Understanding risk for network resilience

Paul Smith, Marcus Schoeller (NEC), and David Hutchison

p.smith@comp.lancs.ac.uk

Multi-service Networks
July 2009

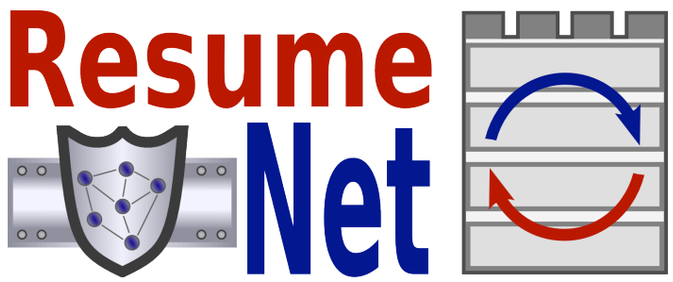

What is network resilience?

- The ability of a networked system to provide an acceptable level service in light of *challenges*
- Challenges
  - Component faults
  - Hardware destruction
  - Human mistakes
  - Malicious attacks
  - ....

Nothing comes for free

- Providing network resilience has a cost
  - We need systems to do this
- We are resource constrained
  - $£€, cpu cycles, bandwidth

Need to prioritise and focus efforts

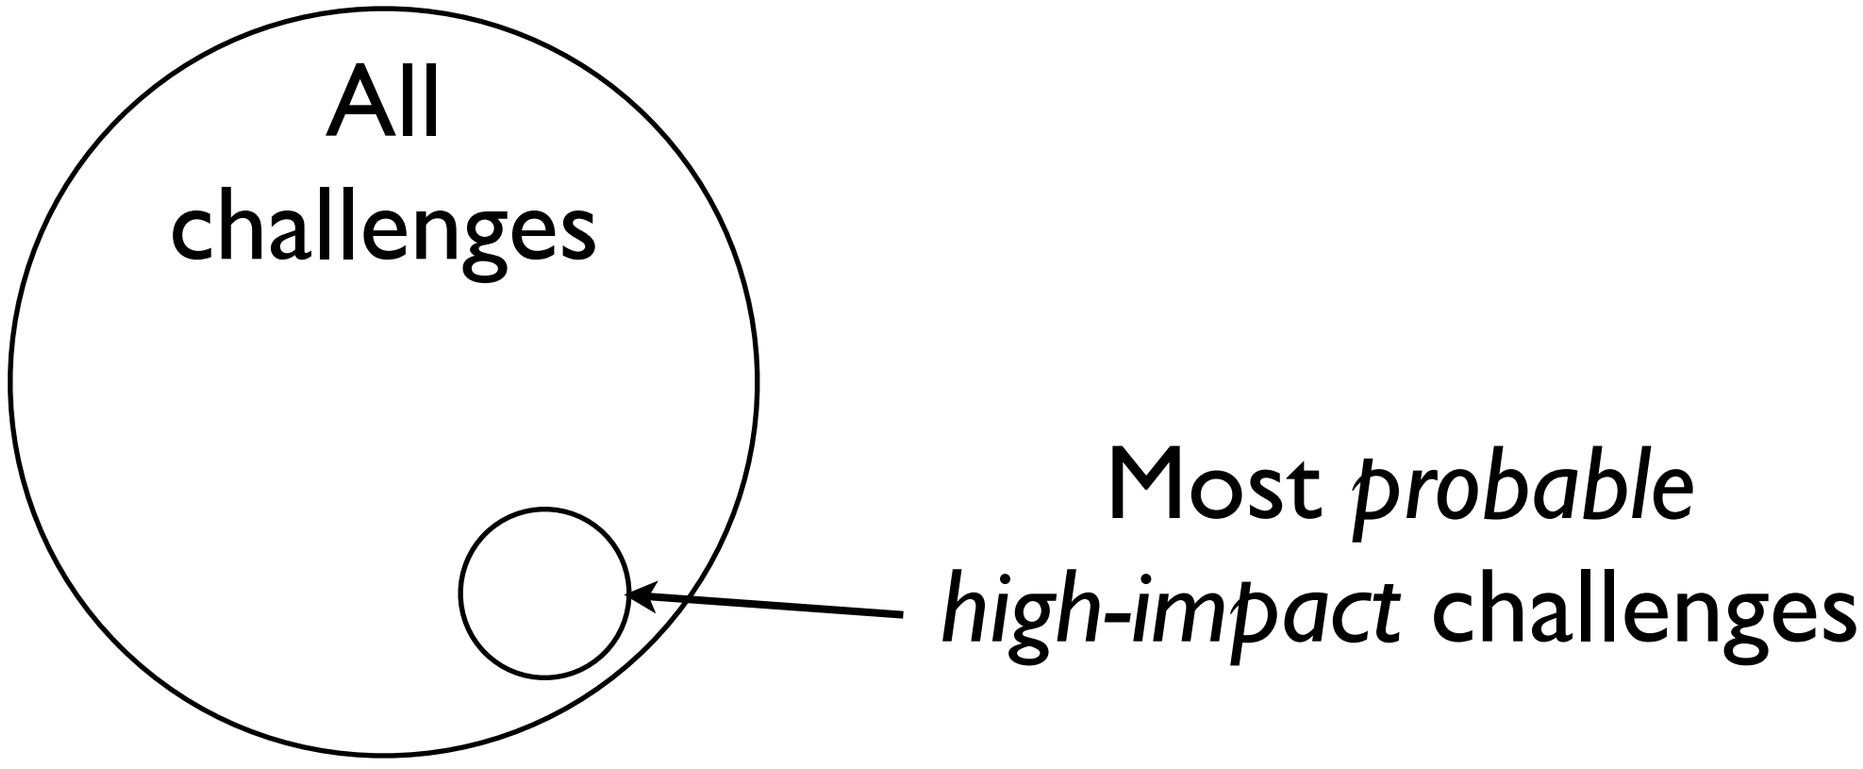

Identifying *critical* challenges

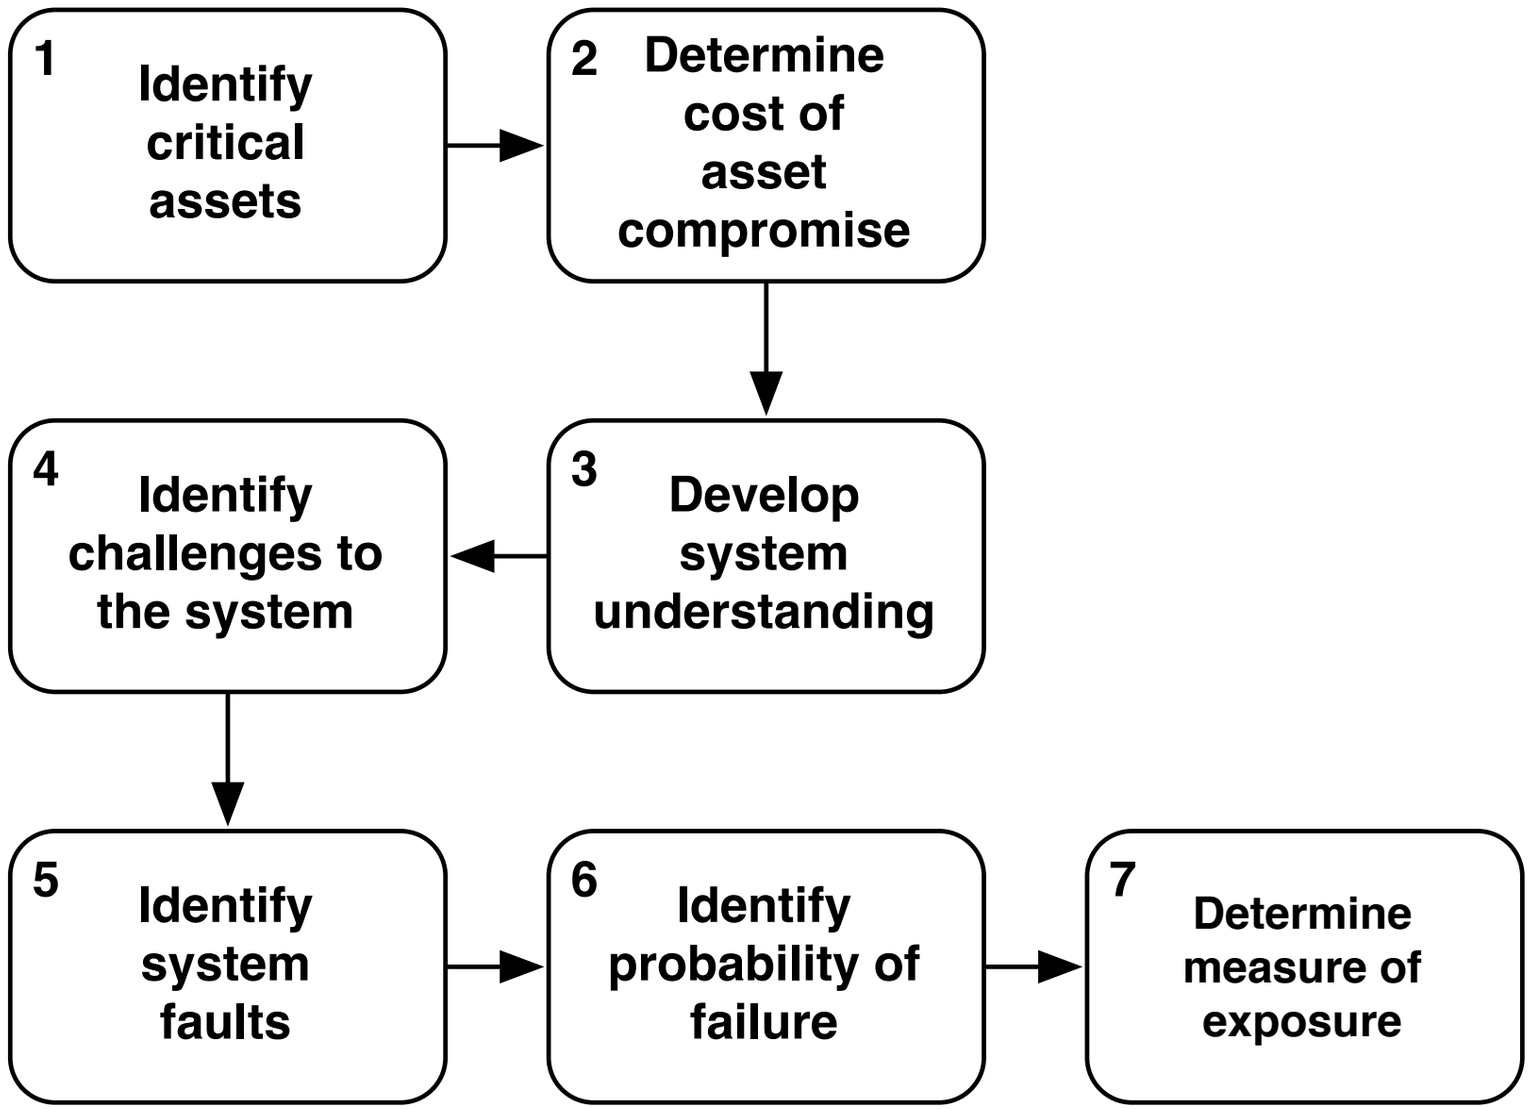

exposure = cost x probability

What's difficult about this?

- Determining reliable measures for challenge occurrence probabilities [and the probability of that leading to failure]
- Quantifying the impact of a challenge

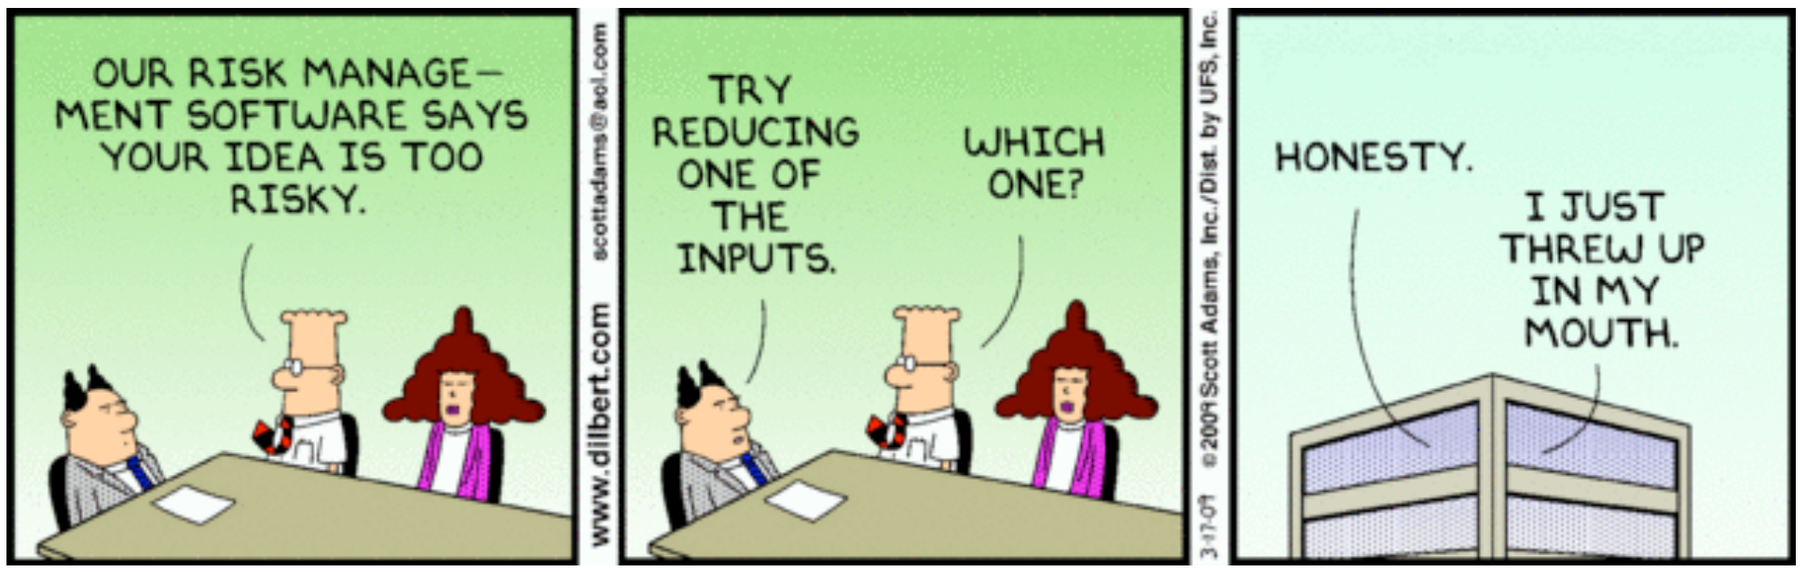

Challenge: Getting reliable numbers

- Off-line analysis
  - Advisories are useful (e.g., www.cert.org)
  - Fault and attack tree analysis
    - Issues of scalability because of complexity
  - Simulation
    - Need to develop good challenge and fault models
- Record monitoring data from on-line system
  - Classify challenges using machine learning
  - Introduces resource and security concerns

Need for on-line *impact* measures for automated mitigation

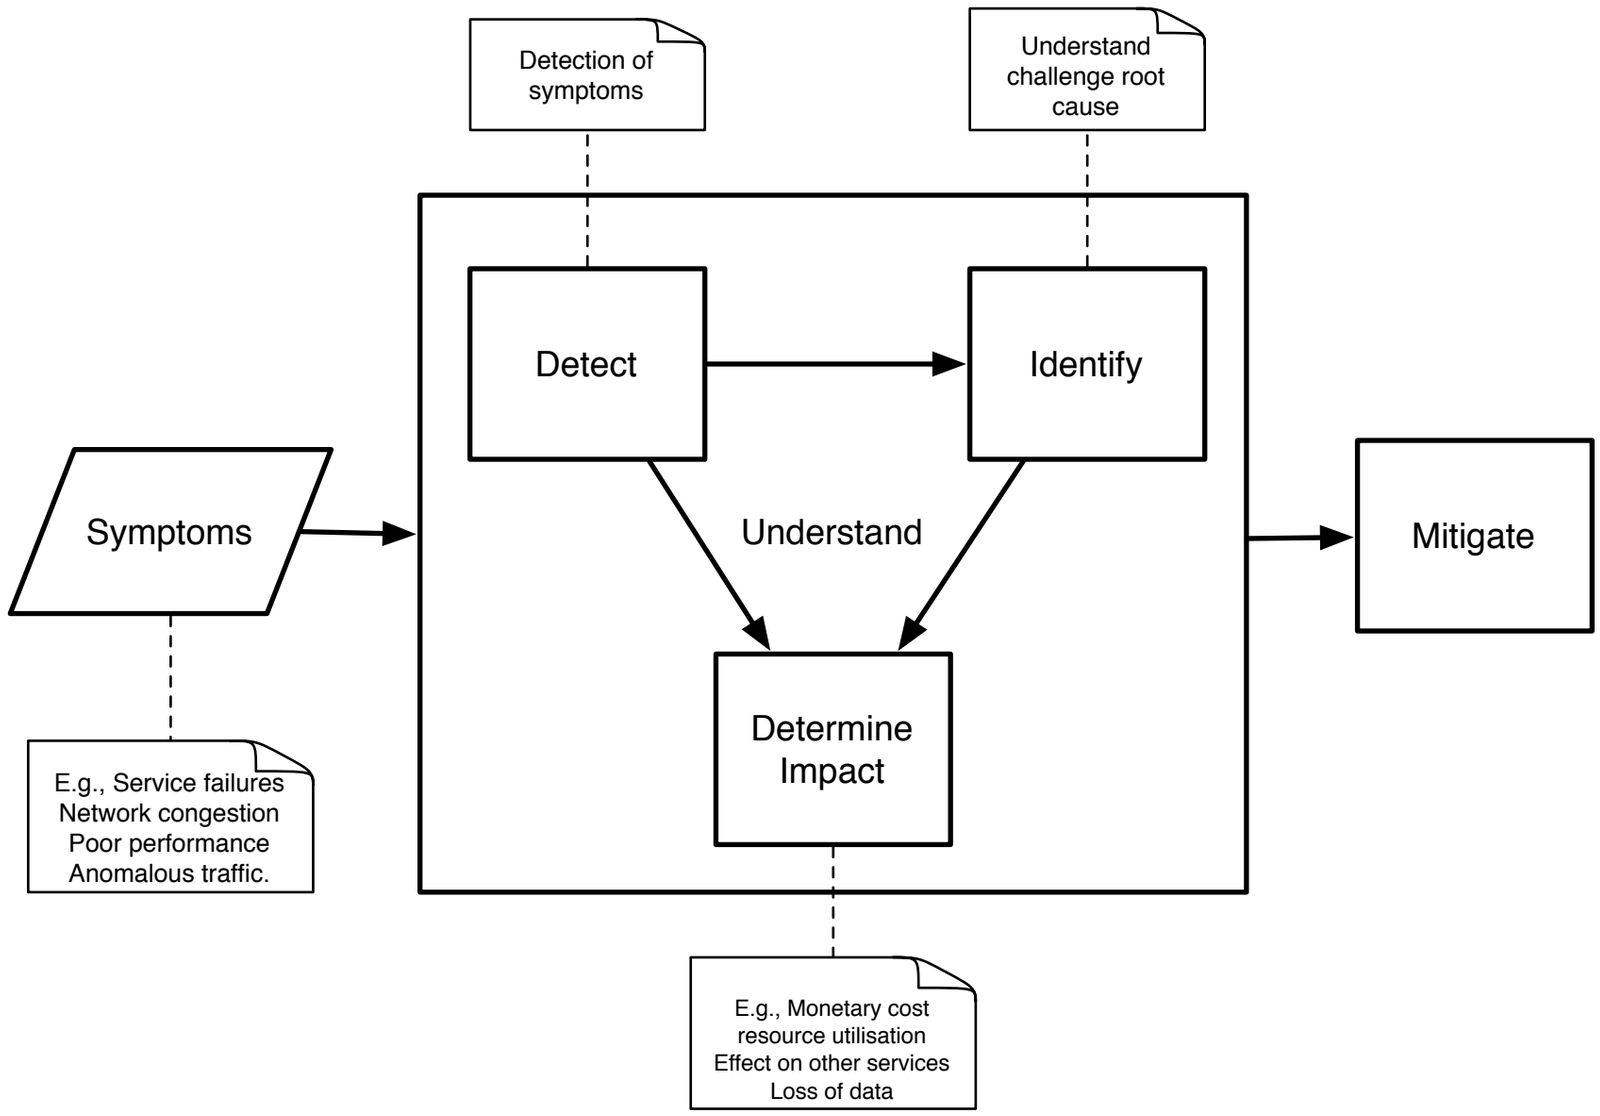

Conclusions

- To make best use of limited resources, we need to determine the *high-impact* challenges
- Getting good numbers for challenge probabilities and their impact is hard
- Some on-line components necessary, which has implications on system design
- For more information see www.resumenet.eu