




Resilience and Survivability for future networking: framework, mechanisms, and experimental evaluation



Deliverable number	1.1
Deliverable name	Understanding Challenges and their impact on network resilience
WP number	1
Delivery date	28/2/2009
Date of Preparation	3/4/2009
Editor	Marcus Schöller
Contributor(s)	Marcus Schöller (NEC), Paul Smith (ULanc), Andreas Fischer (UP)
Internal reviewer	David Hutchison

Summary

This deliverable presents the results of Task 1.2 on “Understanding Challenges”. We call all events, which have the potential to lead to a degradation of delivered service, challenges to the networked system. Understanding these events and how they affect the system is essential to build defensive measures and remediation strategies in WP2 and WP3.

We gather and analyze a variety of challenges from the literature on incidents which have led to system failures or a degradation of service in the past. These incidents encompass classical threats like attacks in the physical or virtual world, large-scale natural and human-caused catastrophes, but also misconfigurations, interoperability problems, and challenging communication environments. Our investigations lead to a taxonomy to classify challenges on their respective nature. This taxonomy is presented and provides the basis for a risk assessment process for resilience. This methodology serves as driver for some prioritization in the treatment of challenges; the assumption is that over the large space of defensive measures and detection algorithms the ones to be adopted and built should make the system resilient against the challenges with the highest impact. Last, a case study demonstrating the application of the proposed risk assessment process is presented. We use a small community mesh network in the north of England for our studies. An excerpt of our findings in this case study exemplifies our methodology and its outcomes.

Contents

1	Introduction	4
2	Related Work	6
2.1	Information Security Risk Assessment	6
2.1.1	CCTA Risk Analysis and Management Method (CRAMM)	6
2.1.2	Operationally Critical Threat, Asset and Vulnerability Evaluation (OC-TAVE)	7
2.2	Related Taxonomies	7
3	Taxonomy	8
3.1	Challenge Categories	8
3.2	Challenge Sheets	10
4	A Risk Management Approach for Resilience	11
4.1	Step 1: Identify critical assets	12
4.2	Step 2: Determine the cost of an asset's compromise	12
4.3	Step 3: Develop system understanding	13
4.4	Step 4: Identify potential challenges to the system	13
4.5	Step 5: Identify system faults	14
4.6	Step 6: Determine the likelihood of a challenge leading to a failure	14
4.7	Step 7: Determine a measure of exposure	14
4.8	Summary	15
5	Usage Scenario – Wray Village WMN	15
5.1	The Wray Village WMN	16
5.2	Assets in the Wray WMN	17
5.3	Identify Potential Challenges	18
5.4	Identify System Faults	18
5.5	Identify Probability of Challenges Leading to a Compromise	20
5.6	Determine exposure values	20
5.7	Wray Analysis Conclusion	22
6	Discussion	22
6.1	Implications for the rest of the work in ResumeNet	23
7	Conclusion	23
A	Challenge Catalog	26

1 Introduction

Omnipresent communication facilities have become an essential part of our daily private and business life, and a requirement for the operation of institutions and governments. We use networks to communicate via voice or video with family and friends, we use the Internet to buy products and services, and manage our finances. Business-to-business communication, customer care, and remote management of facilities are just a few examples of where we use networks in business environments. Governments provide services to their people and rely on communication infrastructure for emergency services and national security. As this dependence constantly increases, challenges to the normal operation of our communication infrastructure and the services deployed on it bear ever more severe consequences.

Therefore, understanding the challenges that affect the services of a system is essential to design and build resilience mechanisms. But as all mechanisms come at a cost, both implementations and operational, a structured procedure to identify the challenges which pose the greatest threats to the system is required. Building defensive measures and remediation mechanisms into the systems for these high-impact challenges provides a way to resolve most efficiently the trade-off between resilience and cost.

We propose to determine what the probable high-impact challenges are in a system by following an approach based on risk assessment. At the basis of the approach lies an understanding of the assets associated with a system. Challenges in and of themselves are meaningless if they are not seen as a threat to specific assets. Jones and Ashenden [JA05] provide a description of what an asset is:

“Assets can be physical, consist solely of information, or be functions that may enable a business process to be carried out.”

Example physical assets include infrastructure, such as router and server devices; information assets include email and customer information records; and functions that enable business processes include network services and protocols. Jones and Ashenden make the following notes about the value of assets:

“The value of an asset often is more than merely the capital costs or operational costs. Value can lie in the embarrassment that would be suffered by an organization if that asset was lost, the danger that would be posed to national security, or loss of business in a commercial organization.”

Assets can be defined at different levels of granularity by the various stakeholders of a system. For example, as we describe in Section 5, users of a network express assets in high-level terms, such as “Internet connectivity” and “privacy of browsing habits”, whereas a network provider will express these in terms of services and systems (e.g., DNS, routers, etc.), for example. To determine the probability of a challenge (e.g., an attack) having an effect on an asset, we need to decompose these expressions of an asset into the constituent software and hardware systems and services that provide the asset.

Fundamental to the availability of assets are the systems, composed of both software services and hardware, that support them. A system/service is considered to operate *normally* when there are no adverse events or conditions present. This loosely corresponds to the conditions for which the system was designed, i.e., when the networked system operates within the desired parameter space and is not under attack, the vast majority of network infrastructure is operational and connectivity is relatively strong. Two sources of events which pose a threat to this normal operation can be distinguished: *faults* of system components and *events* influencing

the system from outside. A fault in a system component triggers an internal fault, e.g., executing a programming bug or the random failure of a hardware component. *Challenges* encompass all outside events which affect the system and lead to an external fault. Examples include, but are not limited to, cyber-attacks, hardware destruction, and failure of the client service. Both classes of adverse conditions will be described in more detail in Section 3.1.

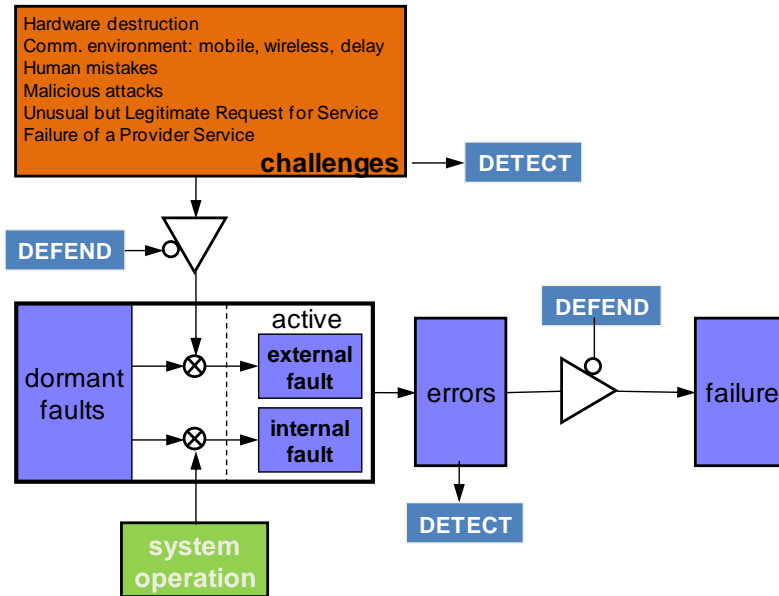


Figure 1: Block diagram of IFIP 10.4 definitions

Manifestations of internal or external faults are called *errors* [ALRL04]. “Errors may cause deviation of a delivered service from the specified service, which is visible to the outside world. The term *failure* is used to denote this type of an error.” [SS04]. These relationships, as captured within the IFIP (International Federation for Information Processing) Working Group 10.4, are depicted in Figure 1, which also shows where defensive measures and detection mechanisms can be applied. Detecting challenges and preventing them from affecting the system is essential to avoid erroneous behavior. But neither can we forecast all possible challenges for a system - a fact that forestalls the design of needed defensive measures - nor does it make sense to build a system, which can defend itself against every challenge due to the associated cost. Therefore, challenges will always affect systems and cause system errors. To contain them within the system and not let them influence other systems is the second line of defense.

Assessing which challenges affect a communication system in which way is essential to decide which defensive mechanisms to incorporate into the system and which remediation strategies to provide. In order to guide this decision we propose a structured risk-based approach (see Section 4) as a way to assess the impact of challenges on a system. Our approach identifies assets, challenges, the impact of the challenges on the assets, and the likelihood that they are expected to happen. The result of this procedure is a prioritized list of challenges for important system assets for which remediation mechanisms have to be put in place if defensive measures are not sufficient on their own. This approach involves a taxonomy that describes and classifies challenges, which is presented in Section 3.

2 Related Work

2.1 Information Security Risk Assessment

There are a number of approaches to understanding the most prominent threats in the area of information security. The aims of these methods are to identify the most probable and high-impact attacks, so that appropriate security mechanisms can be deployed, given a potential set of organizational and monetary constraints.

2.1.1 CCTA Risk Analysis and Management Method (CRAMM)

CRAMM [cra] was originally developed by the UK government in 1985, and is still widely used and licensed solely by Insight Consulting. It offers support for the ISO 270001 standard, and consists of software tools that support a three-stage process:

1. **Asset identification and evaluation**

CRAMM identifies three forms of assets: physical, software, and location assets. The aim of this first phase is to identify the monetary cost to the organisation if these different forms of assets were impaired (e.g., the cost to replace a router, if it was stolen).

2. **Threat and vulnerability assessment**

Having understood the cost associated with assets being impaired, the next step is to consider the probabilities of these costs being incurred. CRAMM aims to solicit both intentional and accidental threats to assets, such as attacks or human error. The result of this phase is a measure of risk.

3. **Countermeasure selection and recommendation**

Finally, CRAMM tools can be used to select appropriate countermeasures from a large database of around 3000 possible options. A key facet of this countermeasure selection is comparing the level of risk measured in phases one and two against the potential security gain of a particular countermeasure.

For our purposes, CRAMM is interesting as it starts with identifying the costs associated with asset impairment, and, in particular, the costs associated with different forms of impairment, e.g., unavailability, loss, or disclosure. When considering network-related assets, such as services, it may be challenging to classify impairments in such a clear way, and consequently associate a cost. Another interesting aspect of CRAMM is the scale of the countermeasure database. If one considers that their already large database is intended to solely address security concerns, one may directly infer that generalizing it for challenges associated with resilience would potentially raise some scalability concerns.

As CRAMM is one of the earliest structured risk assessment approaches ever developed, a number of other strategies have been derived from it, with various differences, including SARA (an accompanying fast-track assessment strategy called SPRINT), and COBRA. The common starting point of these approaches is the understanding of the business risk caused by threats, in other words, determining the assets associated with the information system in question and determining the cost of impairment.

2.1.2 Operationally Critical Threat, Asset and Vulnerability Evaluation (OCTAVE)

The OCTAVE method [ABPW99] is a threat-modelling process intended for use in large corporate, military, and governmental organizations, and developed by the Computer Emergency Response Team (CERT) in Carnegie Mellon University (CMU). It can be used to determine the critical assets of an enterprise and the technical vulnerabilities associated with them. A process is defined for this and a set of principles that OCTAVE-based methods should adhere to.

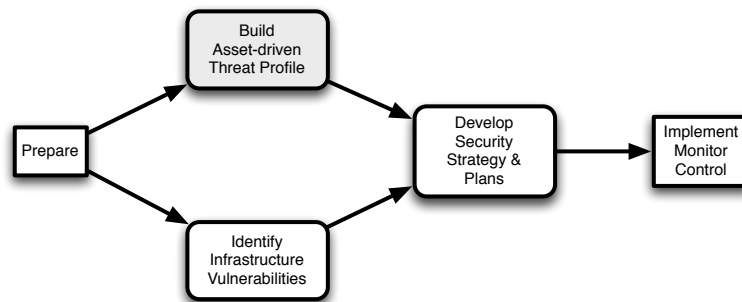


Figure 2: The OCTAVE Process

One such method that adheres to the OCTAVE principles is OCTAVE-S [ADSW03], specifically created for use in businesses with less than 100 employees. The process is carried out by small groups of people, knowledgeable about the network they are assessing; typically employees of the organisation carrying out the assessment. (This differs from CRAMM, where no explicit guidance is given on who drives the risk assessment.) Meetings are held to create asset-based threat profiles, which involve identifying the valuable assets existing in the network and ranking their value according to perceived necessity and importance. The second step is to identify vulnerabilities in the networking infrastructure that could affect the assets that have been identified as most critical. The last step is to take the information from steps one and two, identify specific risks that combine the vulnerabilities and assets, and from that develop security strategies to mitigate potential problems. This process is summarized in Figure 2.

2.2 Related Taxonomies

To our knowledge, no complete taxonomies on challenges and their impact on communication assets have been proposed so far. Nonetheless, taxonomies within our proposed categories already exist. Within the 'malicious attacks' category many taxonomies to classify attacks, the related vulnerabilities, and defense mechanisms have been presented. Igure et al. give an overview of "taxonomies of attacks and vulnerabilities in computer systems" in their survey [IW08]. Although limited to the attack category only, this paper inspired the development of our challenge taxonomy. Mirkovic and Reiher propose a taxonomy to classify DDoS attacks and DDoS defense mechanisms in [MR04]. Although addressing the same challenge domain as the paper from Igure et al., this paper relates attacks with defense mechanisms. Therefore it follows a similar objective as we are doing with our work. In comparison we do not provide a taxonomy of the defensive mechanisms or remediation strategies which can be applied to cope with the challenges but rather inform the system engineer about which assets have to be protected from the impact of challenges to gain the best resilience for the distributed system.

3 Taxonomy

We have developed a taxonomy to systematically document and assess the impact of challenges. This taxonomy provides the basis for our procedure to identify the challenges which pose the biggest threat to the system. We first introduce the challenge categories for the top-level classification. Afterwards, we outline what a challenge sheet defined by our taxonomy looks like and give an example.

The challenge taxonomy uses a two-level hierarchy to classify challenges. On the top-level we classify challenges in categories which reflect the nature of the challenge. The scenario to which this challenge applies is the second-level classification. Some challenges affect a networked system in a very general way, e.g., frequency jamming can be executed on any kind of wireless communication link. Other challenges are very specific to a certain scenario or set-up, e.g., HTTP is challenged by a failure of the TCP service on which it depends; but running HTTP over TCP is not a necessity from an architectural point of view, it is just the most common deployment today as Internet networking uses TCP as the reliable end-to-end transport protocol.

3.1 Challenge Categories

Learning from past failures is essential to build better systems in the future. To systematically group challenges we propose seven categories, as presented in this section. The first four challenge categories are a consequence of being part of the real world, i.e., deterioration, destruction, and physical channel characteristics. Challenge categories 5 and 6 are having their cause in a networked cyberspace, i.e., cyberspace attacks and non-foreseeable system interaction. The last challenge, which might be also seen as overlapping with all previous six categories, results from the composition of systems from subsystems where failures of a subsystem influences other subsystems and the system as a whole. These seven challenge categories are presented in detail, accompanied by some past examples.

Component Faults Internal errors occur during 'normal' operation of the system independent of outside events. They can be caused by software bugs or the deterioration of hardware, for example. In short, these are failures that are brought about by faults in components of the system.

Bugs in soft- or hardware in communication systems are unavoidable. For example, a bug of Cisco's IOS causes a BGP session reset if the length of the AS path exceeds 255 after AS-path prepending [Pep09]. This bug was triggered by a bug from another router vendor not handling configuration parameters correctly, especially missing bound checks [Zmi09]. The result was a tenfold increase in planetary routing instability for an hour and an increase of affected prefixes from 0.45% to 4.76%.

Hardware destruction This category summarizes all challenges where destruction of hardware causes errors or failures. These can be either due to natural causes (e.g., tsunamis, earthquakes or hurricanes) or man-made (e.g., terrorist attacks, fires or cable-cuts).

The Hinsdale Office Fire [Tow88] is an example where destruction of hardware caused significant outages of communication services. Despite having deployed redundant systems for fail-over, a system failure could not be prevented as both the primary and secondary system were physically co-located in the same office building which was destroyed by a fire.

Communication Environment related All challenges that are inherent in the communication environment due to:

- weak, asymmetric, and episodic connectivity of wireless channels
- high-mobility of nodes and subnetworks
- unpredictably long delay paths either due to length (e.g., satellite) or as a result of episodic connectivity

are gathered in the communication environment category.

An (extreme) example of this category is the ad hoc airborne network made up from supersonic jet aircrafts [JPS08]. The high mobility of the jets and drones challenges end-to-end communication since the network nodes are within communication range for a short time-period only. Special designed network and transport protocols are required to establish and maintain communication during aircraft operation. Similar challenges are presented by more conventional ad hoc networks; delay-tolerant, opportunistic, networks have largely emerged as a response to these challenges.

Human Mistakes Human mistakes describe non-malicious errors that are made by people that interact with the system, such as device misconfiguration or not following correct policy. These can become more pernicious if the parties involved try to cover up their mistakes.

Configuration mistakes in firewalls, automatic configuration systems (ACS) or end hosts often lead to degraded network service or prevent any communication. But misconfiguration can also have a larger impact. One example constitute the erroneous BGP advertisements of Pakistan Telecom's upstream providers, which resulted in a hijack of YouTube's Web presence [NCC08].

Malicious Attacks Malicious attacks from intelligent adversaries pose a threat to system performance and form a group of challenges to networked systems.

Cyberspace attacks can be specifically targeted at critical points of the communication system, e.g. Denial of Service attacks (DoS attacks), or be a resource exhaustion attack against a victim including collateral damage to the communication infrastructure, e.g., Distributed Denial of Service attacks (DDoS attacks).

Unusual but Legitimate Demand for Service A non-malicious request for service that is greater (or different along some other dimension) than what is provisioned for, for example, flash crowd events.

Neither the PSTN nor the Internet were designed to cope with the amount of traffic experienced after the 9/11 attacks. Both communication systems experienced severe local degradation of service [LeF01].

Failure of a Provider Service Due to the composition of complex system from multiple services any aforementioned challenge can cause cascade effects. The failure of a provider service must be treated as challenge to the consumer services which depend on the correct behavior of the provider service. As service usage can be vertical, i.e., using a lower layer service, as well as horizontal, i.e., client-server based or peer service, interoperability faults also fall into this category. Last, failing of the provider service due to an unidentifiable challenge is covered in this category.

Packet loss on a wireless link reduces the effective bandwidth of IP over this link and a degradation of the transport service in case of TCP.

3.2 Challenge Sheets

As mentioned before, challenges are classified along two axes. Firstly, along the challenge categories just introduced and secondly according to the scenario the challenge applies to. With respect to this second classification axis every challenge is described by the fields shown below:

- Name** : A descriptive name for the challenge
- Description** : What are the characteristics of the challenge?
- Fault** : Which fault is triggered by the challenge?
- Scope** : Which system/subsystem is affected by the challenge?
- Potential Impact** : What impact can the challenge have on the system?
- Parameters** : What parameters influence the behavior of the challenge?
- Symptoms** : Indicators which can be monitored or service metric affected; not all symptoms are shown in every scenario

We give an example which shows how to use the proposed taxonomy. More examples can be found in Appendix A, which furnishes at least one example for every challenge category. Based upon the taxonomy we define a challenge sheet which summarizes the taxonomy fields into four groups. The first group is called challenge and contains the name of the challenge. The second group is the classification group, followed by the characteristics group which contains information about what is going on, in which subsystem, and what is the impact. The last group, called details, provides the information to build challenge models and challenge detectors.

Challenge	Name	Frequency Jammer
Classification	Category	Malicious attack
	Scenario	Wireless communication
Characteristics	Description	The frequency used for communication is jammed by a) constant, b) periodic, c) interactive, d) arbitrary transmissions of the attacker.
	Scope	MAC on wireless medium
	Potential Impact	Communication among nodes in the vicinity is prevented or severely degraded
Details	Parameters	Duration of interference, period of jamming signal, output signal strength
	Symptoms	MAC layer protocol violation, disrupted link frames, reduced link bandwidth

Table 1: Frequency Jammer Challenge Sheet

It is important to note that a challenge may lead to observable symptoms at multiple layers. For example, a TCP connection running over a jammed wireless link will experience (physical layer) service failures, which result in negative acknowledgments and data retransmission. In the proposed taxonomy these symptoms are not symptoms of the *frequency jammer* challenge but symptoms of a failure of the IP service on which TCP depends. IP itself relies on an

operational MAC layer. From the TCP perspective IP failed to deliver the packet to the receiver. But this can have multiple causes like congestion, node outages, link breaks, or a jammed wireless channel. In addition, unexpected challenges can lead to a failure of a provider service but which can not be detected themselves. As the challenges is unexpected no monitoring capabilities have been build into the system. Therefore, the challenge class *failure of provider service* catches all failures of the provider service independent if the root cause of the challenge can be determined or not. This class of failures introduces recursive failing of services until the error can be contained within a consumer service or propagates to the application.

4 A Risk Management Approach for Resilience

We now describe the process that can be followed to determine the probable high-impact challenges to a networked system and associated assets. It is based upon the taxonomies that are described in Section 3, and the risk management strategies that are described in Section 2. A core motivation for carrying out a risk management process is to make optimal use of the potentially limited resources (e.g., time, computational, and monetary resources) that are used to mitigate risks. They allow informed decisions to be made about the trade-offs associated with managing challenges. This is done by understanding the challenges that will have the highest impact to the system and their frequency of occurrence. Limited resources can then be targeted at mitigating these risks.

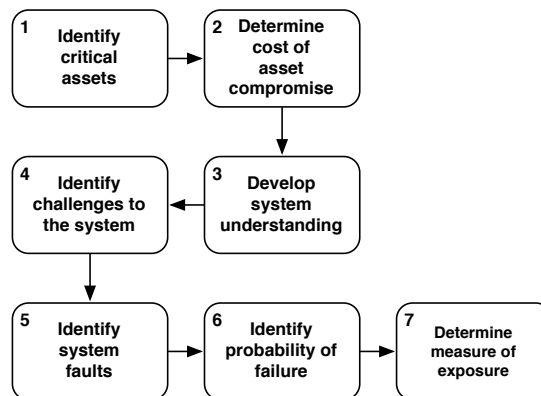


Figure 3: A sequential risk assessment process

The seven-step process we propose for determining the high impact challenges in a particular system context is shown in Figure 3. In summary, the process involves identifying the assets associated with a system and the impact (or cost) of their various forms of compromise. Subsequently, (or in parallel) the understanding of the system is developed, where the services that underpin the assets identified are elicited. Challenges specific to the system context are then identified, using various methods, such as advisories or further system modelling, and the probability of their occurrence is identified. Networked systems are not free of faults; the next step is to try and identify the shortcomings in the system that could lead to a challenge leading to a failure, and therefore the compromise, of an asset. Finally, we determine a *measure of exposure* for each envisaged challenge scenario. In simple words, exposure is the product of the impact of an asset being compromised in some way, and the probability of that event occurring; the latter is a function of the likelihood of a challenge occurrence and the ability of

the system to fail to mitigate it when it occurs. We describe this process in more detail in the following sections.

4.1 Step 1: Identify critical assets

In a similar fashion to OCTAVE [ADSW03] and other risk management strategies, we initially identify the assets associated with the system that are considered *critical*. What is considered a critical asset is very much context-specific, where the context is defined, for example, by stakeholders, the requirements of the organization using the system, and its application. Stakeholders include, for example, customers, providers and governments, and will have different viewpoints on what constitutes a critical asset. Stakeholder viewpoints should be prioritized. Example critical assets include emergency communication services, availability of Web-presence for an eCommerce company, and so on.

For a network provider, physical critical assets will include infrastructure, such as routers, servers, optical fiber, and so on. Information assets are likely to be wide-ranging: routing, billing and accounting, and network performance information are examples. Without these assets, the core business function of a network provider would be seriously impaired.

There are a number of ways of eliciting critical assets – for example, OCTAVE advocates using focus groups that are driven by an organization’s personnel; other approaches employ toolsets that can be used to develop system models that highlight potential assets [cra]. How assets are elicited is beyond the scope of this deliverable. In Section 5, we describe some that were identified using a focus group with members of a rural community that share a common wireless mesh infrastructure.

The result of this activity is a set of assets that are critical to an organization’s successful operation, which could be compromised in some way by a set of potential challenges to a system.

4.2 Step 2: Determine the cost of an asset’s compromise

An asset may be compromised in a number of different ways. For example, as the consequence of a challenge, it could be rendered completely unusable or intermittently available. In the context of information security, compromise is normally considered in terms such as disclosure and loss, for example. The level of degradation and the duration of a compromise will affect the impact a challenge has on assets – long service outages, due to a flash crowd or DDoS attack, will likely have a higher impact than short disruptions, due to a switch-over to a backup configuration.

As with understanding the intrinsic value of an asset, measuring the impact of a compromise can be done quantitatively or qualitatively. For example, a DoS attack could cause a degradation of service on an ISP’s network, leading to SLAs being broken – this could result in reimbursements to customer networks. The same incident could lead to a loss of reputation for the ISP, a qualitative impact. At this step in the methodology we simply consider the different modes of compromise associated with an asset, and their impact on the stakeholder.

4.3 Step 3: Develop system understanding

With an understanding of the critical assets, the next phase we propose is to identify the technical implementation of these assets by means of services. State-of-the-art software engineering approaches decompose the provisioning of assets into multiple sub-systems and services. Such sub-systems and services are often used as components of multiple assets. For example, the asset *Internet connectivity* is composed from services like, TCP/UDP, DNS, IP, ARP, 802.11b. This inherently implies that an asset can be (partly) lost if one or more services composing the asset are affected by a challenge. The result of this phase is an understanding of the services that can be affected by the challenges to be identified in the next step.

4.4 Step 4: Identify potential challenges to the system

Given an understanding of the system, its associated assets and modes of compromise, the next phase is to identify challenges to the system that might occur. We identify two main ways of deriving the challenges that might threaten a system: *system analysis* and *learning from past incidents*. A search functionality for challenges of sub-systems and services identified in phase three can be provided if challenges get documented systematically, i.e., using the taxonomy proposed in Section 3.

System analysis is a mature research area. For example, fault-tree analysis [Ves87] could be used to determine the events that could lead to an undesirable event occurring, such as a challenge. In short, this involves identifying a target undesirable event (the root of the fault tree) and the possible occurrences in the system, combined using Boolean logic, that could lead to the event occurring. If probabilities are associated with nodes in the tree, the probability of the root event occurring can be calculated¹. A similar process can be followed for malicious attacks [Sch99]. Other, more formal, threat-modelling techniques have been proposed such as STRIDE [HLOS06], where architectural models of a system are used as a basis to identify where certain classes of threat (spoofing, tampering, repudiation, information disclosure, denial of service, and elevation of privilege) could be enacted on a system.

Looking at past events can be used a basis for improving distributed systems and strengthen their defensive measures. In the domain of network security many sources to inform the system engineer about threats and pitfalls are available. For example, in the context of malicious attacks, information provided by security advisories [cer, san, mic] can be used to identify prevalent threats. Similar sources of information should be established to document resilience related incidents where the defensive measures and remediation efforts were insufficient to cope with the challenges. In fact, we wish to exploit such learning from past events in the ResumeNet project as a method to evolve the system into a better state. Making this system-internal information available to others could be used as an information source.

The result of this phase of the risk assessment is a set of challenges that are expected to occur in the deployment scenario. As in the security domain, many challenges are not known at system design time. Nonetheless, building defensive measures for the known challenges provides good guidance for the engineer. Providing mechanisms to cope with unexpected events is a complementary effort to this.

¹Related to fault tree analysis, is event tree analysis [Coo07], which takes an inductive, rather than deductive, approach to determining potential faults.

4.5 Step 5: Identify system faults

Creating complex systems that are fault-free and are complete immune to attacks from adversaries is not possible. In this phase, the aim is to identify faults associated with a system that could lead to the critical assets identified in phase one being compromised. In the context of information security, vulnerabilities include poorly configured (or non-existent) firewalls, poorly designed software that could lead to buffer overflow attacks, and poor security processes. However, in the context of resilience, the forms of faults to consider are more wide-ranging. Faults with respect to resilience range from design faults, inappropriate business-processes and interface design that may lead to operational errors, to vulnerability to large-scale destruction of hardware, for example. They may also include inappropriate use of resilience strategies, such as redundancy and diversity.

Looking back on past events can be used to inform us about faults that lead to a service degradation or failure. For example, Hurricane Katrina or the terror attacks of 9/11 showed the impact large-scale disasters can have on a distributed system. Although the Internet as whole has proven relatively resilient to such incidents, local communication was not restored for days and weeks. Side effects of both events have been flash crowds on news sites, which could not cope with the significant increase of requests from people wanting to inform themselves [LeF01]. Erroneous deployment of redundant systems is also a fault of a resilient system. The fire in the Hindsdale office tower disrupted local, long-distance, and mobile telephone service, as well as ATC for Chicago O'Hare International Airport (ORD) [Tow88]. Although a redundant router for fail-over was installed the system failed as both the primary and secondary routers were destroyed by the fire. Similarly, the Baltimore tunnel fire [Var05] showed that virtual redundant links running over the same physical link are a deployment fault and lead to a service failure if challenged by such destructive events.

The system engineer does not usually build a communication system from scratch; existing hard- and software available on the market is normally used. As all these products have faults and additional ones are introduced during composition and deployment, the system engineer has to gain an understanding of these faults. Again, looking back at past events provides guidance here. Systematically documenting challenges, i.e., using a taxonomy like the one we propose, eases the access to required information.

4.6 Step 6: Determine the likelihood of a challenge leading to a failure

Having identified the potential challenges and gained an understanding of the shortcomings of the system, the next phase is to determine the likelihood of a challenge leading to a failure. The probability of a failure is influenced by the nature of a challenge and the system properties, in terms of known faults and the mechanisms that are in place to mitigate challenges. As was shown in Section 1, two lines of defense can be drawn to prevent failures. Defending against the challenge itself, i.e., attack prevention, redundant forwarding paths, data coding schemes, etc., or by containing errors within the service, i.e., error recovery mechanisms (ARQ), store-and-forward mechanisms for DTN, etc. Again looking at past incidents guides the assessment of this probability.

4.7 Step 7: Determine a measure of exposure

With an understanding of the challenges, the likelihood of their occurrence and probability of leading to a system failure, and assessing the impact of assets being compromised, the final

step of the assessment is to determine a measure of exposure (or risk) for each challenge scenario.

In step two of this process the impact of a particular mode of compromise for an asset will have been identified. Information risk management processes vary in the way they represent impact. For example, some give qualitative values (e.g., mild, moderate, severe, or catastrophic), while others aim to associate quantitative values (e.g., potential financial loss, cost of replacement of equipment, etc.). There are pros and cons associated with each approach, which are discussed in [JA05]. We do not want to prescribe one particular strategy, but we would like a numeric value of exposure associated with each challenge scenario. Therefore, we suggest that measures of impact should be the result of a function that yields a value in the range $[0,1]$, where its input is the most suitable qualitative or quantitative measure of impact, depending on available data and context. We determine a measure of exposure using Equation 1, where *challenge_prob* is the probability that a challenge will occur (determined in step four), *compromise_prob* is the probability that a compromise will occur to an asset, which is based upon the likelihood of a system failure, and *impact* is the cost associated with an asset being compromised.

$$exposure = (challenge_prob \times compromise_prob) \times impact \quad (1)$$

It is clear that using this strategy to determine a measure of risk will yield similar values for high-probability, low-impact events and low-probability, high-impact events. Depending on which scenarios the system engineer would like to account for, exposure measures could be filtered out.

4.8 Summary

Step five of the presented risk management approach serves an additional purpose to build resilient systems besides providing the basis to determine the likelihood of a challenge leading to a failure. Having gained an understanding of the faults of the system, it enables the system engineer to design effective defensive measures or remediation strategies. From this design they can derive an estimate on the cost of the required measures that have to be weighed against the costs of losing this asset and its exposure.

The described process is a guideline for system engineers to acquire information about a target networked system, which is needed to design and integrate the resilience mechanisms that can protect the system from the most severe challenges. We used this methodology to identify the challenges in a wireless mesh network. First results of this case study are presented in the next section.

5 Usage Scenario – Wray Village WMN

In this section, we demonstrate the applicability of the risk management approach described in Section 4. We do this in the context of a wireless mesh network (WMN) that is operational in a rural village, called Wray [IBPR08]. Initially, we describe at a high-level the Wray network and subsequently the assets and the manner in which they can be compromised. This corresponds to steps one to three in our risk management approach. We then describe some potential challenges in the Wray context, and assign probabilities of their occurrence (step four in our approach). The probabilities we assign are based upon our intuition and past experience. We

then go on to describe some faults in the Wray network that relate specifically to the challenges we have identified (step five). Based upon knowledge of the system faults and the challenges, we then assign probabilities to a challenge compromising an asset in the network – this is step six in our approach. Finally, we determine a measure of exposure for each challenge scenario.

5.1 The Wray Village WMN

WMNs create a resilient infrastructure using a combination of wireless networking technology and ad-hoc routing protocols that together provide the ability to establish networks in locations with no prior groundwork, and where a wired network would be prohibitively expensive or complex. For example, the community of Wray, situated approximately ten miles from the city of Lancaster in the north-west of England, felt strongly that the lack of broadband Internet connectivity (due to their remoteness) in their village was jeopardizing local businesses, education, and the community itself.

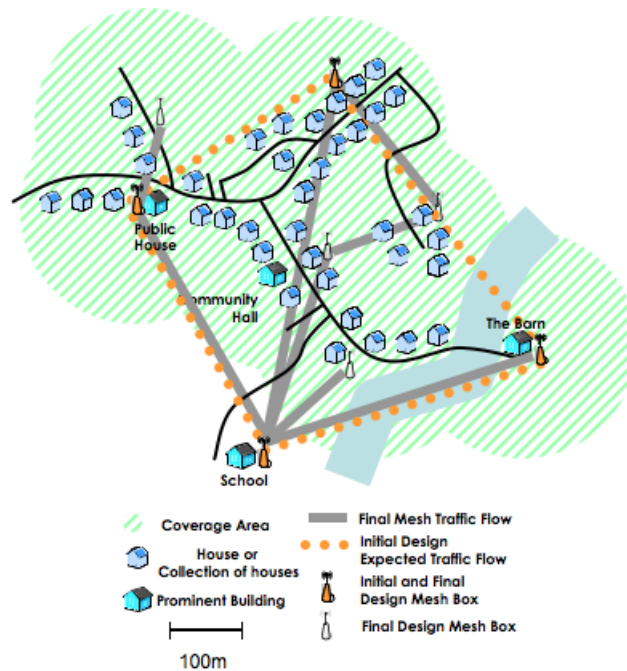


Figure 4: The Wray village wireless mesh network

The WMN that operates in the village is shown in Figure 4. In summary, clients connect wirelessly using IEEE 802.11b to the network via a mesh device, which forwards traffic to a single point in the village (the school) where back-haul is provided to the Internet. At the time of writing, routes are hard-coded on the mesh devices. Each mesh device runs a number of services: DNS, DHCP, NAT, and a firewall. If DNS fails on a mesh device, a server at the University is used. Recall the approach we adopt to identifying risks initially involves eliciting assets associated with a system, we discuss the assets identified in Wray in the following

section. For further details regarding the nature of the Wray deployment, please see [IBPR08].

5.2 Assets in the Wray WMN

There are a number of stakeholders associated with the Wray village network, including its users (who also act as infrastructure providers within the village), researchers at Lancaster University, and the back-haul provider. To understand the nature of the assets from the user group's perspective, Bury et al. [BIR⁺08] held an OCTAVE-style focus group within the village. They found the types of assets described by the villager's to be wide-ranging. Specifically, they identified the following assets which can be grouped into three categories:

1. Security: stored photographs and work documents; unfettered use of the home computer; personal information stored on the computer; and the protection of children from various forms of exploitation or abuse.
2. Safety: of the less computer literate members of the community; personal identity and reputation – protection from theft and damage or unauthorized use; of browsing habits etc.
3. Connectivity: Access to the Internet with a high quality connection that allows remote access, teleconferencing and Skype.

It can be seen that the villagers have a number of data-related assets. Also, they value Internet connectivity, a combination of physical assets, software services, and protocols, amongst other things, and use of the home computer as an asset. As the villagers are not corporate organizations, as are normally considered in information risk management practice, what they consider to be assets (e.g., personal identity and reputation, and dependants) can differ from what Jones and Ashenden [JA05] consider them to be. However, for the most part, these are related to what are conventionally considered assets.

The researchers of Lancaster University have quite different assets. They are investigating the community mesh network from the inside, e.g., active and passive measurements, as well as from the outside, e.g., study of focus groups. Their assets are very specific for this setup and will not be found in comparable scenarios. Therefore, we do not include their assets in this example analysis. The back-haul provider's main assets are more clearly defined and focused on the technical aspects of the network:

1. Security: Protect the infrastructure and deployed services from attacks and misuse
2. Connectivity: provide and maintain connectivity as defined in the SLA, i.e., guaranteed bandwidth, packet loss rate, etc.
3. Services: provide and maintain services to the community mesh network

The provision of connectivity is a central asset for the two main stakeholders of our scenario. Many other assets depend on the correct operation of it. Therefore, we use this asset as an example to illustrate our process.

With an understanding of the critical assets, we need to determine the impact of the different ways in which they could be compromised (step two in our risk management approach). Jabbar et al. [JKHS08] propose a two-dimensional state space to describe the state of a network and associated services: along one dimension the operational state of a network is described as

normal, partially degraded, or severely degraded, and the other dimension the state of a service is described as *acceptable, impaired, or unacceptable*. Unchallenged, the operational state of a network should be *normal* and the state of an associated service *acceptable*. Challenges can lead to a service, such as Internet connectivity, being either *impaired or unacceptable*, depending on the mechanisms in place to resist a challenge and its severity.

If we consider a member of the Wray community who uses the WMN for business purposes, of which there a number of examples, an impaired service over short periods of time (e.g., high delay, some packet loss, etc.) is likely to have some impact on their business through loss of sales, time and competitive advantage. We therefore suggest a medium-to-high impact to this scenario. If Internet connectivity is deemed to be unacceptable for extended periods (e.g., no connectivity at all for a series of days), then the business impact will be much greater; we therefore propose a high impact to this scenario. For each of these levels of impact, we assign a value in the range [0,1], which would be indicative of the severity of the scenario in relation to other assets and their mode of compromise. This is summarized in Table 2.

Asset	State	Impact	Value
Internet connectivity	impaired	medium-to-high	0.75
	unacceptable	high	0.9

Table 2: Measures of impact for the compromised states of Internet connectivity

5.3 Identify Potential Challenges

With an understanding of the system and the assets associated with it, the next phase is to identify the challenges that are likely to occur and their probability of occurrence. Appendix A presents a catalog² of potential challenges that could affect the quality of Internet connectivity in the Wray network. To further our process, we associate approximate probabilities of these challenges occurring and justify them (as shown in Table 3); this is based upon our intuition and past experience in Wray. The probabilities shown are relative to each other, for example, a faulty interface clock is much less likely to occur than signal attenuation. How to determine accurate probabilities of challenges occurring is a matter for further research.

An open question is to determine the timeframes to consider the probabilities of challenges occurring. Should we consider the probability of a challenge occurring within an hour, day, month or year, for example? A way to consider addressing this question is to understand the availability requirements of assets. For example, if an SLA specifies a desired level of availability of an asset (e.g., connectivity) over a period of time (e.g., 99.9995% availability over a monthly billing period), then it might make sense to consider the probability of a challenge occurring within this time period. In our example, we consider the probability of challenges occurring over a single year.

5.4 Identify System Faults

The set of faults associated with a system is potentially very large. To manage this significant space, we identify the faults that could be triggered by the most likely challenges that have been identified. We propose that this should not require complete understanding of a system. For example, we could use historical data to determine with a certain probability that a fault may

²This is not intended to be comprehensive.

Challenge	Prob.	Reason
Faulty interface clock	0.00005	High quality of device manufacturing process leads to few faults of this sort.
Broken antenna	0.01	Antenna are exposed to high winds and inclement weather conditions, some exist at street level (i.e., not on roofs), so could be maliciously vandalized. Operational experience suggests aerials have to be changed on occasion.
Signal attenuation	0.1	Often inclement conditions, (old) thick-walled buildings, past experience of vehicles obstructing antenna make this a high probability.
Routing loop	0.01	Mesh devices are configured by non-experts, though with increasing operational experience, with known issues occurring when misconfigurations occur. However, small hop counts make routing loops unlikely.
De-association attack	0.005	Requires access to locality, which is unlikely as Wray is rural. Also, requires high levels of expertise and there is little incentive for this form of attack in the context.
ARP storm	0.005	Some file sharing activity on the network and relatively novice user-base could lead to malicious code infection. However, the network's address space is private, which makes external attacks difficult to achieve.
TCP connection establishment failure	0.1	Sustained signal attenuation problems could lead to this challenge occurring.

Table 3: Challenge probabilities and justifications

exist in an operating system implementation, or a network operator would mis-configure one of their routers. In relation to the signal attenuation challenge identified above, we envisage the following example system faults (or shortcomings):

- **A lack of redundant paths**
 Redundant network paths potentially using different types of connection (e.g., wired links) could reduce the impact of signal attenuation on communication. In Wray, these are not available.
- **Inability of wireless devices to dynamically change channels**
 In the presence of attenuation caused by other wireless devices, such as Bluetooth handsets or other wireless access points, a wireless interface could shift onto an orthogonal channel. This capability is not available in Wray.
- **Routing tables are hard-coded**
 Hard-coding routing tables ensures that alternative routes that could circumvent the cause of signal attenuation is not possible.

With this knowledge, we can understand how well the network in Wray might perform in the presence of signal attenuation. In other words, we can determine the probability that a challenge will lead to an asset becoming compromised, which is the next stage in our approach.

5.5 Identify Probability of Challenges Leading to a Compromise

With an understanding of the challenges that are likely to occur and the system faults, we can determine the probability of a particular form of compromise to an asset occurring. Note, this differs from the probability of a challenge occurring – a challenge may occur and have no impact on a system and associated assets because appropriate defense mechanisms are in place. For the challenges we have identified in Table 3, if they occur, the probability they will cause the Internet connectivity asset to become impaired, unacceptable, or remain unaffected is presented in Table 4. As with the previous probability figures, these are estimates and determining more realistic values are a matter for further work.

Challenge	Impaired	Unacceptable	Normal
Faulty interface clock	0.7	0.2	0.1
Broken antenna	0.1	0.8	0.1
Signal attenuation	0.8	0.1	0.1
Routing loop	0.1	0.8	0.1
De-association attack	0.2	0.8	0.0
ARP storm	0.4	0.2	0.4
TCP connection establishment failure	0.2	0.8	0.0

Table 4: Probabilities of effects of challenges on Internet connectivity

It can be seen in Table 4 that if a faulty interface clock exists on a system, 70 % of the time it will lead to an impaired service, 20 % to an unacceptable service, and 10 % of the time will have no effect. The figures should be derived from an understanding of the system and its associated faults, and the nature of the challenge. For example, in the Wray WMN, mesh devices are not equipped with redundant wireless network cards, so if a clock on the interface becomes faulty, this will almost certainly lead to a reduced service.

5.6 Determine exposure values

The final phase of our proposed approach is to determine a measure of exposure for each challenge and compromise scenario. To do this, we initially determine the probability of a challenge leading to an asset being compromised in a particular way by calculating the product of the probability of a challenge occurring (from Table 3) and the probability of that challenge leading to a compromised state (from Table 4). We then take the product of this value and the impact of a particular compromise occurring to an asset (from Table 2) to gain a measure of exposure. This is shown in Equation 1.

A way of depicting the results of this process is shown in Figure 5; we call this an exposure graph. The connections on the left-hand side of the graph (from the *asset* column to the *compromise* column) are annotated with the impact a particular compromise to an asset could have on the studied stakeholder. The annotations on the connections of the right-hand side of the graph show the product of the challenge probability and the challenge leading to a particular compromise.

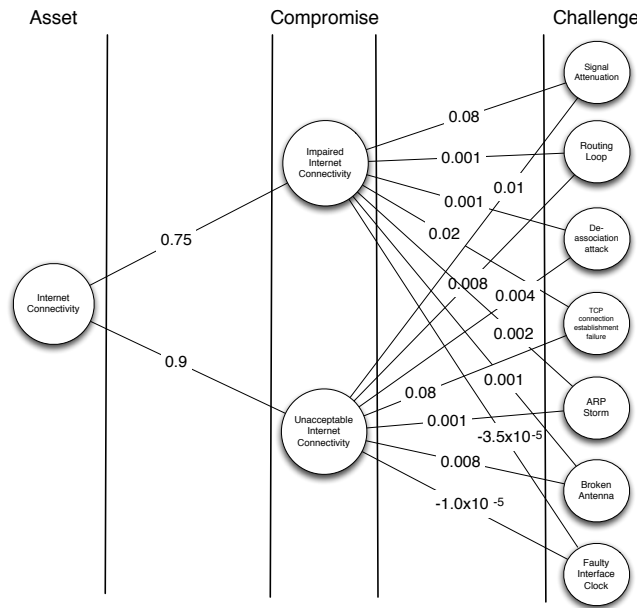


Figure 5: An example exposure graph for the Wray WMN use-case study

Table 5 shows the sorted exposure values for the different challenge scenarios in relation to the Internet connectivity asset. It can be seen that TCP connection establishment failures and signal attenuation are the two most significant challenges that could lead to impaired or unacceptable service. Following this, broken antenna and routing loops caused by misconfigurations are the next most significant challenge. An engineer wishing to improve the resilience of the Wray network could use this information to select appropriate defense mechanisms, e.g., introduce multiple paths to avoid service degradation due to signal attenuation, or improve known faults, for example.

Challenge	Compromise	Exposure
TCP connection establishment failure	Unacceptable	-7.2×10^{-2}
Signal Attenuation	Impaired	-6.0×10^{-2}
TCP connection establishment failure	Impaired	-1.5×10^{-2}
Signal Attenuation	Unacceptable	-9.0×10^{-3}
Broken Antenna	Unacceptable	-7.2×10^{-3}
Routing loop	Unacceptable	-7.2×10^{-3}
De-association attack	Unacceptable	-3.6×10^{-3}
ARP storm	Impaired	-1.5×10^{-3}
ARP storm	Unacceptable	-9.0×10^{-4}
Broken Antenna	Impaired	-7.5×10^{-4}
Routing loop	Impaired	-7.5×10^{-4}
De-association attack	Impaired	-7.5×10^{-4}
Faulty Interface Clock	Impaired	-2.6×10^{-5}
Faulty Interface Clock	Unacceptable	-9.0×10^{-6}

Table 5: The challenge scenarios sorted on exposure for the Internet connectivity asset

5.7 Wray Analysis Conclusion

In this section, we have presented a summarized version of the risk management approach described in Section 4. In the Wray village WMN there are a number of stakeholders, with differing assets that are wide-ranging. Some of these assets are not those typically considered in an enterprise setting. However, they are supported by a set of underlying services that can be compromised in various ways. By understanding the nature of the system and its faults, we can determine the probability of a challenge leading to a particular asset being compromised in a certain way. This, combined with the impact of an asset being impaired, yields a measure of exposure that can be used by engineers to influence where defense and mitigation mechanisms should be deployed. The probabilities we present in the section are not provably sound – deriving meaningful probabilities is an issue for further work. We have presented a risk management process, additional work is required to build tools to support this potentially laborious task.

6 Discussion

While our risk management approach sets forth a concrete base for assessing quantitatively the impact of challenges, there are certain aspects of it that necessitate further exploration:

Determining reliable measures for challenge occurrence probabilities Further work has to be carried out on estimating reliably the occurrence probabilities of challenges. On the one hand, this includes the analysis (and assumes availability) of monitoring information. First results from our study scenario in Section 5, reveal rather unexpected incidents: everyday at around 9am a link goes down for a couple of minutes. On-site investigation identified a milk truck causing the challenge, as it was parking at the same spot on its tour obstructing the line of sight of a wireless link. On the other hand, information from other sources has to be gathered and analyzed: as humid weather conditions and heavy winds are common phenomena in the north of England, signal attenuation or broken antennas are likely events; whereas crime statistics show that vandalism will hardly cause broken antennas.

Approaches to determining system faults As we discuss in Section 4, the range of potential system faults in relation to resilience is potentially vast. This is especially true when we consider complex interconnected systems (such as the Internet) that are made up of opaque components – consider a typical operating system or a peer autonomous system. Looking at historical data will help considerably here.

Quantifying the impact of a challenge on a system and its assets An important part of the approach we adopt is to quantify the impact of challenges on the systems and services that support an asset, i.e., determine the likelihood of a failure occurring. To do this, we could use historical data to gain an inference about the probability of a failure occurring in light of a challenge. However, the context of past events may be different to that of the system under examination, which may lead to an inaccurate understanding. One way to improve confidence in the probability of a failure occurrence is to develop simulations of the system, including (statistical) fault models, and exercise challenges against it. Before developing simulations of challenges, we first need to enhance our challenge taxonomy further, and, in particular, the challenge sheets. An ideal formulation would involve directly translating a high-level description of a challenge, akin to our challenge sheets,

into a simulation model. Analytical models would also be of use in this task, where analysis can track or reasonably abstract the complexity of the problem.

6.1 Implications for the rest of the work in ResumeNet

This risk assessment procedure complements the technical aspects of the proposed $D^2R^2 + DR$ strategy, which will be continuously validated in Task 1.1 of the project. The presented results add a new dimension to this work.

As already mentioned in Section 4, which describes our risk management approach, the analysis of past events is a main information source to understand potential challenges of a system. The background operation control loop of the proposed $D^2R^2 + DR$ strategy, namely the *diagnose* and the *refine* phases, are based on the storage of historical data. In order to refine the system over time, one has to analyze the collected data from logged incidents, the selected countermeasures, and the validation results. Making such information available to system engineers could prove useful and lead to a better understanding of potential challenges and their probability of occurrence. In Tasks 2.2 and 2.3, data storage mechanisms to gather and structure such information will be investigated. However, making the information available is not only a technical problem; it also entails privacy and operational issues. As we monitor communication on multiple layers and at various locations within the network, this data might contain private information about the system's users or details about the resources, services, and policies of the infrastructure provider. A well balanced and law-abiding solution (where privacy law is handled very differently in different countries) has to be found.

Understanding the critical assets of the stakeholders and their technical implementation provides a solid basis for operational policies for resilience. Task 1.4 will define policies, which guide the decision process of the *remediation* phase. Such policies are important when multiple services are affected by a challenge and a decision must be made for which of these service remediation procedures should be executed, and which services will be dropped.

This work could also be used in Task 2.2.1 on *understanding normal behavior*. One of our resilience principles is that it is necessary to understand the normal behavior of a system to determine when it is challenged, the effects of a remedy, and when a challenge has abated. The set of aspects of a system (e.g., protocol and system behavior) that could be used to understand normal behavior is very large, and is probably not tractable. How do we manage this space? By understanding the probable challenges that will occur to a system (essentially the aim of this work), and then understand how they perturb a system, we can then select an appropriate set of facets (signals) to understand its normal behavior. By doing this, we know the signals we have selected are perturbed when the system is exposed to the most probable and high-impact challenges.

7 Conclusion

In this deliverable we presented a methodology to identify the challenges which have the most severe impact on a networked communication system. This task is essential in order to enhance the system with defensive measures and remediation mechanisms to increase its resilience, while keeping costs reasonable. On the one hand, our methodology is based on a challenge taxonomy to classify and describe these challenges. On the other hand, the communication system is defined by a set of assets, i.e., the values of the system we have to protect. For

every combination of *asset, challenge* pair, we calculate an exposure value which represents the likelihood of the challenge affecting the asset and the related loss of value. Thereby, an ordered list of *asset, challenge* pairs is produced. Enhancing the system with mechanisms to prevent the challenges from affecting the system follows this ordered list. Mechanisms can be added as long as the overall costs of all mechanisms stay below a predefined cost limit. We used the community mesh network deployed in the village of Wray, Lancashire, UK, to illustrate our methodology and to suggest the feasibility of this approach.

References

- [ABPW99] C. Alberts, S. Behrens, R. Pethia, and W. Wilson. Operationally Critical Threat, Asset, and Vulnerability Evaluation Framework, version 1.0. Technical Report CMU/SEI-99-TR-017, Carnegie Mellon University, June 1999.
- [ADSW03] C. Alberts, A. Dorofee, J. Steven, and C. Woody. Introduction to the OCTAVE Approach. Technical report, Carnegie Mellon University, 2003.
- [ALRL04] A. Avizienis, J. C. Laprie, B. Randell, and C. Landwehr. Basic concepts and taxonomy of dependable and secure computing. *Dependable and Secure Computing, IEEE Transactions on*, 1(1):11–33, 2004.
- [BIR⁺08] Sara Bury, Johnathan Ishmael, Nicholas J. P. Race, Paul Smith, and Mark Rouncefield. Towards an understanding of security concerns within communities. *Wireless and Mobile Computing, Networking and Communication, IEEE International Conference on*, 0:478–483, 2008.
- [cer] CERT. <http://www.cert.org>.
- [Coo07] M. J. Cooper. *Event Tree Analysis*. Brunel Technical Press, 2007.
- [cra] CRAMM. <http://www.cramm.com>.
- [HLOS06] Shawn Hernan, Scott Lambert, Tomasz Ostwald, and Adam Shostack. Uncover Security Design Flaws Using The STRIDE Approach. *MSDN Magazine*, November 2006.
- [IBPR08] J. Ishmael, S. Bury, D. P. Pezaros, and N J.P. Race. Deploying Rural Community Wireless Mesh Networks. *IEEE Internet Computing Magazine*, 12(4):22–29, July–August 2008.
- [IW08] V. Ijure and R. Williams. Taxonomies of attacks and vulnerabilities in computer systems. *Communications Surveys & Tutorials, IEEE*, 10(1):6–19, 2008.
- [JA05] Andy Jones and Debi Ashenden. *Risk Management for Computer Security – Protecting Your Network and Information Assets*, chapter 11. Elsevier, 2005.
- [JKHS08] Abdul Jabbar, Manasa K., David Hutchison, and James P.G. Sterbenz. A framework to quantify network resilience. ITTC IAB poster, Available on-line at: <http://www.ittc.ku.edu/resilinet/posters/ITTC-IAB-poster-2008-metrics.pdf>, June 2008.

- [JPS08] Abdul Jabbar, Erik Perrins, and James P.G. Sterbenz. A cross-layered protocol architecture for highly-dynamic multihop airborne telemetry networks. In *International Telemetry Conference (ITC) 2008*, San Diego, CA, October 2008.
- [LeF01] William LeFebvre. CNN.com: facing a world crisis, 2001. <http://www.tcsa.org/lisa2001/cnn.txt>.
- [mic] Microsoft Security Notification Service. <http://www.microsoft.com/security>.
- [MR04] Jelena Mirkovic and Peter Reiher. A taxonomy of ddos attack and ddos defense mechanisms. *SIGCOMM Comput. Commun. Rev.*, 34(2):39–53, 2004.
- [NCC08] RIPE NCC. YouTube Hijacking: A RIPE NCC RIS case study, 2008. <http://www.ripe.net/news/study-youtube-hijacking.html>.
- [Pep09] Ivan Pepelnjak. Oversized AS paths: Cisco IOS bug details, 2009. <http://blog.ioshints.info/2009/02/oversized-as-paths-cisco-ios-bug.html>.
- [san] SANS Institute. <http://www.sans.org>.
- [Sch99] Bruce Schneier. Attack trees. *Dr Dobb's Journal*, 24(12), December 1999.
- [SS04] M.I. Steinder and A.S. Sethi. A survey of fault localization techniques in computer networks. *Science of Computer Programming*, 53(2):165–194, November 2004.
- [Tow88] Patrick A. Townson. The great fire. *Telecom Digest*, vol.8 iss.76, 10 May 1988, May 1988. <http://massis.lcs.mit.edu/archives/history/fire.in.chicago.5-88>.
- [Var05] Kazys Varnelis. The centripetal city: Telecommunications, the internet, and the shaping of the modern urban environment. *Cabinet Magazine* 17, 2004/2005.
- [Ves87] William E. Vesely. *Fault Tree Handbook*. Nuclear Regulatory Commission, ISBN-10: 0160055822, 1987.
- [Zmi09] Earl Zmijewski. Longer is not always better, 2009. <http://www.renesys.com/blog/2009/02/longer-is-not-better.shtml>.

A Challenge Catalog

Challenge	Name	Faulty interface clock
Classification	Category	Component Fault
	Scenario	any
Characteristics	Description	The clock on the interface is running out of sync compared to other interface clocks. Therefore, signal sampling at this interface can not be done properly.
	Scope	PHY layer
	Potential Impact	reduced bandwidth
Details	Parameters	Deviation of clock
	Symptoms	Corrupted frames

Table 6: Faulty interface clock

Challenge	Name	Broken antenna
Classification	Category	Hardware Destruction
	Scenario	Wireless networks
Characteristics	Description	The antenna is damaged accidentally or maliciously preventing the reception and transmission of signal. The challenge lasts until the antenna is repaired.
	Scope	PHY Layer
	Potential Impact	Loss of connectivity
Details	Parameters	Duration
	Symptoms	No incoming frames

Table 7: Broken Antenna

Challenge	Name	Signal attenuation
Classification	Category	Communication Environment
	Scenario	Wireless communication
Characteristics	Description	The signal strength received decreases due to attenuation caused by rain, fog, obstacles or similar. This effectively reduces the communication range.
	Scope	MAC layer
	Potential Impact	Degraded connectivity
Details	Parameters	Duration of attenuation period, attenuation factor
	Symptoms	Erroneous MAC frames

Table 8: Signal attenuation

Challenge	Name	Routing loop
Classification	Category	Human Mistake
	Scenario	any network
Characteristics	Description	A mis-configuration of a router using static routes leads to packets looping between two to more routers until they are discarded when the TTL is decremented to zero.
	Scope	Network layer
	Potential Impact	Loss of network connectivity
Details	Parameters	
	Symptoms	TTL equals zero at intermediate node, TTL exceeded error messages

Table 9: Routing loop

Challenge	Name	De-association attack
Classification	Category	Malicious attack
	Scenario	
Characteristics	Description	The attacker sends spoofed de-association packets to the access point preventing the victim to connect to the access point. The victim has connectivity only for short intervals: the time between a successful association request and the next de-association message.
	Scope	MAC layer
	Potential Impact	Degraded connectivity
Details	Parameters	Attack packet interval
	Symptoms	Increased association resets at access point, association failures at victim

Table 10: Distributed denial of service attack (DDoS)

Challenge	Name	ARP storm
Classification	Category	Unusual but legitimate request for service
	Scenario	any
Characteristics	Description	Worm propagation in local networks can cause a significant increase in ARP request and response messages.
	Scope	MAC
	Potential Impact	Degraded bandwidth
Details	Parameters	Packets per second, length of packet
	Symptoms	Increase in number of ARP packets, frequent changes of ARP table

Table 11: ARP storm

Challenge	Name	TCP connection establishment failure
Classification	Category	Failure of provider service
	Scenario	TCP
Characteristics	Description	Failures on the PHY (broken antenna), MAC (signal attenuation) or network layer prevent the TCP handshake. Hence TCP gives up the connection establishment after two re-tries and reports a failure to the application.
	Scope	Transport layer
	Potential Impact	No access to remote service
Details	Parameters	
	Symptoms	Missing SYN/ACK message, Failure report from dependent service

Table 12: TCP connection establishment failure