Strategies for Network Resilience: Capitalising on Policies

Paul Smith[1], Alberto Schaeffer-Filho[1], Azman Ali[1], Marcus Schöller[2], Nizar Kheir[3], Andreas Mauthe[1], and David Hutchison[1]

[1] Computing Department, Lancaster University, UK
{p.smith,asf,a.ali,andreas,dh}@comp.lancs.ac.uk
[2] NEC Laboratories Europe, Heidelberg, Germany
marcus.schoeller@neclab.eu
[3] France Télécom R&D Caen, 14066 CAEN, France
nizar.kheir@orange-ftgroup.com

Abstract. Networked systems are subject to a wide range of challenges whose nature changes over time, including malicious attacks and operational overload. Numerous mechanisms can be used to ensure the resilience of networked systems, but it can be difficult to define how these mechanisms should be configured in networks that support many services that have differing and shifting requirements. In this paper, we explore the potential benefits of using policies for defining the configuration of mechanisms for resilience. We discuss some of the difficulties of defining configurations, such as identifying conflicts, and highlight how existing policy frameworks could be used or extended to manage this complexity.

1 Introduction

The cost of failure of communication systems can be extremely high, as we depend on them to support many aspects of our daily life. Developing strategies to ensure the resilience of networked systems is of primary importance, but the challenges these systems are subject to are wide-ranging and change over time. These include component faults and mis-configurations, as well as operational overload and malicious behaviour from intelligent adversaries.

In this paper, we explore the use of policies to define configurations of mechanisms that can ensure the resilience of networked systems. The configuration of resilience mechanisms via policies is considered in the context of a general high-level strategy for resilience, called $\mathsf{D}^2\mathsf{R}^2$ + DR – *Defend*, *Detect*, *Remediate*, *Recover*, and *Diagnose* and *Refine* [1]. We believe that using policies is beneficial for a number of reasons: we de-couple the implementation of the mechanisms from the strategy used to enable resilience, which is a desirable property considering the changing nature of challenges. Furthermore, policy frameworks may assist in tackling a number of challenging problems in defining resilience strategies for multi-service networks: we are specifically interested in deriving concrete configurations from high-level requirements, identifying conflicting configurations, and evolving configurations over time in response to the changing nature of challenges and requirements.

B. Stiller and F. De Turck (Eds.): AIMS 2010, LNCS 6155, pp. 118–122, 2010.

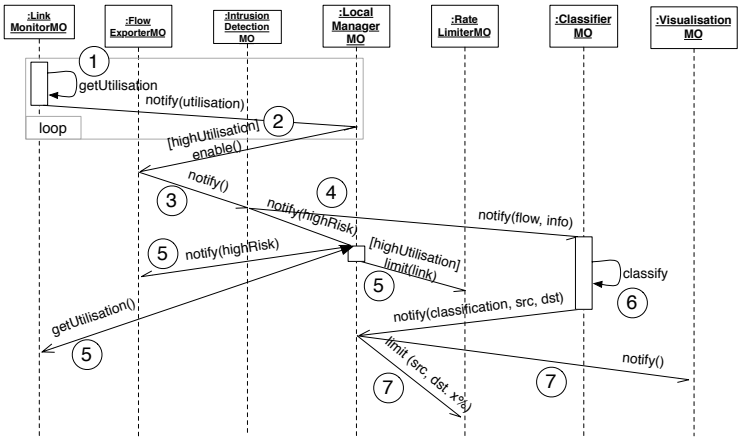

Fig. 1. Strategy for network traffic volume resilience

2 Policy-Based Resilience Strategy

One of the key problems related to resilience in networks is to discriminate operational overload due to legitimate service requests from malicious attacks, such as a *Distributed Denial of Service* (DDoS), and then apply adequate countermeasures [2]. To confront these challenges, we defined a strategy which relies on a number of mechanisms that must co-operatively enforce the resilience of the network, including a *flow exporter*, a *rate limiter*, and an *anomaly classifier* (Fig. 1). We use policies to configure and coordinate the interactions between these mechanisms. For example, specific root causes will require distinct remediation strategies, and, when a flow is classified as a possible DDoS attack, a preventive rate limiting action may be applied (*Step 7*). By having a resilience strategy implemented with the aid of policies, as opposite to having it hardcoded, one can easily change it by adding or removing policies, thereby permitting the modification of the strategy during run-time. This is of particular importance to us, as strategies for resilience are subject to frequent modifications, due to changes in requirements, context changes or new types of challenge.

3 Complexities of Defining Configurations

We highlight here where support can be found in policy-based management frameworks to address the complexities of defining configurations for resilience.

3.1 Deriving Configurations from High-Level Requirements

We assume policies will realise a high-level requirement to ensure resilience, e.g., in terms of the availability of a server farm and the services it provides. However, it is not clear if a resilience strategy such as the one in Fig. 1 is sufficient to ensure that a given high-level goal, e.g., defined in a SLA, is met. Moreover, complex

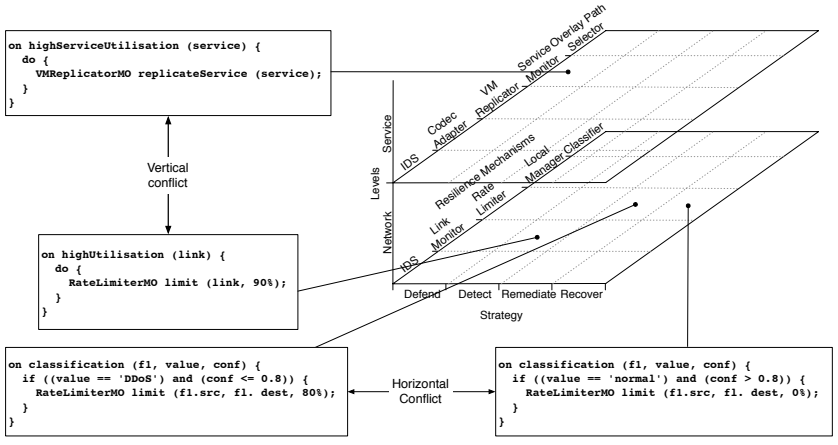

Fig. 2. Defining configurations for resilience is a multi-level problem, with vertical configuration conflicts across levels, and horizontal conflicts along the $\mathsf{D}^2\mathsf{R}^2$ strategy

scenarios would make deriving concrete policies by hand intractable. We seek to derive implementable policy configurations from high-level specifications and intend to build on techniques, such as [3], which apply goal elaboration and refinement of QoS requirements into the policy configuration of routers.

3.2 Identifying and Resolving Conflicting Configurations

In complex multi-service networks, conflicts can lead to the resilience requirements of a set of services being met unnecessarily at the expense of another set, or no requirements being met for any service. Conflicts can manifest in a number of ways: ***vertically*, across protocol levels**, in the presence of concurrent challenges – e.g., a flash crowd and a DDoS attack. Because of a DDoS attack, rate limiting may be started on routers; a *network-level* mechanism. During a flash crowd, a service could be replicated to another server farm; a *service-level* mechanism. However, due to the naïve rate limiting, replicating a service could make the resource starvation situation worse. Another type of conflict may occur ***horizontally*, along the $\mathsf{D}^2\mathsf{R}^2$ strategy**. Consider an attack targeted at both a server farm and a corporate customer. Attack traffic could saturate the links that provide access to a core network, making *push back* of malicious traffic to the Internet gateways desirable. Detection mechanisms at the server farm may determine that a node has ceased to behave maliciously, and initiate a *recovery* configuration for that node by stopping a rate limiter. However, the node may still be behaving maliciously in relation to the corporate customer, and recovery could inadvertently disengage the remediation configuration for that network. Policies that demonstrate these conflicts are shown in Fig. 2.

Policy analysis can help to ensure the correct specification of resilience strategies, in particular in terms of *dominance* and *coverage* checks [4]: the former could be applied in multi-level analysis, to ensure that mechanisms at one level

do not render mechanisms at another level redundant, the latter could be used for the analysis of configurations at the same level, e.g., conditions or range of values where mechanisms are not co-ordinated properly.

3.3 Learning Resilience Behaviour

Resilience configurations will need to evolve over time because the nature of attacks may change and new customer agreements may cause high-level priorities to shift. Furthermore, a strategy may prove to be sub-optimal or incorrect. To assist with this, we can benefit from existing research on policies. Typically, policy-based learning relies on the use of logical rules for knowledge representation and reasoning, as policies can be easily translated into a logical program [5]. Rules can be iteratively amended to better reflect resilience practices, based on how successful previous attempts to mitigate a challenge were. Similarly, the system must be able to learn entire new rules, for example, that during the football league final, high link utilisation is better remediated with the replication of the server streaming the live match, rather than simply rate limiting link capacity.

4 Conclusions

Network resilience is difficult to ensure because the configuration of systems is complex, spans across several levels, and is subject to a wide range of challenges. Policies provide flexibility in the configuration of the components that implement this strategy, as forms of detection and remediation are subject to frequent modifications. We examined the applicability of policies to mitigate high traffic volume challenges and highlighted how policy-based approaches can assist in making the problem more tractable. Future work will investigate how these policy techniques can be extended.

Acknowledgment

The research presented in this paper is partially funded by the European Commission in the context of the Research Framework Program Seven (FP7) project ResumeNet (Grant Agreement No. 224619). This work has also been supported by the EPSRC funded India-UK Advance Technology Centre in Next Generation Networking. The authors are grateful to Angelos Marnerides for insights relating to detection approaches.

References

1. Sterbenz, J.P., Hutchison, D., Çetinkaya, E.K., Jabbar, A., Rohrer, J.P., Schöller, M., Smith, P.: Resilience and Survivability in Communication Networks: Strategies, Principles, and Survey of Disciplines. Computer Networks: Special Issue on Resilient and Survivable Networks, COMNET (to appear, 2010)

2. Peng, T., Leckie, C., Ramamohanarao, K.: Survey of network-based defense mechanisms countering the DoS and DDoS problems. ACM Comp. Surv. 39(1) (2007)
3. Bandara, A.K., Lupu, E., Russo, A., Dulay, N., Sloman, M., Flegkas, P., Charalambides, M., Pavlou, G.: Policy Refinement for IP Differentiated Services Quality of Service Management. IEEE Trans. on Network and Service Management 3(2) (2006)
4. Agrawal, D., Giles, J., Lee, K.W., Lobo, J.: Policy ratification. In: POLICY 2005, Washington, DC, USA, pp. 223–232. IEEE Computer Society Press, Los Alamitos (2005)
5. Corapi, D., Ray, O., Russo, A., Bandara, A., Lupu, E.: Learning rules from user behaviour. In: 2nd Int. Workshop on the Induction of Process Models (2008)