percentage of decrease in revenue.
- Criticality of impacted and dependent services;
- Number of / increase in helpdesk calls or incident tickets;
- Reputation damage.

Resilience metrics

In this section, we present a number of identified resilience metrics. It is not the intention of this document to include the full list of baseline metrics.

Where possible, each metric is associated with target values. Due to the specific nature of different network services existing and the importance of those services to their customers, it is very difficult to include target values for all metrics. The study of target values should be included in future work.

When using the example taxonomy of section 0, the metrics can be categorized according to Figure 12.

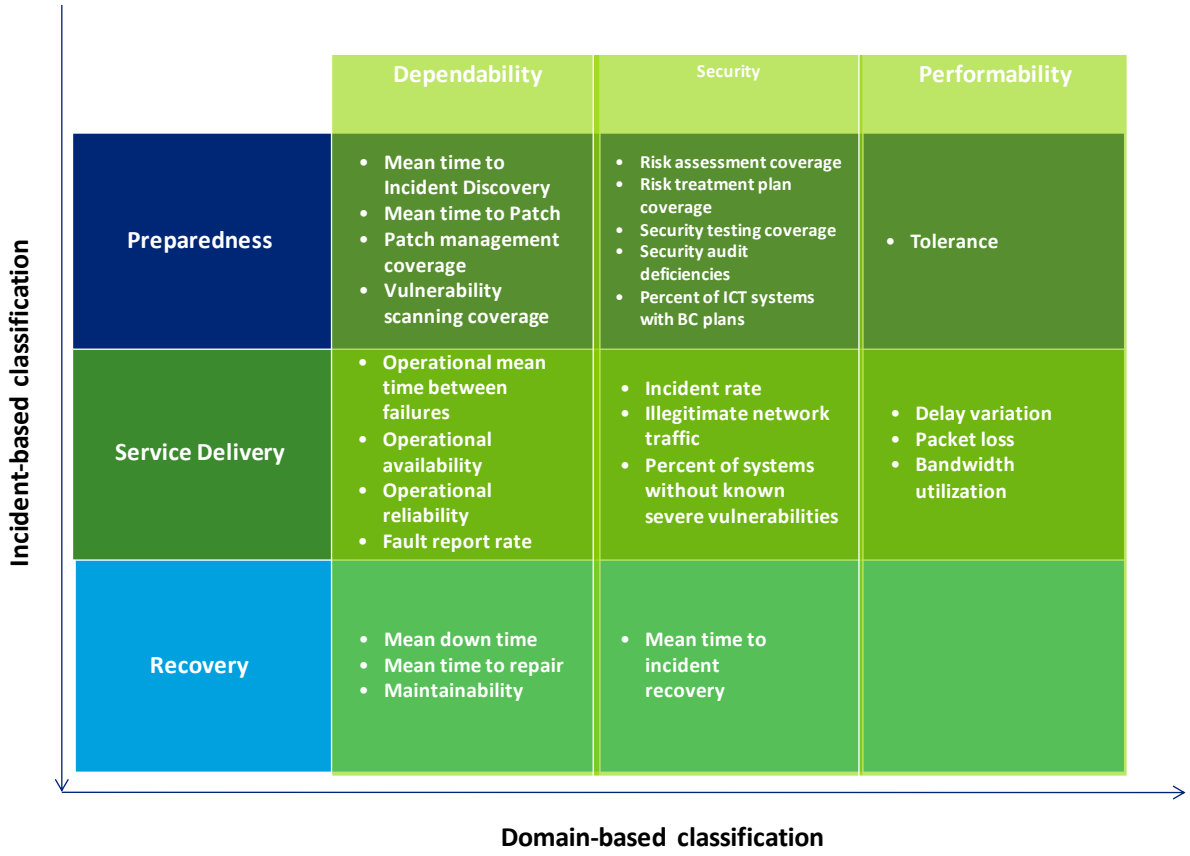

Figure 12: Metrics categorized in the taxonomy of section 0

Preparedness phase

This section will specify a number of resilience metrics that belong to the preparedness phase. The included metrics in this dimension measure how well systems and services are prepared to cope with challenges and faults.

The metrics covered were selected on their ability to be implemented pragmatically and the feasibility of accurate measurements.

More specifically we present following metrics in more detail:

- Mean time to Incident Discovery (see 1.1.1.6);
- Mean time to Patch (see 1.1.1.7);
- Patch management coverage (see 1.1.1.8);
- Vulnerability scanning coverage (see 1.1.1.9);
- Tolerance (see 1.1.1.10);
- Risk assessment coverage (see 1.1.1.11);
- Risk treatment plan coverage (see 1.1.1.12);
- Security testing coverage (see 1.1.1.13);
- Security audit deficiencies (see 1.1.1.14);
- Percent of the ICT systems with business continuity plans (see 1.1.1.15).

Metrics that are not elaborated in this document but can be useful to measuring the preparedness phase of resilience are:

- Percentage of ICT systems for which availability requirements have been specified;
- Percentage of ICT systems for which recovery procedures have been defined and implemented;
- Collateral damage;
- Percentage of contracts with subcontractors and partners that provably encompass suitable clauses with respect to information security;
- Percentage of subcontractors and partners for which compliance with contractual information security agreements has provably been evaluated or tested.

1.1.1.6 Mean Time to Incident Discovery

Metric name	MTTID: Mean Time To Incident Discovery
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	Mean time to incident discovery characterizes how effective organisations are in the detection of incidents, by measuring the average elapsed time between the initial occurrence of an incident and the discovery thereof. The MTTID metric expresses how well an organisation is prepared against incidents: a higher score means that incidents are discovered fast (on average) compared to organisations with a lower score.
Objective	Generally, the faster an organisation can detect an incident, the less damage it is likely to incur, as such the MTTID metric serves as a leading indicator of resilience in an organisation: a short MTTID is a sign of an organisation with effective security incident monitoring and detection, which will aid in maintaining an acceptable level of service.
Measurement method	MTTID is the amount of time, in hours, that elapsed between the Date of Occurrence and the Date of Discovery for a given set of incidents, divided by the number of incidents. $$MTTID = \frac{\sum_i (Date\_of\_Discovery_i - Date\_of\_Occurrence_i)}{Number\_of\_incidents}$$ The calculation can be averaged over a time period, and grouped per types of incidents, business units, or incident severity. Unit of the metric is hours/incident. The NIST incident handling guide [44] recommends apply granularity to this metric by using following incident categories: • Denial of Service—an attack that prevents or impairs the authorized use of networks, systems, or applications by exhausting resources • Malicious Code—a virus, worm, Trojan horse, or other code-based malicious entity that successfully infects a host • Unauthorized Access—a person gains logical or physical access without permission to a network, system, application, data, or other IT resource • Inappropriate Usage—a person violates acceptable use of any network or computer policies • Multiple Component—a single incident that encompasses two or more

	incidents. Usage of a common categorization scheme will allow for an accurate view when aggregating data across companies in a sector or region.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	MTTID values should trend lower over time. The value of '0 hours' indicates hypothetical instant detection times. There is evidence the metric result may be in a range from weeks to months (2008 Verizon Data Breach Report). Because of the lack of experiential data from the field, no consensus exists on the range of acceptable goal values for MTTIDs.
Reporting format	Reporting of the incident rate should be per category and based on the hours/incident value.

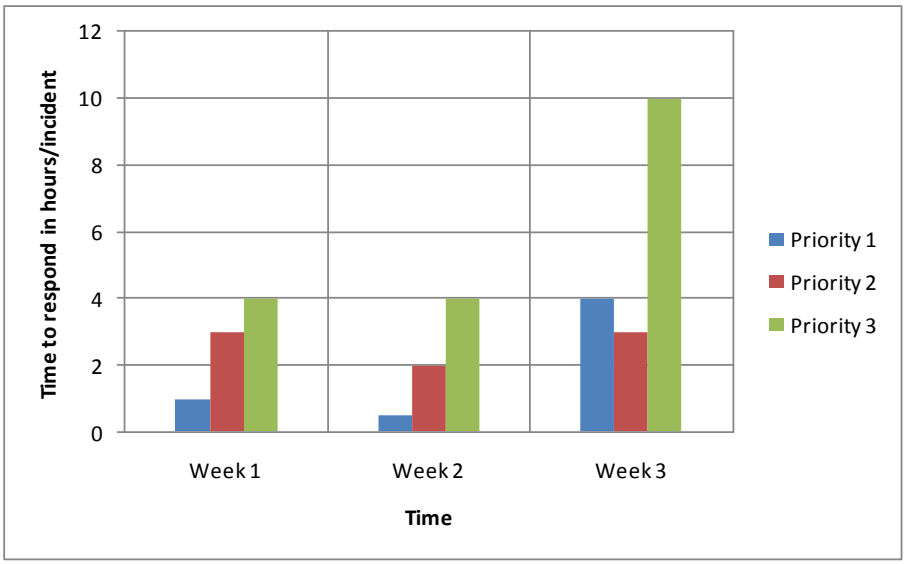

Figure 13: MTTID reporting example

1.1.1.7 Mean time to Patch

Metric name	MTTP: Mean Time To Patch
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	Mean Time to Patch (MTTP) characterizes the effectiveness of the patch management process by measuring the average time taken from date of patch release to installation in the organisation for patches deployed during the metric time period. This metric serves as an indicator of the organisation's overall level of exposure to vulnerabilities by measuring the time the organisation takes to address ICT systems known to be in vulnerable states that can be remediated by security patches. This is a partial indicator as vulnerabilities may have no patches available or occur for other reasons such as system configurations.
Objective	Mean Time to Patch (MTTP) measures the average time taken to deploy a patch to the organisation's ICT systems. The more quickly patches can be deployed, the lower the mean time to patch and the less time the organisation spends with systems in a state known to be vulnerable.
Measurement method	MTTP is calculated by determining the number of hours between the Date of Availability and the Date of Installation for each patch completed in the current scope, for example by time period, criticality or business unit. These results are then averaged across the number of completed patches in the current scope. $$MTTP = \frac{\sum_i Date\_of\_installation_i - Date\_of\_availability_i}{Count\_Of\_Completed\_Patches}$$ The unit is hours per patch.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	MTTP values should trend lower over time. Most organisations put patches through test and approval cycles prior to deployment. Generally, the target time for MTTP will be a function of the criticality of the patch and business criticality of the technology. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for Mean Time to Patch exists.
Reporting format	MTTP can be calculated over time, typically per-week or per-month. To gain insight into the relative performance and risk to one business unit over another,

MTTP may also be calculated for different patch criticalities and cross-sections of the organisation, such as individual business units or geographies.

1.1.1.8 Patch management coverage

Metric name	Patch management coverage
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	Patch management coverage measures the relative amount of an organisation's ICT systems that are managed under a patch management process such as an automated patch management system.
Objective	The higher the percentage of technologies managed under an automatic patch system, the timelier and more effectively patches are deployed to reduce the number and duration of exposed vulnerabilities. This metric also serves as an indicator of the ease with which security-related changes can be pushed into the organisation's environment when needed.
Measurement method	Patch management coverage is calculated by dividing the number of the ICT systems under patch management by the total number of ICT systems within the organisation. $$PMC = \frac{\sum Count\_Of\_ICT\_Systems\_With\_Patch\_Management}{Count\_Of\_ICT\_Systems}$$ This metric can be calculated for subsets of ICT systems such as by asset criticality or business unit. The metric is relative and unitless: it is the percentage of the ICT systems under patch management.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	Patch management coverage values should trend higher over time. Given the difficulties in manually managing ICT systems at scale, having technologies under patch management systems is preferred. An ideal result would be 100% of technologies. However, given incompatibilities across ICT technologies and systems this is unlikely to be attainable. Higher values would generally result in more efficient use of security resources. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for PMC exists.
Reporting format	Reporting of the patch management coverage should be the percentage on a time-scale to show the evolution.

1.1.1.9 Vulnerability scanning coverage

Metric name	Vulnerability scan coverage
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	Vulnerability Scan Coverage (VSC) indicates the scope of the organisation's vulnerability identification process. Scanning of ICT systems known to be under the organisation's control provides the organisation the ability to identify open known vulnerabilities on their ICT systems. Percentage of ICT systems covered allows the organisation to become aware of areas of exposure and proactively remediate vulnerabilities before they are exploited.
Objective	Vulnerability Scanning Coverage (VSC) measures the percentage of the organisation's ICT systems under management that were checked for vulnerabilities during vulnerability scanning and identification processes. This metric is used to indicate the scope of vulnerability identification efforts.
Measurement method	Vulnerability Scanning Coverage is calculated by dividing the total number of ICT systems scanned by the total number of ICT systems within the metric scope such as the entire organisation: $$VSC = \frac{Number\_of\_scanned\_ICT\_systems}{Total\_number\_of\_ICT\_systems} * 100$$ The metric is expressed as a percentage.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	VSC values should trend higher over time. Higher values are obviously better as it means more ICT systems have been checked for vulnerabilities. A value of 100% means that all the ICT systems are checked in vulnerability scans. For technical and operational reasons, this number will likely be below the theoretical maximum.
Reporting format	Organisations can use this metric to evaluate their risk position in terms of concentrations of unknown vulnerability states of ICT systems. In combination with other vulnerability metrics, it provides insight on the organisation's exposure to known vulnerabilities. The results of the coverage metric indicate the: • Scope of the vulnerability scanning activities • Applicability of other metric results across the organisation

- Relative amount of information known about the organisation's vulnerability

1.1.1.10 Tolerance

Metric name	Tolerance
Source	This metric definition is based on the definitions from ResiliNets [42].
Description	The tolerance of a service is the permissible limit or limits of variation in a measured value of an ICT system's property before the service level changes from normal to degraded or unacceptable. It can be measured in 3 different aspects: • **Fault tolerance** is the ability of an ICT system to tolerate faults such that service failures do not result. Fault tolerance generally covers single or at most a few random faults and thus measures an aspect of resilience. • **Traffic tolerance** is the ability of an ICT system to tolerate unpredictable offered load without a significant drop in carried load (including congestion collapse), as well as to isolate the effects from cross traffic, other flows, and other nodes. The traffic can either be unexpected but legitimate such as from a flash crowd, or malicious such as a DDoS (Distributed Denial-of-Service) attack. • **Disruption tolerance** is the ability of an ICT system to tolerate disruptions in connectivity among its components. Disruption tolerance is a superset of tolerance of the environmental challenges: weak and episodic channel connectivity, mobility, delay tolerance, as well as tolerance of power and energy constraints.
Objective	Tolerance expresses the extent to which the ICT system can continue to operate when facing challenges or faults (this resistance to challenges is an integral part of the definition of resilience). When the tolerance limit is exceeded, the ICT system will no longer operate at an acceptable level. The tolerance metrics expresses the ability of the ICT system to endure these challenges while remaining at the acceptable service level.

Measurement method	This report cannot provide exhaustive definitions of measurement methods as the tolerance metric should be regarded as a concept metric that can be extended to the specific properties of different systems. This section will enumerate a number of tolerance metrics for specific services but does not constitute an exhaustive list.
	• **Fault tolerance**: Fault tolerance can be provided via:
	  ○ **Redundancy**: Multiple ICT systems exist that can take over from each other;
	  ○ **Replication**: Multiple ICT systems communicate continuously with each other to make sure that share the same state before fail-over;

- **Diversity**: An ICT system can have multiple components that provide the same function. If one of those components fails due to certain circumstances, another component can take over that function.

An example fault tolerance metric for a certain network is the number of ICT systems which are build in a redundant way (i.e. for each redundant system, at least one additional equivalent system that can take over in case of failure) versus the total number of ICT systems in that network. The example metric can be formulated as:

$$Fault\_tolerance = \frac{Amount\_of\_redundant\_ICT\_systems}{Total\_number\_of\_ICT\_systems}$$

The unit of this metric is a percentage, expressing the resilience of that network as a whole against the failure of one of its subsystems. Alternatively said, it expresses the amount of systems that can fail without impacting the network as a whole.

- **Traffic tolerance:** Traffic tolerance is the ability of an ICT system to tolerate unpredictable offered load without a significant drop in carried load.

The ability of a network to handle additional increases in traffic is an example metric for traffic tolerance. It can be measured by measuring an operational parameter measuring the load increase placed on a system versus a specific parameter that measures the level of service.
A possible metric in the increase in additional data traffic (expressed in bits/s²) versus the increase in delay of that network path (expressed in s).

$$Traffic\_tolerance = \frac{Increase\_in\_traffic}{Increase\_in\_delay}$$

The unit of this metric is a percentage, expressing the resilience of that network's service parameters against a variation in its operational parameters.

- **Disruption tolerance:** Disruption tolerance is the ability of an ICT system to tolerate disruptions in connectivity among its components.

An example of disruption tolerance is the number of ICT systems in a network that are protected against power outages (for example: by having a battery backup) versus the total number of ICT systems in that

	network. The example metric can be formulated as: $$Disruption\_tolerance = \frac{Amount\_of\_ICT\_systems\_protected}{Total\_number\_of\_ICT\_systems}$$ The unit of this metric is a percentage, expressing the resilience of that network as a whole against the failure of one of its subsystems by a power outage. Alternatively said, it expresses the amount of ICT systems that can survive without grid power without impacting the network as a whole.
Frequency	Given the instantaneous nature of the load, the different tolerance metrics can be measured in real-time. Other metrics are more static and require an update when significant changes happen to the operational aspects of these systems (for example: the number of systems protected against power outages).
Target values	Tolerance should be as high as possible and can be dictated by the criticality of the services this network support. For example: the disruption tolerance of the PSTN backbone (Public Switched Telephone Network) will be higher than the disruption tolerance of a single phone at a residential customer. Given the generic nature of the metric proposed and the many different implementations that are possible, no target values are specified.
Reporting format	Tolerance metrics can be calculated over time, typically per-week or per-month. To gain insight into the relative performance and risk to one service over another, tolerance may also be calculated for different network services or geographies.

1.1.1.11 Risk assessment coverage

Metric name	Risk assessment coverage
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	This metric reports the percentage of ICT systems that have been subjected to risk assessments.
Objective	Risk assessment coverage indicates the percentage of ICT systems that have been subject to a risk assessment at any time.
Measurement method	The metric is calculated by dividing the number of ICT systems that have been subject to any risk assessments by the total number of ICT systems in the organisation: $$RAC = \frac{ICT\_systems\_which\_undergone\_a\_risk\_assessment}{Total\_of\_ICT\_systems} * 100$$ This metric is expressed as percentage of ICT systems that have undergone a risk assessment.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	RAC values should trend higher over time. A higher result would indicate that more ICT systems have been examined for risks. Most security process frameworks suggest or require risk assessments for ICT systems deployed in production environments. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for Risk Assessment Coverage exists.
Reporting format	This metric can be used to evaluate their risk posture in terms of ICT systems that have undergone a risk assessment. A better understanding of the quantity of ICT systems that have not been exposed to a risk assessment allows the organisation to evaluate their level of unknown risk associated with these ICT systems. With metric results for different dimensions, it is possible to identify and evaluate concentrations of risk, such as for results for critical ICT systems or ICT systems containing confidential information.

1.1.1.12 Risk treatment plan coverage

Metric name	Risk treatment plan coverage
Source	Not applicable
Description	This metric reports the percentage of ICT systems for which risk treatment plans have been documented.
Objective	Risk treatment plan coverage indicates the percentage of ICT systems for which risk treatment plans have been documented, relative to the amount of ICT systems for which a risk assessment has been performed (see 1.1.1.11). The existence of risk mitigation plans indicate that a risk analysis has been performed and the risks have been evaluated. Some of the risks might have been accepted by the organisation while mitigation measures or controls have been put in place for others.
Measurement method	The metric is calculated by dividing the number of ICT systems for which risk treatment plans have been documented by the number of ICT systems in the organisation for which a risk assessment has been performed: $$RTPC = \frac{ICT\_systems\_with\_risk\_treatment\_plans}{ICT\_systems\_which\_undergone\_a\_risk\_assessment} * 100$$ This metric is expressed as percentage of ICT systems for which a risk treatment plan exists.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	RTPC values should trend higher over time and ideally reach 100%. A higher result would indicate that more ICT systems have a documented risk treatment plan. Most security process frameworks suggest or require risk treatment plans for ICT systems deployed in production environments. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for RTPC exists.
Reporting format	This metric can be used to evaluate their risk posture in terms of ICT systems. A better understanding of the quantity of ICT systems that have no risk treatment plan allows the organisation to evaluate their level of unknown risk associated with these ICT systems. With metric results for different dimensions, it is possible to identify and evaluate concentrations of risk, such as for results for critical ICT systems or ICT

systems containing confidential information.

1.1.1.13 Security testing coverage

Metric name	Security testing coverage
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	This metric indicates the percentage of the organisation's ICT systems have been tested for security risks.
Objective	This metric tracks the percentage of ICT systems in the organisation that have been subjected to security testing. Testing can consists of manual or automated white and/or black-box testing and generally is preformed on ICT systems post-deployment (although they could be in pre-production testing). Studies have shown that there is material differences in the number and type of ICT system weaknesses found. As a result, testing coverage should be measured separately from risk assessment coverage.
Measurement method	This metric is calculated by dividing the number of ICT systems that have had post-deployment security testing by the total number of deployed ICT systems in the organisation: $$STC = \frac{ICT\_systems\_which\_undergone\_security\_testing}{Total\_of\_deployed\_ICT\_systems} * 100$$ This metric is expressed as percentage of ICT systems that have undergone a security testing.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	STC values should trend higher over time. Generally, the higher the value and the greater the testing scope, the more vulnerabilities in the organisation's ICT systems will be identified. A value of 100% indicates that every ICT system has been subject to post-deployment testing. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for Security Testing Coverage exists.
Reporting format	This metric can be used to evaluate the degree to which ICT systems have been tested for weaknesses during the post-development phase (dimensions could be used to expand this metric to cover various stages of the development lifecycle). Quantifying the ICT systems not subjected to security testing allows the organisation to evaluate their application risk.

1.1.1.14 Security audit deficiencies

Metric name	Security audit deficiencies
Source	Not applicable
Description	This metric indicates the average number of deficiencies found in internal and external security audits performed in the past 12 months.
Objective	This metric tracks the deficiencies found in internal and external security audits performed in the past 12 months. It indicates how the level of security within ICT systems is maintained throughout its lifecycle, when systems are added and removed. For the metric to reflect reality (and the measurements to be comparable), the security audits must be performed applying consistent standards and procedures.
Measurement method	This metric is calculated by averaging the number security deficiencies found in the past 12 months over the number of audits performed. $$SAD = \frac{\sum_i Security\_deficiencies\_found\_in\_audit_i}{Total\_of\_audits\_performed\_in\_past\_12\_months} * 100$$ This metric is expressed as the average number of security deficiencies found during audits occurred in the past 12 months. The metric does not include deficiencies found during audits older than 12 months (even if they are still unresolved).
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	SAD values should trend lower over time. The lower the value (while maintaining at least an equal testing scope), the less deficiencies were detected in the organisation's ICT systems.
Reporting format	This metric can be used to evaluate the degree to which ICT systems contain deficiencies detected during security audits (dimensions could be used to expand this metric to cover various stages of the development lifecycle). Organisations can choose to report the internal audit and external audit deficiencies as different categories.

1.1.1.15 Percent of ICT systems with business continuity plans

Metric name	Percent of the systems with business continuity plans
Source	This metric is based on [35] and the ISO27004 standard.
Description	Percent of ICT systems with Business Continuity plans measures the validity of the business continuity management of an organisation. The metric evaluates the number of ICT systems for which business continuity plans have been adequately (a) documented & (b) proven by suitable testing within the past 12 months.
Objective	Percent of ICT systems with Business Continuity plans measures the percentage of systems that have established Business Continuity plans which are documented and are compliant with the organisational standards. For those systems where appropriate Business Continuity plans are in place, the metric also verifies whether testing has occurred in the past 12 months to ensure that the plan is validated against a recent state of the system. If the plan has not been tested, there is a significant risk of incoherencies or omissions. Since the metric involves both the creation of new Business Continuity plans and the testing of the existing ones, the Percent of ICT systems with Business Continuity plans metric value will vary over time. Organisations can use this metric to gauge their relative level of risk when the business continuity is endangered.
Measurement method	Percent of ICT systems with Business Continuity plans is calculated by counting those ICT systems for which Business Continuity plans exist and were recently tested divided by the total number of ICT systems which exist in the organisation. $$PSBCP = \frac{Number\_of\_ICT\_Systems\_With\_Tested\_BC\_Plans}{Total\_Number\_of\_ICT\_Systems} * 100$$ The metric is unitless and expressed as a percentage.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	The metric values should trend higher over time. An ideal result would be 100%. It should be noted that, if no actions are taken on the business continuity plans, the metric value will can start decreasing (due to the 'suitable testing with the past 12 months' criterion). Higher values would generally result in better preparation against outages or

	other incidents. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values exists.
Reporting format	Reporting of the metric should be the percentage on a time-scale to show the evolution.

Service Delivery

This section will specify a number of resilience metrics that belong to the service delivery phase. Metrics in this dimension measure the difference in service level before, during and after the fault or challenge.

More specifically we present following metrics in more detail:

- Operational mean time between failures (see 1.1.1.16);
- Operational availability (see 1.1.1.17);
- Operational reliability (see 1.1.1.18);
- Fault report rate (see 1.1.1.19);
- Incident rate (see 1.1.1.20);
- Illegitimate network traffic (see 1.1.1.21);
- Percent of ICT systems without known severe vulnerabilities (see 1.1.1.22);
- Delay variation (jitter) (see 1.1.1.23);
- Packet loss (see 1.1.1.24);
- Bandwidth utilization (see 1.1.1.25).

Metrics that are not elaborated in this document but can be useful to measuring the service delivery phase of resilience are:

- Application-specific metrics (see 1.1.1.26);
- Percentage of ICT systems that are monitored, measured, managed and reported on 24x7 basis;
- Percentage of ICT systems assets that is covered by a maintenance contract;
- Percentage of ICT systems assets that is obsolete on a lifecycle basis;
- Percentage of ICT systems assets for which fault tolerance is implemented.

1.1.1.16 Operational mean time between failures

Metric name	Operational MTBF: mean time between failures
Source	This metric definition is adopted from the IEEE Standard Glossary of Software Engineering Terminology [39].
Description	Operational MTBF is a basic indicator of reliability for **fault tolerant** ICT systems. For obvious reasons the ability of the ICT system to recover from failures is a prerequisite here. Operational MTBF expresses the expected time between consecutive failures in an ICT system. It is important to note how a failure is defined: We define a failure as the transition from the normal service level to impaired or even unacceptable service level.
Objective	This metric indicates the predicted time between different failures of an ICT system during operation.
Measurement method	Operational MTBF is defined as the mean value of the length of time between consecutive failures, computed as the ratio of the cumulative observed time to the number of failures under stated conditions, for a stated period of time in the life of an item. It is calculated as the sum of the operational periods divided by the number of observed failures (the operational period is defined as the difference in time between the moment the service starts operating at the normal service level until the moment the service fails). Note that the duration of the failure has no impact on the metric value. $$OperationalMTBF = \frac{\sum_i operational\_periods_i}{number\_of\_failures}$$ Operational MTBF is reported as an absolute value in hours.
Frequency	Operational MTBF should be monitored on real-time basis.
Target values	Target values depend highly on the criticality of the service and the topology of the system. For example: If a service is very critical, the operational MTBF targets will be higher compared to a normal service. As an example, the operational MTBF target for an Internet service for large corporations will be higher than the target for Internet service for residential customers.
Reporting	Operational MTBF is reported as an absolute time value versus the target value

format	for different services.

1.1.1.17 Operational availability

Metric name	Operational availability
Source	This metric definition is based on the definitions from [12].
Description	Operational availability is defined as the percentage of time an ICT system is available to end users.
Objective	The goal of the metric is to indicate the observed availability, which is the probability that an ICT system is not failed or undergoing a repair action when it is requested for use.
Measurement method	Operational availability is calculated as the percentage of the mean time that an ICT system is running at the normal service level over the total time. Two intermediate concepts are introduced, needed for the calculation of the operational availability terms: • Mean Time Between Maintenance Actions (MTBMA): The mean time between maintenance actions (corrective and preventive maintenance). • Mean Down Time (MDT): The mean time that an ICT system is non-operational, including preventive/corrective maintenance actions. A more extended MDT definition can be found in 1.1.1.27. $$Operational\_availability = \frac{MTBMA}{MTBMA + MDT}$$ The unit of MTBMA and MDT should be the same (hours, seconds …) while the operational availability Is expressed as a percentage.
Frequency	Operational availability should be monitored on real-time basis.
Target values	Target values for operational availability are impossible to specify for a generic ICT system. They are specified in the service level specification of the service provider. The difference between the operational availability and the availability as specified in service level specification should be monitored.
Reporting format	Operational availability is measured in a predefined time window. For example: 99,9% operational availability measured on a yearly basis allows for a consecutive unavailability of 8,76 hours whereas the same operational availability in a measurement window of 1 month would only allow for 0,744 hours of consecutive service unavailability.

	Availability reporting is done in function of the measurement window (e.g. reporting of the availability per month for all months of the year).

1.1.1.18 Operational reliability

Metric name	Operational reliability
Source	This metric definition is based on [40].
Description	The operational reliability of an ICT system is the ability of to perform its required functions under stated conditions (i.e. operate at the normal service level) for a specified period of time.
Objective	Reliability indicates the probability that an ICT system will perform its required function for a specific period of time t, referred to as 'mission time'. Calculating operational reliability includes a dimension of mission time for calculating the results (this is not the case for availability, where only the probability of the ICT system being available for end-users at a certain moment in time is calculated).
Measurement method	The operational reliability of a system is a function of the Operational Mean Time between Failures' (MTBF) and a mission time t. Mission time is defined as the time between the time where the ICT system starts operating at the normal service level and the time at which the ICT system fails. Failure is defined as functioning below the acceptable service level. The expected reliability R(t) is modelled with the exponential distribution, which describes random failures: $$R(t) = e^{-t / OperationalMTBF}$$ The probability R(t) indicates the probability that an ICT system will run for a specified mission time 't'. Operational MTBF and mission time t is specified in the same time dimension, i.e. hours, seconds, days... The operational MTBF and mission time t have the same unit of time measurement (e.g. hours, years …), while expected reliability is expressed as a unitless probability.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	Target values depend highly on the criticality of the service and the topology of the ICT system. However, as soon as the metric is below e^(-1) (= 0,3678 = 1/e), the ICT system or service has been running longer than the mean time between failure: This means, on average, the service would have encountered a failure and failure has become more imminent.

Reporting format	Reliability should be monitored on a monthly basis. The figure below shows the expected reliability curve.

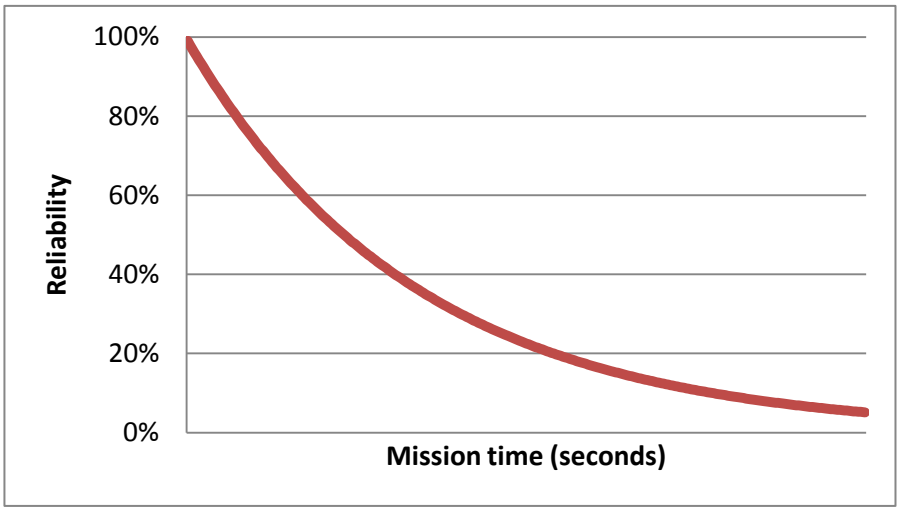

Figure 14: Operational reliability curve

1.1.1.19 Fault report rate

Metric name	Fault report rate
Source	Not applicable
Description	The fault report rate metric measures of the number of faults occurring in a given time period.
Objective	The fault report rate indicates the number of detected faults during the metric time period. In combination with other metrics, such as operational availability, this can indicate the degree to which the ICT system can overcome occurring faults and maintain the normal service level.
Measurement method	To calculate the fault report rate metric, the number of faults in a given time period are counted. Additional grouping could occur per category or organisational departments for example. $$Fault\_Re\ port\_Rate = \frac{Amount\_of\_faults\_per\_category}{Length\_of\_time\_window}$$ The time window is expressed as an absolute unit of time (e.g. hours or days) while the number of faults is an absolute number, indicating how many faults have occurred in the past time window. The fault report rate is expressed as faults per time period.
Frequency	The fault reporting and follow-up should happen on a continuous basis and at least daily.
Target values	No specific target can be set, as the metric value will depend on the categories of the faults that are taken into account in this metric. A target should be set on the variation of faults that occur (to trigger alarms). For example, the ratio of faults occurring per time window versus the number of ICT systems should be closely monitored and should be almost constant.
Reporting format	Reporting of the fault report rate should be per category and in a time-series plot. For example, faults can be categorized according to severity or impact. This allows for a more granular overview of the different faults and a better understanding of the different fault types occurring. An example of a reporting with a fault categorization and an assigned priority for resolution is shown below.

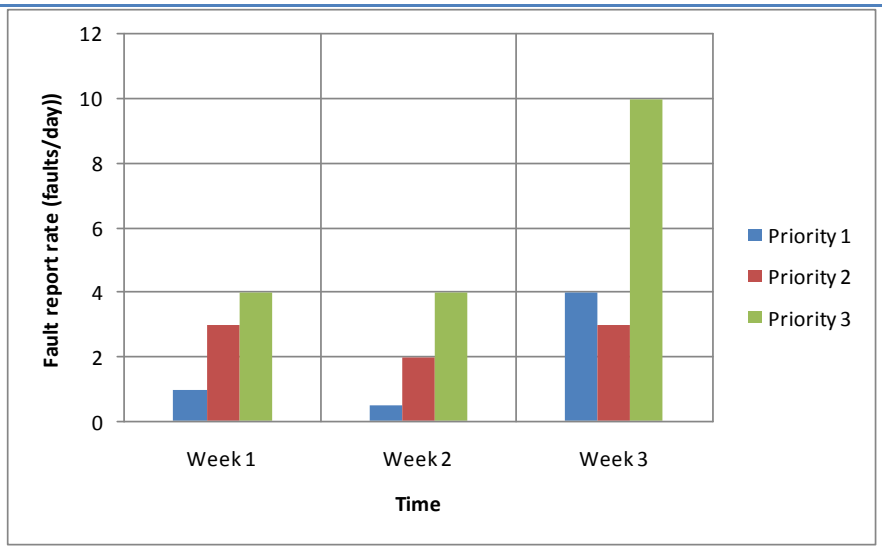

Figure 15: Fault report rate reporting example

1.1.1.20 Incident rate

Metric name	Incident Rate
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	The incident rate metric measures the number of security incidents that occur in a given time period from selected incident categories.
Objective	The incident rate indicates the number of detected security incidents the organisation has experienced during the metric time period. In combination with other metrics, this can indicate the level of threats, the effectiveness of security controls and/or incident detection capabilities.
Measurement method	To calculate the incident rate metric, the number of security incidents in a given time period are counted, additional grouping could occur per incident category or organisational departments for example. $$Incident\_Rate = \frac{Amount\_of\_incidents\_per\_category}{Length\_of\_time\_window}$$ The time window is expressed as an absolute unit of time (e.g. hours or days) while the number of incidents is an absolute number, indicating how many incidents have occurred in the past time window. **Note:** In a network of ICT security systems, it is possible that each security device reports an attack at the very same time, although only one attack is ongoing (for example: an incident on the outer firewall and an incident on the IDS system can indicate the very same event). This can result in a skewed view of the amount of incidents that occurs on the network.
Frequency	The incident management and follow-up should happen on a continuous basis and at least daily.
Target values	No specific target can be set, as the metric will also depend on the categories of incidents that are taken into account in this measure. A target should be set the variation of incidents that occur (to trigger alarms). Incident rate values should trend lower over time – assuming perfect detection capabilities. The value of "0" indicates hypothetical perfect security since there were no security incidents. Because of the lack of experiential data from the field, no consensus on range of acceptable goal values for Incident Rate exists.
Reporting	Reporting of the incident rate should be per category and in a time-series plot.

format	Example of a reporting format with an incident categorization per incident priority:

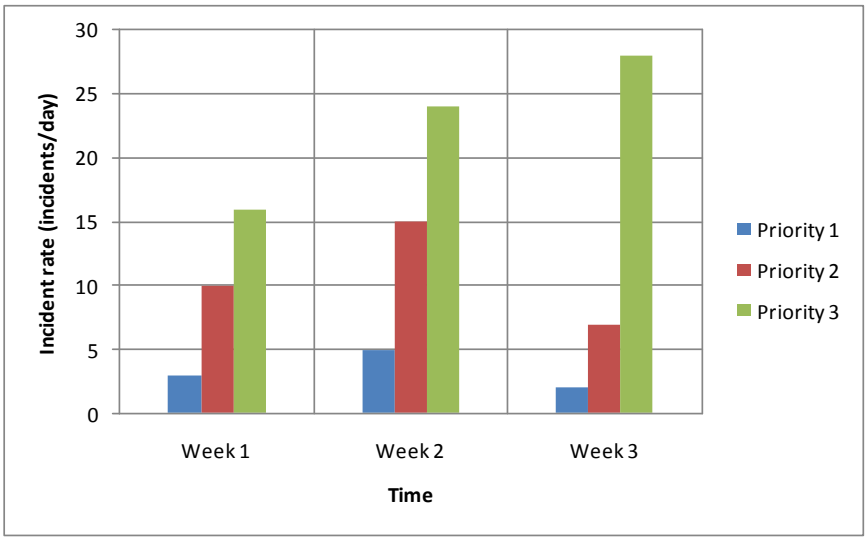

Figure 16: Sample incident rate report

1.1.1.21 Illegitimate network traffic

Metric name	Illegitimate traffic
Source	Not applicable
Description	The illegitimate traffic metric measures the ratio of malicious, spam or unauthorized traffic versus all traffic on the network. A high metric value indicates an increased presence of malicious entities or infected ICT systems on the network.
Objective	The metric indicates the resistance against unauthorized traffic that tries to enter on the network.
Measurement method	Different measures of illegitimate traffic are available, depending on the category of illegitimate traffic: • **Spam traffic**: Observed spam messages divided by the total number of e-mail messages during a specific timeframe. This metric can be extracted from the anti-spam defence systems of an organisation (if installed). The metric unit is number of spam messages per day and % of spam messages on the total amount of mail messages received (within a predefined time period). $$Spam\_Traffic = \frac{Amount\_of\_spam\_messages\_received}{Total\_amount\_of\_messages\_received}$$ • **Observed malicious and unauthorized traffic**: By using network anomaly detection systems that recognise certain Command-and-Control botnet traffic or illegal protocols on the network, the offending entities can be singled out and the source of the illegitimate traffic can be taken away. Measurement can be done by calculating the ratio of malicious traffic to the total traffic on the network, based on certain places in the network (examples: inside a company's DMZ or at the external firewall). Another metric of malicious traffic is to take the number of offending hosts on the network, compared to the total number of hosts. The metric unit is a percentage of malicious traffic or hosts on the total amount of malicious traffic or hosts received (within a predefined time period). $$Malicious\_Traffic = \frac{Amount\_of\_malicious\_traffic\_received}{Total\_amount\_of\_traffic\_received}$$
Frequency	Weekly, Monthly, Quarterly, Annually

Target values	Defining a target value is difficult: There is a lack of experimental data to base the target values on, For spam messages, a possible reference to help setting target values is the SenderBase e-mail scanning system: the SenderBase e-mail reputation platform reports 85% as the average percentage spam in all e-mail traffic seen by mail scanner systems around the globe at the time of writing.[12]
Reporting format	Reporting of the illegitimate traffic should be per category. Time-plots of the spam traffic received and the malicious traffic stream rates can give an indication on the infected systems.

[12] *http://www.senderbase.org/home/detail_spam_volume*

1.1.1.22 Percent of ICT systems without known severe vulnerabilities (PSWKSV)

Metric name	Percent of systems without known severe vulnerabilities
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	Percent of ICT systems without known severe vulnerabilities (PSWKSV) measures the organisation's relative exposure to known severe vulnerabilities. The metric evaluates the percentage of ICT systems scanned that do not have any known high severity vulnerabilities.
Objective	Percent of ICT systems without known severe vulnerabilities (PSWKSV) measures the percentage of ICT systems that when checked were not found to have any known high severity vulnerabilities during a vulnerability scan. Vulnerabilities are defined as "High" severity if they have a CVSS[13] base score of 7.0-10.0. Since vulnerability management involves both the identification of new severe vulnerabilities and the remediation of known severe vulnerabilities, the percentage of ICT systems without known severe vulnerabilities will vary over time. Organisations can use this metric to gauge their relative level of exposure to exploits and serves as a potential indicator of expected levels of security incidents (and therefore impacts on the organisation). This severity threshold is important, as there are numerous informational, local, and exposure vulnerabilities that can be detected that are not necessarily material to the organisation's risk profile. Managers generally will want to reduce the level of noise to focus on the greater risks first. This metric can also be calculated for subsets of systems, such as by asset criticality of business unit
Measurement method	Percent of ICY systems without known severe vulnerabilities is calculated by counting those systems that have no open high severity level vulnerabilities, measured against a vulnerability database (for example, the NIST's National Vulnerability Database or an organisation's internal vulnerability database). A severe vulnerability could be defined as a vulnerability where the Vulnerability Status !='Open' & CVSS Base Score >= 7.0. This result is then divided by the total number of systems in the scanning scope.

[13] *http://www.first.org/cvss/cvss-guide.html*

	$$PSWKSV = \frac{Number\_of\_ICT\_Systems\_Without\_Known\_Severe\_Vulnerabilities}{Number\_of\_Scanned\_ICT\_Systems} * 10$$ The metric is unitless and expressed as a percentage.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	PSWKSV values should trend higher over time. It would be ideal to have no known severe vulnerabilities on systems; therefore, an ideal target value would be 100%. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for Percent of Systems Without Known Severe Vulnerabilities exists.
Reporting format	Different categorizations are useful [47], including: • Vulnerability score by operating system, application, or organisation division – this metric provides a high level measurement of how the organisation is doing, cut across several dimensions. • Most vulnerable applications, with a breakdown into vulnerability score by application version – this metric helps highlight old, vulnerable versions of software that should be upgraded or eliminated. • ICT systems scanned within the last 'X' days – this metric shows how many ICT systems is being scanned in a timely fashion. • Unowned devices and unapproved applications – this metric is very useful to track 'unowned' devices that may be rogue devices or simply contractor/consultant systems, as well as the trend of applications that are not specifically allowed on the network.

1.1.1.23 Delay variation (jitter)

Metric name	Delay variation (jitter)
Source	This metric definition is based on [48].
Description	One-way delay variation (jitter) is usually introduced in network nodes as an effect of queuing. Delay variation or jitter represents the variation in one-way delay (latency).
Objective	One-way delay variation is a metric that describes the level of disturbance of packet arrival times with respect to an 'ideal' pattern, typically the pattern in which the packets were sent. Such disturbances can be caused by competing traffic (i.e. queuing), or by contention on processing resources in the network.
Measurement method	Jitter is calculated as the average difference between the one-way delays for a selected pair of packets [49] and is expressed in milliseconds (ms).
Frequency	Sub-hourly measurements are needed to monitor the right service level
Target values	Time-sensitive traffic such as voice and video has a maximum jitter of 1 ms. [50]. This value above is based on application-based requirements (for voice data) but does not indicate a target level for the network resilience. Instead, resilience based on the one-way delay can be measured in several ways: • Either as the metric being close to zero (if the jitter is constant, the network is considered to be resilient to challenges and faults in the network). A network that has fluctuating jitter values should be considered to be less resilient to challenges and faults on the network. • Another way of setting resilience targets is how many service users experience the jitter with respect to the service level specification for how much % of service time. For example, a target can be that 99% of all service users have sub-1 ms jitter for 99,9% of the time for the voice traffic class.
Reporting format	Service providers typically provide network links with Quality of Service enabled: This means that the service provider guarantees different service levels per traffic class. In order to make correct measurements, packets of each traffic class that has been agreed with the provider must be sent to have a correct performance overview. Reporting is typically done per traffic class and plotted as a time-series to show the evolution.

1.1.1.24 Packet loss

Metric name	Packet loss
Source	This metric definition is based on [48].
Description	Packet loss is the probability of a packet to be lost in transit from a source to a destination. There are two main reasons for packet loss: **Congestion:** When the offered load exceeds the capacity of a part of the network, packets are buffered in queues. Since these buffers are also of limited capacity, severe congestion can lead to queue overflows, which lead to packet drops. Severe congestion could mean that a moderate overload condition holds for an extended amount of time, but could also consist of the sudden arrival of a very large amount of traffic (burst). **Errors:** Another reason for loss of packets is corruption, where parts of the packet are modified in-transit. When such corruptions happen on a link (due to noisy lines etc.), this is usually detected by a link-layer checksum at the receiving end, which then discards the packet. The upper-layer protocols (UDP/TCP) are responsible for re-transmitting the packets.
Objective	This metric is an important indicator for Voice-over-IP / video conferencing traffic quality, as voice/video traffic is very sensitive to delay and retransmission of lost packets causes delays. The result of packet loss is usually degradation in sound or image quality. Packet loss is an indicator for network disturbance and how resilient a network is when experiencing congestion or errors. A low packet loss metric indicates a high resilience against faults and challenges.
Measurement method	Packet loss can be actively measured by sending a set of packets from a source to a destination and comparing the number of received packets against the number of packets sent. Ping (ICMP echo) is an example of a tool that implements this procedure.
Frequency	Sub-hourly measurements are needed to monitor the right service level
Target values	Target values are very dependent on application requirements. For example, time-sensitive traffic such as voice and video have a maximum packet loss of 3%. [50] To measure resilience based on the packet loss, it can be measured in several

	ways: • Either as the metric being close to zero (if the packet loss is close to 0, the network is considered to be resilient to challenges and faults in the network). A network that has fluctuating packet loss values should be considered to be less resilient to challenges and faults on the network. • Another way of setting resilience targets is how many service users experience the packet loss with respect to the service level specification for how much % of service time. For example, a target can be that 99% of all service users have a maximum of 1% packet loss for 99,9% of the time for the voice traffic class.
Reporting format	Service providers typically provide network links with Quality of Service enabled. This means that the service provider guarantees different service levels per traffic class. In order to make correct measurements, packets of each traffic class that has been agreed with the provider must be sent to have a correct performance overview. Reporting is typically done per traffic class and plotted as a time-series to show the evolution.

1.1.1.25 Bandwidth utilization

Metric name	Bandwidth utilization: peak, average, variance
Source	Not applicable
Description	Bandwidth or channel utilization or throughput is the rate of message delivery over a communication channel. Bandwidth utilization can be measured in absolute figures or as a percentage of the maximum system bandwidth. The following types of bandwidth measures are relevant to the resilience of a system: • **Peak**: Peak bandwidth utilization is maximum throughput measured by a system within a specified timeframe. The value is the throughput measured over a short period of time. This is an indicator of the maximum bandwidth consumed within a certain period. • **Average**: The average channel utilization, also known as bandwidth utilization efficiency, in percentage is the achieved throughput related to the maximum throughput of a digital communication channel. For example, if the throughput is 5.5 Mbit/s in a 10 Mbit/s network connection, the channel utilization is 55%. • **Variance**: The arithmetic average of the squared distance from the mean bandwidth throughput. Typically, bandwidth to external parties such as WAN-traffic is contracted in the Service Level Specification as a Committed Information Rate (CIR). This is the bandwidth that is always guaranteed by the service provider. The CIR should be considered as the normal level of service contracted between the client and the provider. Any bandwidth utilization that is below the CIR and is not caused by the service provider's client not sending traffic is considered as a degraded service of the network.
Objective	Bandwidth measurements indicate the performance of the network. However, if bandwidth measurements would result in sub-average values, it does not always indicate a problem: it could be that the network is working fine but the devices on it are not sending or receiving data. When the network encounters challenges or faults, a low variation of bandwidth and constant average bandwidth indicate a high level of resilience.
Measurement method	Bandwidth utilization can be measured on 3 aspects: • Peak bandwidth utilization: Highest throughput measured between 2 network nodes in a predefined time window. • Minimum bandwidth utilization: Lowest throughput measured between 2

	network nodes in a predefined time window. • Average bandwidth utilization: Average throughput measured between 2 network nodes in a predefined time window. The metric unit is bits/second, kilobits/second (1x10^3 bits/s), megabits/second (1x10^6 bits/s) or gigabits/second (1x10^9 bits/s).
Frequency	Sub-hourly measurements are needed to monitor the right service level
Target values	Bandwidth is dependent on the needs of the entity for its network. It is very hard to give generic values for bandwidth.
Reporting format	Service providers typically provide network links with Quality of Service enabled: This means that the service provider guarantees different service levels per traffic class. In order to make correct measurements, packets of each traffic class that has been agreed with the provider must be sent to have a correct performance overview. Reporting is typically done per traffic class and plotted as a time-series to show the evolution.

1.1.1.26 Application-specific metrics

Depending of the criticality of a certain network service for an organisation, application-specific resilience metrics can prove to be very useful.

Examples of these applications are:

- Voice-over-IP
- HTTP traffic
- E-mail

As an example, a metric for speech quality in Voice-over-IP environments can be defined called 'Speech Quality'.

Speech quality measures instantaneous voice quality when Voice-over-IP traffic is sent over networks. It quantifies the quality by standardized ITU-T methods, based on a five-point category-judgement scale (more information can be found in [51]). ITU-T uses the term "Mean Opinion Score" (MOS) as the metric name.

The MOS metric indicates the resilience of the voice traffic to degraded network conditions (In an optimal situation, the speech quality should be as high and constant as possible in face of varying network conditions).

Recovery phase

This section will specify a number of resilience metrics that belong to the recovery phase. Metrics in this dimension revolving about how fast a service/network can recover from faults or challenges.

More specifically we present following metrics in more detail:

- Mean down time (see 1.1.1.27);
- Mean time to repair (see 1.1.1.28);
- Maintainability (see 1.1.1.29);
- Main time to incident recovery (see 1.1.1.30).

Metrics that are not elaborated in this document but can be useful to measuring the recovery phase of resilience are:

- Mean Time to Vulnerability Mitigation;
- Ratio of successful attempts to execute recovery.
    - By type of recovery (fix vs. rollback);
    - By service type;
    - By recovery method: Automatic, semi-automatic or manual.

1.1.1.27 Mean down time

Metric name	MDT: Mean down time
Source	This metric definition is based on [51].
Description	Mean down time (MDT) is the average time that an ICT system is non-operational. This includes all non-operational time associated with repair, corrective and preventive maintenance and includes any logistical or administrative delays. The difference between MDT and MTTR (mean time to repair) is that MDT includes any and all delays involved; MTTR looks solely at repair time.
Objective	This metric indicates the average time between the occurrence of a failure to the restoration of the normal service level. A higher value would indicate that a failure is likely to impact the service for a longer time, hence indicating a lower resilience (lower resistance to faults and challenges).
Measurement method	MDT is the total non-operational time divided by the total number of outages during a given period of time. $$MDT = \frac{\sum_i (Non\_Operational\_Time_i)}{Number\_of\_outages}$$ MDT is expressed as an absolute value in seconds or hours.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	No specific target values can be given, as this is highly specific per organisation.
Reporting format	Reporting of the Mean down time should be per category and in a time-series plot.

1.1.1.28 Mean time to repair

Metric name	Mean time to repair
Source	This metric definition is based on [39].
Description	The expected or observed time required to repair an ICT system and return it to normal operations. The difference between MDT and MTTR is that MDT includes any and all delays involved; MTTR looks solely at repair time.
Objective	This metric indicates the average time from start of the repair until the return to the operational state of an ICT system. A high value indicates that repair times are on average long and the service down time will be longer.
Measurement method	MTTR is the total corrective maintenance time divided by the total number of corrective maintenance actions during a given period of time. $$MTTR = \frac{\sum_i (Maintenance\_time_i)}{Number\_of\_maintenance\_actions}$$
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	No specific target values can be given, as this is highly specific per organisation.
Reporting format	Reporting of the MTTR should be per category and in a time-series plot.

1.1.1.29 Maintainability

Metric name	Maintainability
Source	This metric definition is based on [39].
Description	Maintainability is the ease with which an ICT system can be modified to correct faults, improve performance or other attributes, or adapt to a changed environment. It is defined as the probability of performing a successful repair action within a given time t. For example, if a particular ICT system has a 90% maintainability for a repair time t equal to 1 hour, this means that there is a 90% probability that the component will be repaired within 1 hour.
Objective	Maintainability indicates the probability that an ICT system will be repaired within a time t, referred to as 'maximum repair time'. Calculating maintainability includes a dimension of maximum repair time for calculating the results. A high maintainability metric value indicates that the ICT system can be easily and quickly restored to operational status. It does not provide a metric for measuring the resilience of an ICT system in an operational state but rather indicates the speed in which an ICT system can be restored to an acceptable level of service.
Measurement method	The maintainability of an ICT system is a function of a variable Mean Time To Repair (note that the MTTR value is updated every time a failure occurs) and repair time t. It is modelled using the exponential distribution, which describes random repair time: $$M(t) = 1 - e^{-t/MTTR}$$ The metric is expressed as a percentage.
Frequency	Weekly, Monthly, Quarterly, Annually
Target values	For a specific desired repair time t (agreed with the customer in the service level specification), the maintainability target value is 100% (meaning a certitude of reparation can be given as the value of the metric depends on the repair time that has already elapsed). The value of 100% is unattainable (cfr. the formula and Figure 17). An example of a stated maintainability goal is a 90% probability that

	maintenance repair times will be completed in 8 hours or less with a maximum repair time of 24 hours. This requires an ICT system's MTTR of 3.48 hours.
Reporting format	As an example, we plotted the maintainability for varying maximum repair times t where the MTTR is 100 seconds.

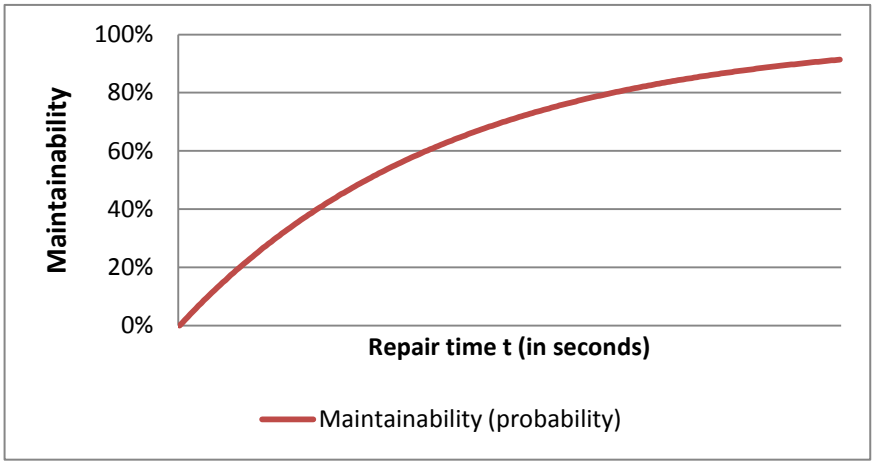

Figure 17: Maintainability curve

1.1.1.30 Mean time to incident recovery

Metric name	MTTIR: Mean time to incident recovery
Source	This metric is adopted from 'The CIS security metrics - Consensus Metric Definitions v1.0.0' [5].
Description	Mean time to incident recovery (MTIR) characterizes the ability of the organisation to return to a normal state of operations. This is measured by the average elapse time between when the incident occurred to when the organisation recovered from the incident.
Objective	Mean time to incident recovery measures the effectiveness of the organisation to recovery from security incidents. The sooner the organisation can recover from a security incident, the less impact the incident will have on the overall organisation.
Measurement method	MTTIR is measured by dividing the average elapsed time between the incident occurrence and the recovery to normal service level over the number of incidents. This calculation can be averaged over a time period $$MTTIR = \frac{\sum_i (Date\_of\_Recovery_i - Date\_of\_Occurrence_i)}{Number\_of\_incidents}$$ Unit of the metric is a time over the number of incidents, for example hours/incident.
Frequency	Weekly, Monthly, Quarterly, Annually.
Target values	MTTIR values should trend lower over time. There is evidence the metric result will be in a range from days to weeks (2008 Verizon Data Breach Report). The value of '0' indicates hypothetical instantaneous recovery. Because of the lack of experiential data from the field, no consensus on the range of acceptable goal values for Mean Time to Incident Recovery exists.
Reporting format	Reporting of the incident recovery time should be per category and based on the hours/incident value.

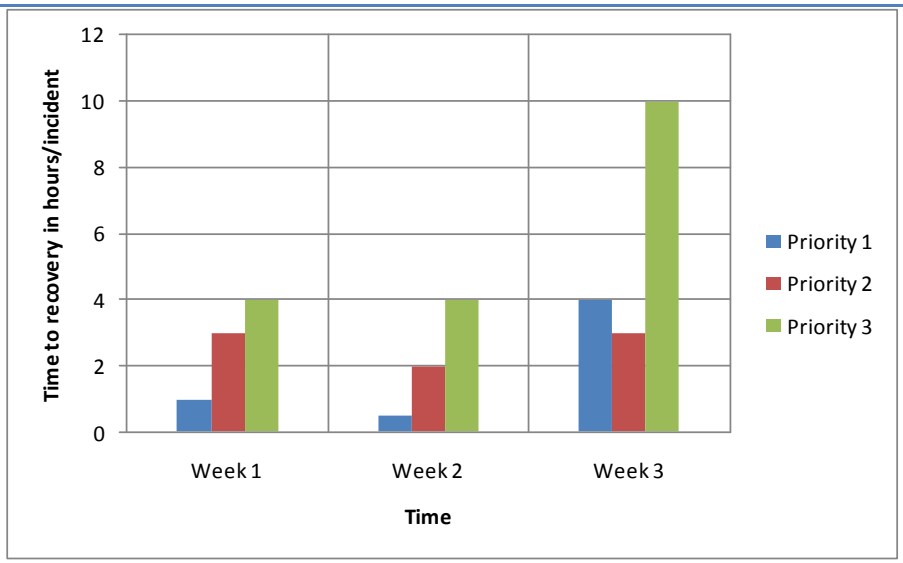

Figure 18: Sample Time to recovery report

Design-based metrics

Referring to the time-based event-oriented classification (section 0), a distinction can be made between the 'preparedness' phase on one side and the service delivery and recovery phase on the other side.

Preparedness includes all the actions and measures taken to prevent an incident from happening or to diminish the impact of the incident to the service level. In other words, in the preparedness phase, the aim is to measure how well a system is prepared to faults and challenges.

Section 0 describes the different preparedness metrics identified in this document, it is important to note that there also exist a number of theoretical counterparts to the preparedness metrics – in this report, they are called preparedness indicators or design-based resilience metrics (examples are given below in this paragraph). To reflect the difference with their theoretical counterpart, the names of these theoretical metrics are preceded by the adjective 'operational' (for example: operational availability).

These theoretical indicators are, as described in this section, calculated during the design phase of the ICT systems and are time-independent. Therefore, we consider the all indicators in the preparedness phase to be **design-based resilience indicators** while the service delivery and recovery phases are true metrics, which are to be measured during the operation of the service and which are time-dependent.

A number of these theoretical performance indicators are presented in this paragraph. These indicators are expressed as calculated probabilities. It is important to note that, while the calculation here is based on theoretical numbers and probabilities, they do present value in describing resilience.
For example: a network that has redundant path will have a higher Mean Time between failures as a system, compared to a network with the same components but without redundant paths.

Below, these theoretical performance indicators are presented in their theoretical interpretation. No target values nor measuring frequencies are provided, as the values of these indicators depend on the design of the system/server but are not time-dependent. Given the static and theoretical character of these indicators, measurement frequency is a term that does not apply.

More specifically we present following indicators in more detail:

- Expected mean time between failures (see 0);
- Expected availability (see 0);
- Expected reliability (see 0);
- Link/node failure (see 0).

This section will conclude by illustrating metrics on example topologies.

Expected mean time between failures

Indicator name	Expected MTBF: Mean Time Between Failures
Source	This metric definition is adopted from the IEEE Standard Glossary of Software Engineering Terminology [39].
Description	Expected MTBF is a basic indicator of reliability for **fault tolerant** ICT systems. For obvious reasons the ability of the ICT system to recover from failures is a prerequisite here. Expected MTBF expresses the expected time between consecutive failures in an ICT system. It is important to note how a failure is defined: We define a failure as the transition from the normal service level to impaired or even unacceptable service level. **Important note:** A clear distinction should be made between the Expected MTBF and the Operational MTBF as defined in section 1.1.1.16. The difference between these metrics is the source data that is used: While the operational MTBF uses historical data, the expected MTBF uses vendor-provided statistics on the MTBF of the equipment to calculate the MTBF of a certain network service.
Objective	This metric indicates the predicted time between different failures of an ICT system during operation.
Calculation method	Expected MTBF is defined as the mean value of the length of time between consecutive failures, computed as the ratio of the cumulative observed time to the number of failures under stated conditions, for a stated period of time in the life of an item. It is calculated as the sum of the operational periods divided by the number of observed failures (the operational period is defined as the difference in time between the moment the service starts operating at the normal service level until the moment the service fails). Note that the duration of the failure has no impact on the metric value. $$ExpectedMTBF = \frac{\sum_i operational\_periods_i}{number\_of\_failures}$$ For hardware components, expected MTBF is usually a technical specification provided by the equipment vendor. If not provided, it can be empirically determined by measuring and averaging the mission time of a service in a controlled environment. There are two different ways to calculate the Expected MTBF: Government programs use calculations per the latest version of MIL-

	HDBK-217 [45], while commercial programs use the Telcordia SR-332 method [46]. For ICT systems, expected MTBF can be calculated from the expected MTBF values of the different components, depending on the redundancy in the topology and the type of connections (series/parallel). These calculations are complex and are beyond the scope of this document. The formulas are documented in [41]. Expected MTBF is reported as an absolute value in hours. An MTBF calculation may result in an anticipated failure rate of once every year, but it should be clear that MTBF is an average. An MTBF of once for each ten years could also mean twice in five years, or two failures in the first few weeks of operation, with correct operation for the remaining years.
Frequency	This is a design-based indicator and is time-independent: It does not need to be measured but is rather calculated during the design phase.
Target values	Target values depend highly on the criticality of the service and the topology of the ICT system. For example: If a service is very critical, it could either be built by a few components with a very high expected MTBF (thus very reliable) or by multiple redundant components that could have a lower expected MTBF (if in this case, failure of a component is compensated by an active redundant component).
Reporting format	The expected MTBF of an ICT system is calculated during the design phase and should be recalculated when components are added or removed or changes in topology occur. Other than that, the expected MTBF values are static and do not require periodic reporting.

Expected availability

Indicator name	Expected availability
Source	This metric definition is adopted from the IEEE Standard Glossary of Software Engineering Terminology [36].
Description	Expected availability is indicative for both the reliability (how long will an ICT system run without failures) and maintainability (if the ICT system breaks down, how easy is it to repair) properties of an ICT system. Reliability expresses how long an ICT system will run without failures, while maintainability indicates how easy it is to repair a system.
Objective	The goal of the metric is to indicate the probability that the ICT system is operating properly when it is requested for use. That is, expected availability is the probability that an ICT system is out of service when it needs to be used.
Calculation method	It is defined using the already introduced Expected Mean Time Between Failure (Expected MTBF) and Mean Time To Repair (MTTR) metrics. Just to shortly recall: The Expected MTBF is defined as the average time that an ICT system can operate flawlessly between 2 failure events. The MTTR is the time that it costs to repair a failed ICT system or system component after a failure event. Expected availability is calculated as: $$ExpectedAvailability_{COMPONENT} = \frac{ExpectedMTBF_{COMPONENT}}{ExpectedMTBF_{COMPONENT} + MTTR_{COMPONENT}}$$ The Expected availability of the ICT system can be calculated from the component expected availability values, depending on the redundancy in the topology and the type of connections (series/parallel). These calculations become quickly very complex and specific software tools are available to assist in these calculations. The formulas are documented in [41]. The unit of expected MTBF and MTTR should be the same (hours, seconds …) while the expected availability is expressed as a percentage.
Frequency	This is a design-based indicator and is time-independent: It does not need to be measured but is rather calculated during the design phase.

Target values	Target values depend highly on the criticality of the service and the topology of the ICT system.
Reporting format	The expected availability of an ICT system is calculated during the design phase and should be recalculated when components are added or removed or changes in topology occur. Other than that, the expected availability values are static and do not require periodic reporting.

Expected reliability

Indicator name	Expected reliability
Source	This metric definition is based on [40].
Description	The expected reliability of an ICT system is the ability of to perform its required functions under stated conditions (i.e. operate at the normal service level) for a specified period of time.
Objective	Expected reliability indicates the probability that an ICT system will perform its required function for a specific period of time t, referred to as 'mission time'. Calculating expected reliability includes a dimension of mission time for calculating the results (this is not the case for availability, where only the probability of the system being available for end-users at a certain moment in time is calculated).
Calculation method	The expected reliability of an ICT system is a function of the Expected Mean Time between Failures (Expected MTBF) and a mission time t. Mission time is defined as the time between the time where the service starts operating at the normal service level and the time at which the service fails. Failure is defined as functioning below the acceptable service level. The expected reliability R(t) is modelled with the exponential distribution, which describes random failures: $$R(t) = e^{-t/ExpectedMTBF}$$ The probability R(t) indicates the probability that an ICT system will run for a specified mission time 't'. Expected MTBF and mission time t is specified in the same time dimension, i.e. hours, seconds, days,... The expected MTBF and mission time t have the same unit of time measurement (e.g. hours, years, …), while expected reliability is expressed as a unitless probability.
Frequency	This is a design-based indicator and is time-independent: It does not need to be measured but is rather calculated during the design phase.
Target values	Target values depend highly on the criticality of the service and the topology of the ICT system. However, as soon as the metric is below e^{-1} (= 0,3678 = 1/e), the network or service has been running longer than the mean time between failure: This means, on average, the service would have encountered a failure and failure has become more imminent.

Reporting format	Reliability should be monitored on a monthly basis. The figure below shows the expected reliability curve.

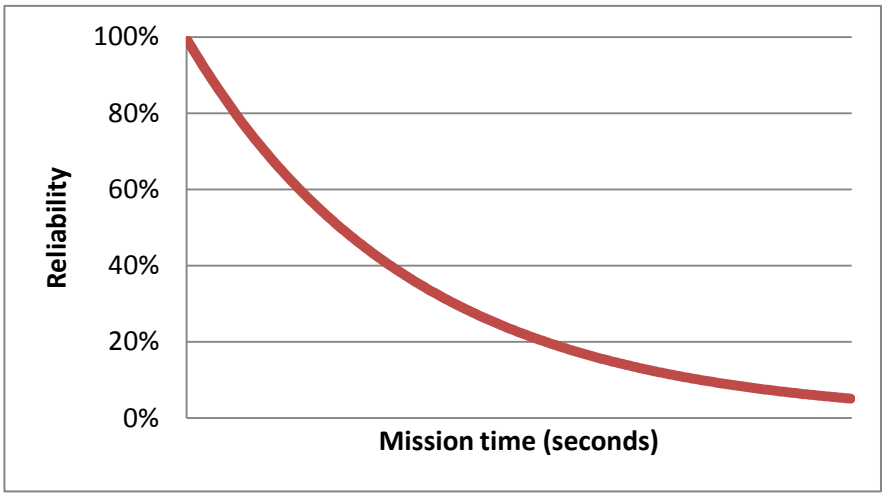

Figure 19: Expected reliability curve

Link/node failure

Indicator name	Link/node failure
Source	This metric definition is based on [53].
Description	Link/node failure is an indicator for the robustness of a network to link and/or network nodes failures.
Objective	The resilience of a network is expressed by investigating the change of a specific network performance indicator (e.g. bandwidth or packet loss) in value over time, when the network system is exposed to challenges. In this indicator, the challenges are the (partial) failure of a link, node or specific component within each node. It is an indicator for the robustness against stress of the network topology.
Calculation method	The link/node failure indicator is expressed as a network performance parameter (bandwidth, packet loss...) in function of the number of links, network nodes or components of the network nodes that are removed. This indicator cannot be calculated: Data must be collected either empirically or via simulation of the network topology. The data is collected by varying the number of links removed while measuring the network response parameter. During this process, a number of performance curves will be developed. After all possible combinations have been tested, the best, worst and average case curves will be determined. After the data collection, an 'envelope' can be determined which is confined by the best case and worst case curves. It is important to note that this envelope determines the upper and lower boundaries of the performance impact for a given number of link/node failures. Using the envelope, the effect of various failures can be shown visually and resilience against network degradation is done by comparing metric envelopes. Envelopes can be developed for a multitude of challenges (random number of challenges, fire, misconfiguration, earthquake, intentional attacks).
Frequency	This is a design-based indicator and is time-independent: It does not need to be measured but is rather calculated during the design phase.
Target values	Target values depend highly on the criticality of the service and the topology of the system.
Reporting	Reporting can be done either via the envelope figures.

format

The example shown below uses the number of links removed as the challenge parameter in function of an unspecified performance parameter m.

The figures are taken from [53].

The envelope is developed by measuring the impact on performance parameter m by removing a number of links. Every measurement is shown as a separate curve (Figure 20).

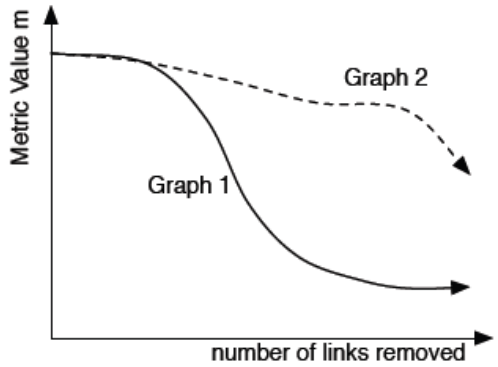

Figure 20: Empirical determination of the impact of the metric value m by removing links

After every possible combination of link failures has been measured, the best, average and worst case curves become apparent (Figure 21).

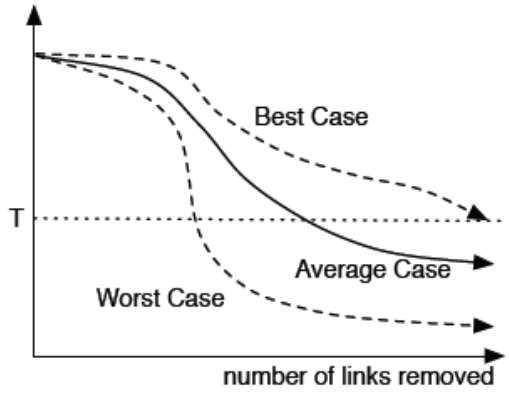

Figure 21: The best, average and worst case scenarios becomes apparent

Using the 3 curves as displayed in Figure 21, the envelope can be determined. The envelope can be used to determine the upper and lower boundaries of the impact on performance parameter m for a specific number of link/node failures.

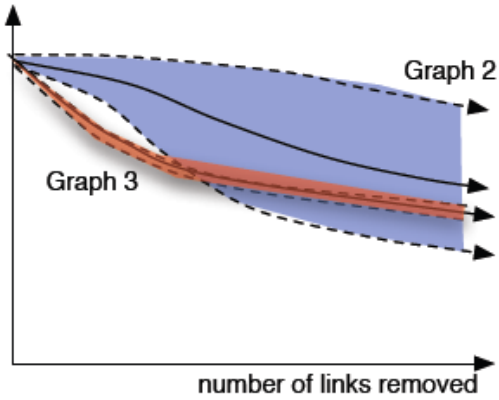

Figure 22: The envelope is defined by the boundaries of the best and worst case

Several research papers document the application of topology metrics to the Internet network or to large-scale research networks.

- CAIDA (The Cooperative Association for Internet Data Analysis) has performed a number of simulations on complex, highly interconnected, large-scale networks (a simulated Internet). The goal was to measure topological resilience in face of node and link failures. [54]
- A paper from UCLA has investigated the resilience of Internet nodes against BGP prefix hijacking. [55]
- Simulations of link failures and their impact to the network service parameters have been done for the GEANT2, Sprint and AT&T by the ResiliNets project. [56]

The complexity of these research methods are beyond the scope of this document.

Case studies

This section will demonstrate the significance of the design-based metrics presented above on a few example topologies. While the expected MTBF and expected availability are in most cases known for individual components and links, they must be calculated to represent the expected MTBF and expected availability for a network of devices and links. These calculations are called composition rules.

The composition rules depend on the system architecture and topology (e.g. series versus parallel connections) and can become complex very quickly. The formulas are documented in [41]. In this section, we will describe a few case studies to demonstrate the usage of the design-based metrics and the composition rules.

The topologies used in the case study will represent a small network with 2 servers and a number of network devices in between. We will use different topologies to illustrate the effect of series and parallel composition on the expected MTBF and expected availability.

Following assumptions are made (assumptions and topologies are taken from [57]):

- The expected network device MTBF is 45.000 hours;
- The expected MTTR is 4 hours;
- The calculated parameters only apply to the network hardware and make abstraction of software operation, cable ruptures and human errors;
- The availability is measured per year.

1.1.1.31 Topology 1 – No redundancy

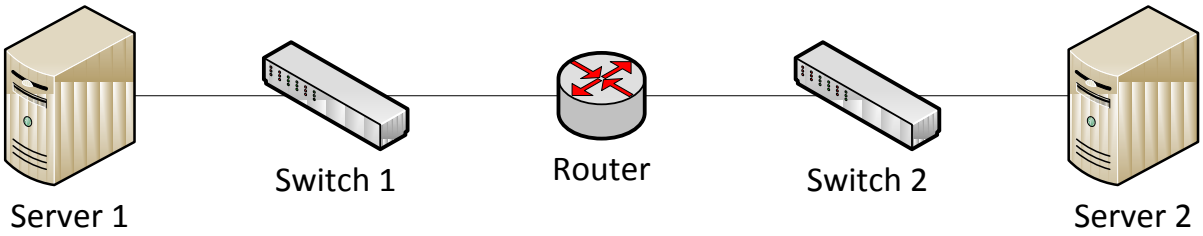

Figure 23: Topology 1

The first topology has no redundant network devices or connections. As a result, failures of any network device results in the failure of the entire system and loss of connectivity between the 2 servers.

Using the formula for series composition from [41], we calculate the expected availability of the system as the product of the expected availabilities of each of the 3 network nodes:

$$ExpectedAvailability_{Topology1} = ExpectedAvailability_{Switch1} * ExpectedAvailability_{Router1} * ExpectedAvailability_{Switch2}$$

Using the formula in 0, the expected availability is

$$ExpectedAvailability = \frac{ExpectedMTBF}{ExpectedMTBF + ExpectedMTTR}$$

The expected availability of each network device is 4 hours (MTTR) divided by the sum of 45.000 hours (MTBF of a network device) and 4 hours (MTTR). This amounts to 99,991% expected availability per year for each network component.

Combining the expected MTBF metrics to calculate the expected MTBF of topology 1 using the formula presented in this paragraph, the expected availability of the system is 99,973%.

Using an availability window of 1 year (cfr. assumptions), we can calculate the expected system MTBF:

- 1 year is equal to 8760 hours
- This means that topology 1 will have an expected annual uptime of 99,973% * 8760 hours or 8757,66 hours. This is the expected MTBF of the system.

1.1.1.32 Topology 2 – Redundant routers

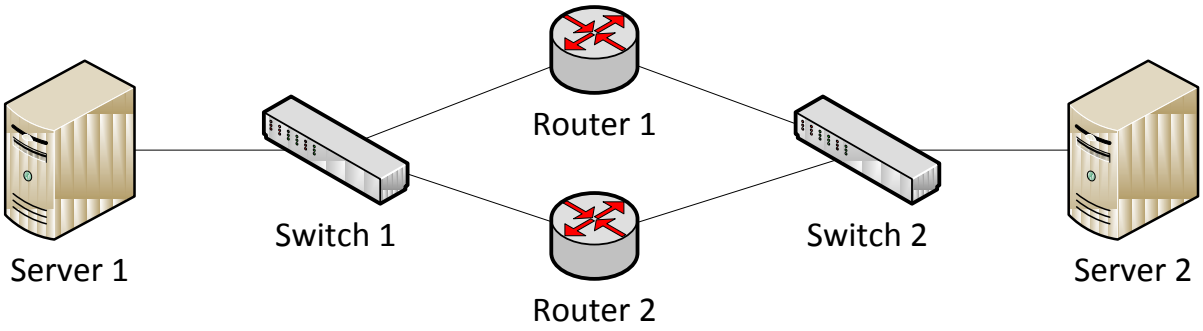

Figure 24: Topology 2

The second topology has redundant routers. The calculations remain the same as for topology 1 except that in this case, both routers must fail before connectivity between the 2 servers is lost. The probability of one of the routers operating can be calculated as 1 minus the probability of both routers being defective.

Using the formula for parallel composition from [41], we calculate the expected availability of the system as the product of the expected availabilities of each of the 4 network nodes:

$$ExpectedAvailability_{Topology2} = ExpectedAvailability_{Switch1} * (1 - ExpectedAvailability_{Router1} * ExpectedAvailability_{Router2}) * ExpectedAvailability_{Switch2}$$

Combining the expected MTBF metrics to calculate the expected MTBF of topology 2 using the formula presented in this paragraph, the expected availability of the system is 99,946%. Due to the redundancy in the setup of topology 2, its availability is higher compared to topology 1.

Using an availability window of 1 year (cfr. assumptions), we can calculate the expected system MTBF:

- 1 year is equal to 8760 hours;
- This means that topology 2 will have an expected annual uptime of 99,98% * 8760 hours or 8758,44 hours. This is the expected MTBF of the system.

1.1.1.33 Topology 3 – Redundant routers and switches

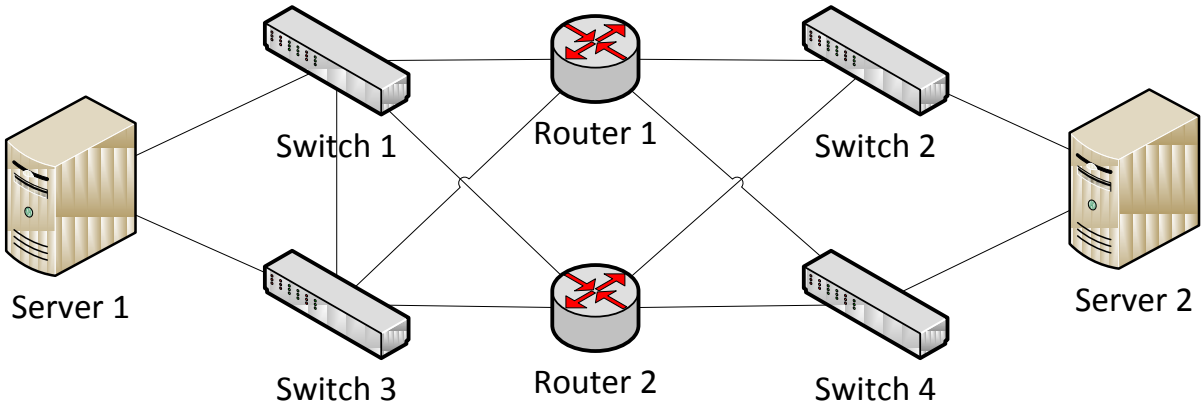

Figure 25: Topology 3

The third topology has redundant routers and redundant switches on each side of the routers. The calculations remain the same as for topology 2 except that in this case, the switches are calculated in the same way as in topology 2. Before connectivity between the 2 servers is lost:

- Switch 1 or switch 3 must be defective;
- AND either router 1 or router 2 must be defective;
- AND either switch 2 or switch 4 must be defective.

Using the formula for parallel composition from [41], we calculate the expected availability of the system as the product of the expected availabilities of each of the 4 network nodes:

$$ExpectedAvailability_{Topology1} = (1 - ExpectedAvailability_{Switch1} * ExpectedAvailability_{Switch3}) * (1 - ExpectedAvailability_{Router1} * ExpectedAvailability_{Router2}) * (1 - ExpectedAvailability_{Switch2} * ExpectedAvailability_{Switch4})$$

Combining the expected MTBF metrics to calculate the expected MTBF of topology 3 using the formula presented in this paragraph, the expected availability of the system is 99,999%. Due to the redundancy in the setup of topology 3, its availability is even higher compared to already partially redundant topology 2.

Using an availability window of 1 year (cfr. assumptions), we can calculate the expected system MTBF:

- 1 year is equal to 8760 hours;
- This means that topology 2 will have an expected annual uptime of 99,99% * 8760 hours or 8759,99 hours. This is the expected MTBF of the system.

Bibliography and references

[1] 'Measurement Frameworks and Metrics for Resilient Networks and Services: Challenges and Recommendations', ENISA;

[2] 'Gaps in standardization related to resilience of communication networks', ENISA;

[3] 'Stock taking report on the technologies enhancing resilience of public communication networks in the EU Member States', ENISA;

[4] 'Virtual Working Group on network providers' resilience measures and good practices- Activities Report – Paris' – ENISA, 29 Oct. 2009;

[5] 'The CIS security metrics - Consensus Metric Definitions v1.0.0', The Center of Internet Security, 2009;

[6] 'Performance Measurement Guide for Information Security – NIST Special Publication 800-55 Revision1', National Institute of Standards and Technology, 2008;

[7] 'Directions in Security Metrics Research – NISTIR 7564', National Institute of Standards and Technology, 2009;

[8] 'ISO/IEC 21827 - Information technology — Security techniques — Systems Security Engineering — Capability Maturity Model® (SSE-CMM®)', 2008;

[9] 'ISO/IEC 27004 - Information technology — Security techniques — Information security management — Measurement', 2009

[10] 'Critical Thinking: Moving from Infrastructure Protection to Infrastructure Resilience', George Mason University CIP Program, 2007;

[11] 'Measuring Cyber Security and Information Assurance' State of the Art Report, Information Assurance Technology Analysis Center (IATAC), 2009;

[12] 'Complete Guide to Security and Privacy Metrics: Measuring Regulatory Compliance, Operational Resilience, and ROI', Debra S. Hermann, 2007;

[13] 'Information Assurance Measures and Metrics - State of Practice and Proposed Taxonomy', Department of Computer Science Mississippi State University, 2002;

[14] 'D2.2 State of the Art', Assessing, Measuring, and Benchmarking Resilience (AMBER), 2009;

[15] 'CERT - FISMA and Metrics', Samuel A. Merell, 2007;

[16] 'Organisational Resilience: Security, Preparedness, and Continuity Management Systems – Requirements with Guidance for use – ASIS SPC.1-2009', ASIS International, 2009;

[17] 'Telecommunications Resilience Good Practice Guide – re-20040501-00393', National Infrastructure Security Co-ordination Centre (NISCC), 2006;

[18] 'Structure-Based Resilience Metrics for Service-Oriented Networks', Daniel J. Rosenkrantz, 2004;

[19] 'Security Metrics: Replacing Fear, Uncertainty, and Doubt', Andrew Jaquith, 2007;

[20] 'Information Security Management Metrics: A Definitive Guide to Effective Security Monitoring and Measurement', Krag Brotby, 2009;

[21] 'Emergency Communications (EMTEL); Overview of Emergency Communications Network Resilience and Preparedness – ETSI TR 102 445 v1.1.1', European Telecommunications Standards Institute (ETSI), 2006;

[22] 'D1.5a First interim strategy document for resilient networking)', ResumeNet, 2009;

[23] 'D2.1a First draft on defensive measures for resilient networks', ResumeNet, 2009;

[24] 'D6.3 Report of technical work in WP2 and WP3 during the 1st year', ResumeNet, 2009;

[25] 'D6.4b Periodic progress report', ResumeNet, 2009;

[26] 'The ARECI Study – Availability and robustness of electronic communications infrastructures – Final Report', Alcatel-Lucent, 2007;

[27] 'ETSI EG 202 057-4 V1.2.1', ETSI, 2008;

[28] ITU-T Y.1541, ITU-T, 2006;

[29] 'CERT Resilience Management Model, Version 1.0' - CMU/SEI-2010-TR-012, CMU/SEI, 2010;

[30] 'Service Level Specification Semantics and Parameters <draft-tequila-sls-02.txt>', IETF, 2002;

[31] 'NISTIR 6025 – Metrology for Information Technology (IT)', MEL/ITL Task Group on Metrology for Information Technology (IT), 1997;

[32] 'IEEE Standard Glossary of Software Engineering Terminology', IEEE, 1990;

[33] 'Measure, Metric, or Indicator: What's the Difference?', Bryce Ragland, 1995;

[34] 'Process Control System Security Metrics – State of Practice', I3P – Institute for Information Infrastructure Protection, 2005;

[35] http://www.isaca.org.uk/northern/Docs/ISACANC09_VPPres.pdf;

[36] ITU.T Recommendation X.805 (http://www.itu.int/ITU-T/worksem/ngn/200505/presentations/s5-zelstan.pdf

[37] 'A framework to quantify network resilience and survivability', Abdul Jabber, 2010 (http://kuscholarworks.ku.edu/dspace/bitstream/1808/6770/1/Jabbar_ku_0099D_11013_DATA_1.pdf);

[38] 'A Framework to Quantify Network Resilience and Survivability', James Sterbenz, 2010 (http://www.ittc.ku.edu/resilinets/presentations/sterbenz-hutchison-smith-cetinkaya-hameed-jabbar-rohrer2010.pdf);

[39] IEEE 90 – Institute of Electrical and Electronics Engineers, IEEE Standard Computer Dictionary: A Compilation of IEEE Standard Computer Glossaries. New York, NY: 1990

[40] 'Mean Time Between Failure: Explanation and Standards', Wendy Torell & Victor Avelar, APC, 2004 (http://www.ptsdcs.com/whitepapers/57.pdf);

[41] 'Probabilistic R&M Parameters And Redundancy Calculations', SARS, (http://www.sars.org.uk/BOK/Applied%20R&M%20Manual%20for%20Defence%20Systems%20%28GR-77%29/p4c06.pdf);

[42] ResiliNets initiative - https://wiki.ittc.ku.edu/resilinets/Main_Page;

[43] ResiliNets initiative - https://wiki.ittc.ku.edu/resilinets/Related_Work#Related_Disciplines;

[44] The NIST incident handling guide (NIST Special Publication 800-61 Revision 1): http://csrc.nist.gov/publications/nistpubs/800-61-rev1/SP800-61rev1.pdf;

[45] MIL-HDBK-217: http://snebulos.mit.edu/projects/reference/MIL-STD/MIL-HDBK-217F-Notice2.pdf;

[46] Telcordia SR-332: http://telecom-info.telcordia.com/site-cgi/ido/docs.cgi?ID=SEARCH&DOCUMENT=SR-332&;

[47] http://www.ncircle.com/index.php?s=solution_vulnerability-management-faq;

[48] http://kb.pert.geant.net/PERTKB/WebHome;

[49] RFC 3393: IP Packet Delay Variation Metric for IP Performance Metrics (IPPM);

[50] ETSI EG 202 057-4;

[51] ITU-T P.800: Methods for objective and subjective assessment of quality (http://www.itu.int/rec/T-REC-P.800-199608-I/en);

[52] 'Reliability-Centered Maintenance (Rcm) For Command, Control, Communications, Computer, Intelligence, Surveillance, And Reconnaissance (C4isr) Facilities', Headquarters, Department Of The Army, http://140.194.76.129/publications/armytm/tm5-698-2/entire.pdf;

[53] 'A Computational Approach to Multi-Level Analysis of Network Resilience', Christian Doerr and Javier Martin Hernandez;

[54] 'CAIDA: Topological Resilience in IP and AS Graphs '(http://www.caida.org/research/topology/resilience/);

[55] 'Understanding Resiliency of Internet Topology Against Prefix Hijack Attacks' (http://irl.cs.ucla.edu/papers/hijack-dsn.pdf);

[56] 'Modelling and Analysis of Network Resilience' (http://www.ittc.ku.edu/resilinets/papers/Sterbenz-Cetinkaya-Hameed-Jabbar-Rohrer-2011.pdf);

[57] High Availability Network Operations, Cisco Press, 2001;

PO Box 1309, 71001 Heraklion, Greece, Tel: +30 2810 391 280
www.enisa.europa.eu