

# Resilient SIEM Systems: How to Prevent the Naked Guard Syndrome

**Nuno Ferreira Neves**

**Univ. of Lisboa, Portugal**

**<http://www.di.fc.ul.pt/~nuno>**

>MASSIF<

## Security Information and Event Management (SIEM) Systems

- SIEM systems offer various capabilities for the collection and analysis of security information in networked infrastructures
  - integrating a large range of security and network tools
  - allowing the correlation of thousands of events and the reporting of attacks and intrusions in near real-time
- Main components
  - **Sensors:** collect information about the local environment and help on the responses; Can be: signature or anomaly-based IDS; vulnerability scanners; network profiling; inventory management
  - **Collectors:** gather and normalize the events generated by the sensors and any external systems; can be deployed standalone or in a Sensor
  - **Management server (or SIEM core engine):** event correlation and real-time monitoring; risk assessment; reporting and data mining; network profiling and inventory management

## Why do you need resilient operation?

- Decided to attack the mostly used open-source SIEM system, considering adversaries with different capabilities
  1. Found vulnerabilities (e.g., CSRF) that allowed removal of collectors
  2. Steal authentication credentials and authenticate into the engine
  3. Carry out session-hijacking between collectors and engine and send fake events
  4. More subtly, delay certain events to prevent correlation with other events, and avoid alarms
  5. Consistently perform TCP session RESET to cause a DoS
  6. ...
  
- Typical protection in existing tools based on current technology
  - SSL channels + some level of authentication
  - But, this is not enough!!!

## Why do you need resilient operation?



Instead of having a very feeble guard  
which is actually very watchful



One needs a watchful guard  
with a strong body armor !



## RATIONALE for an SIEM RESILIENT INFRASTRUCTURE

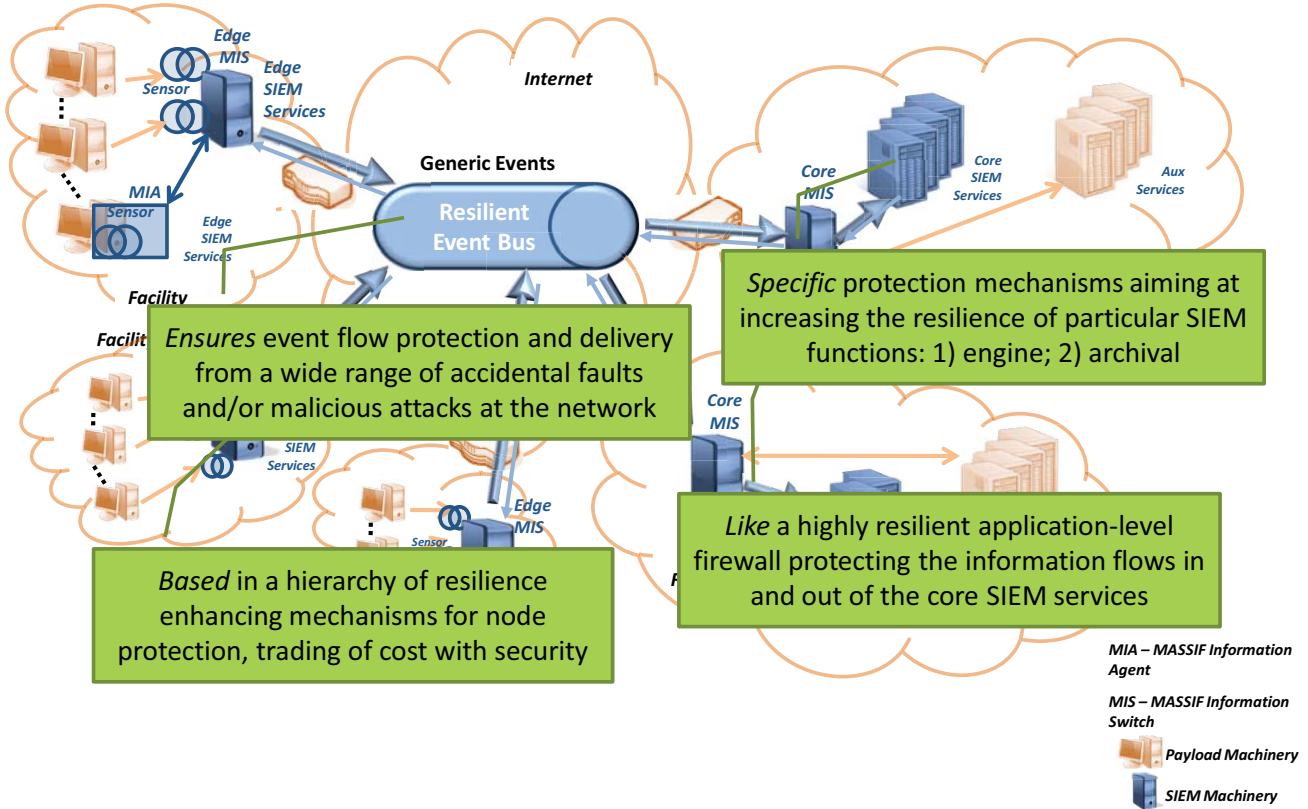
- **Complement classical security techniques with resilience mechanisms**
  - largely based on prevention, human intervention and ultimately disconnection
  - need for achieving tolerance, automation, and availability under attack
- **Promote automatic control of macroscopic information flows**
  - between layers of increasing resilience, from unprotected edge facilities up to the necessarily protected core processing units
- **Reconcile resilience with legacy preservation**
  - interfere as least as possible with the target system
  - SIEM integration should be as seamless and as transparent as possible
- **Avoid single points-of-failure**
  - at the edge level: protection the event collection
  - at the communication level: integrity and timeliness of the information flow
  - at the core processing level: availability and integrity of event processing
- **Secure timeliness in the presence of faults and attacks**
  - in different grades of real-time, from edge to core
  - detecting timing failures when timeliness enforcement is impossible

## Overview of the resilient MASSIF architecture

- **Main characteristics**
  - the architecture is laid down as a *sort of overlay* on the target, so as to preserve legacy but allow seamless integration
  - modeled pretty much as a SCADA system, producer-consumer system upstream, with low bandwidth commands downstream
- **Resilience procurement based on**
  - protecting crucial processing units
  - making the dissemination infrastructure itself resilient
  - implementing all functions in a modular way around conceptual devices called MASSIF Information Switches (MIS) and MASSIF Information Agents (MIA), respectively in HW and in SW

# General MASSIF Architecture

## Resilience solutions

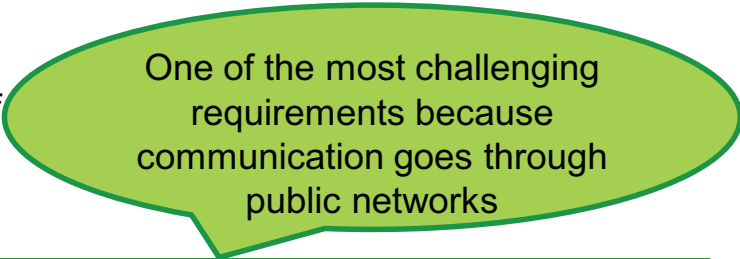


# FOCUSING ON PROTECTING THE COMMUNICATIONS

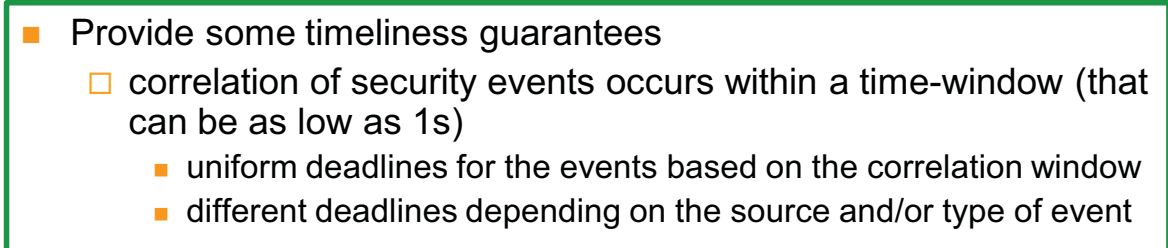


## Protecting the Communications

- Secure
  - integrity and authentication
  - confidentiality
- Reliable transmission of
  - accidental
  - malicious



One of the most challenging requirements because communication goes through public networks


- 
- Provide some timeliness guarantees
    - correlation of security events occurs within a time-window (that can be as low as 1s)
      - uniform deadlines for the events based on the correlation window
      - different deadlines depending on the source and/or type of event



## Objective

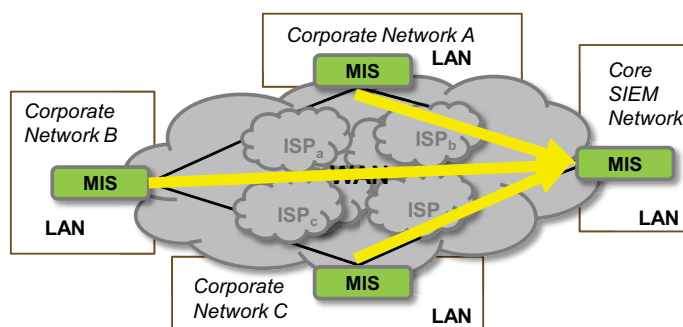
- Design a **practical** solution for timely and reliable communication with **high probability** taking into consideration
  - compatible with current deployments: should allow seamless integration without requiring major changes to the operation and organization of existing networks
  - no Internet changes: should not assume or require any special support from the underlying network, and therefore, timeliness has to be attained with best effort IP channels
  - cost consciousness: should take advantage of existing redundancy (e.g., due to over-provisioned multihoming connections), but should avoid the use of more links or costly dissemination operations (like flooding)
  - incremental integration: allow for a transition period with existing and new mechanisms

## Network Considerations (1)



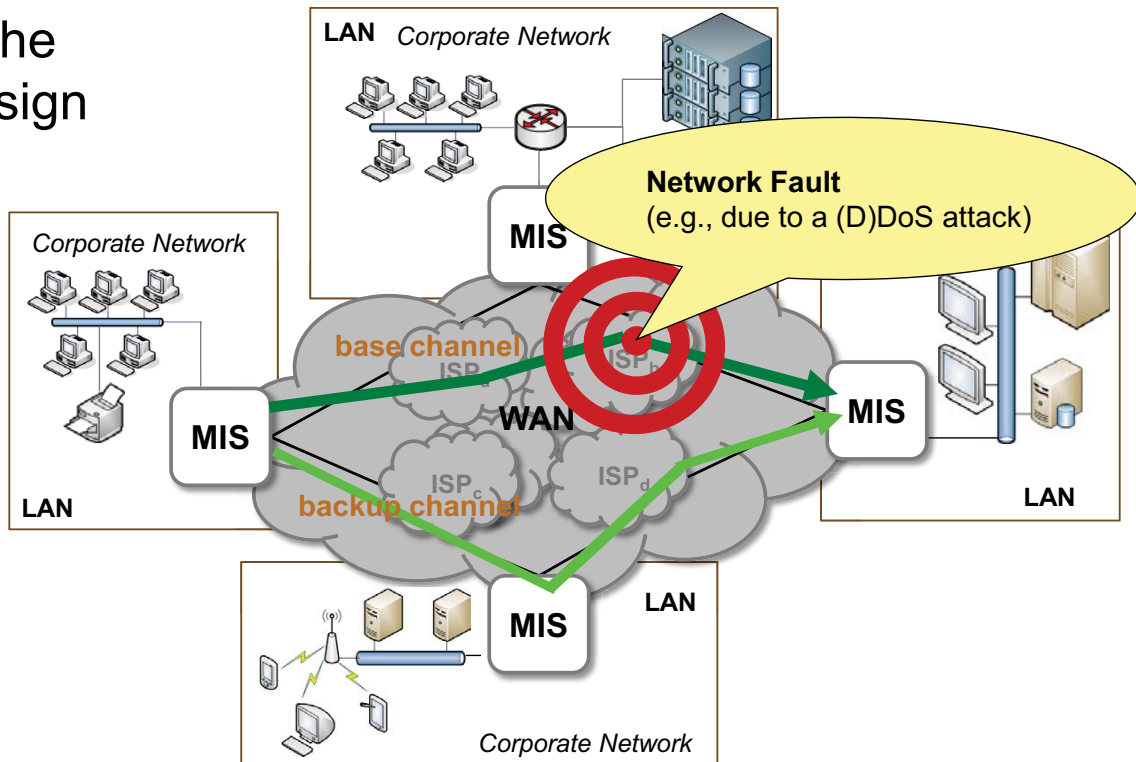
- There is a MIS at the boarder of each of the virtual n LANs
- MIS are reliable (due to replication, self-healing, ...)
- MIS can be multihomed, and therefore be connected by a number redundant ISPs
  - contracted for link independence to get some assurances of communication availability
  - but these connection can not be changed

## Network Considerations (2)



- Messages (that matter) have associated deadlines
- Closed environment with well-defined sources of legal traffic
- Mainly static
  - LANs are not added or removed too often
  - every LAN can be authenticated

# Rationale of the Design




## MIS Overlay Routing

- Fundamental ideas
  - One-hop source routing
  - Base channel + backup channels
  - Allow for a maximum number of retransmissions over diverse channels while ensuring the deadlines
- Three components of the solution
  - Measurements
  - Number of tries
  - Transmission strategy


Do not attempt to minimize latency, but deliver messages just-in-time

# Basic Measurements (1)

- Top level routing decisions are mainly based on the *latency* among MIS nodes and on *spatial redundancy*
  - Transmission time (TXT) between two MIS



- Channel correlation: number of routers shared by two channels



# Example of the Number of Tries

Channels table on a MIS  $i$  to reach MIS  $j$

0	TXT = 22
1	TXT = 36
2	TXT = 42
3	TXT = 60
4	TXT = 81
5	
bad channels	...
9	

$m$ , deadline = 100

One try:  $22 < 100$   
 Two tries:  $2 \cdot 36 + 22 = 94 < 100$   
 Three tries:  $2 \cdot 42 + 2 \cdot 36 + 22 = 178 > 100$

**Tries = 2**

$m'$ , deadline = 300

One try:  $22 < 300$   
 Two tries:  $2 \cdot 36 + 22 = 94 < 300$   
 Three tries:  $84 + 72 + 22 = 178 < 300$   
 Four tries:  $120 + 84 + 72 + 22 = 298 < 300$   
 Five tries:  $162 + 120 + 84 + 72 + 22 = 450 > 300$

**Tries = 4**

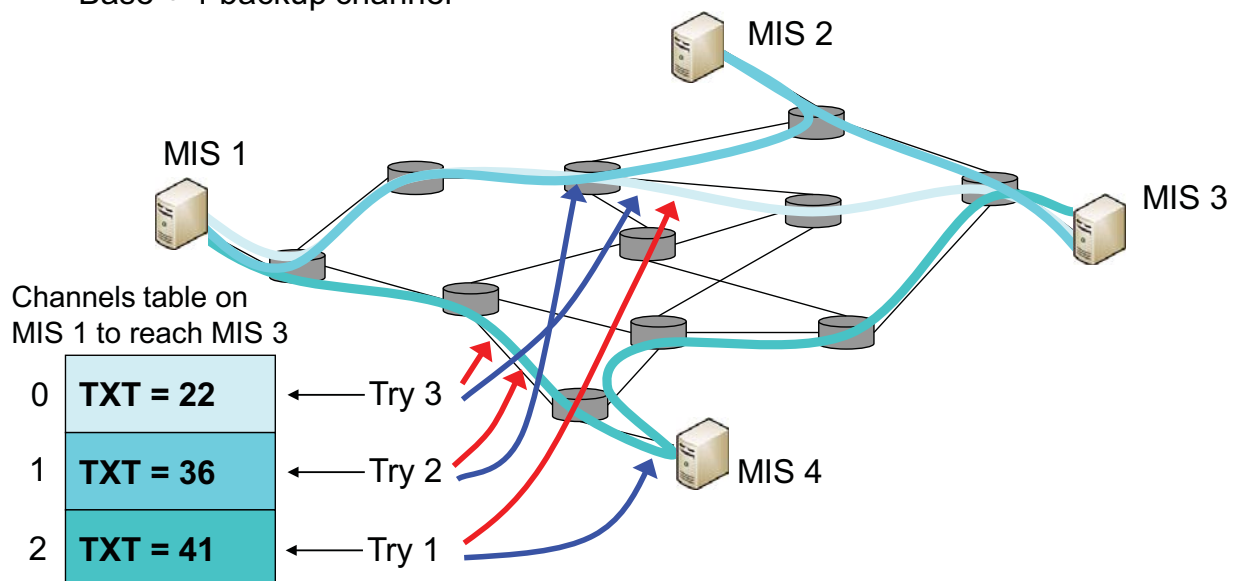


## Overlay Routing Strategy

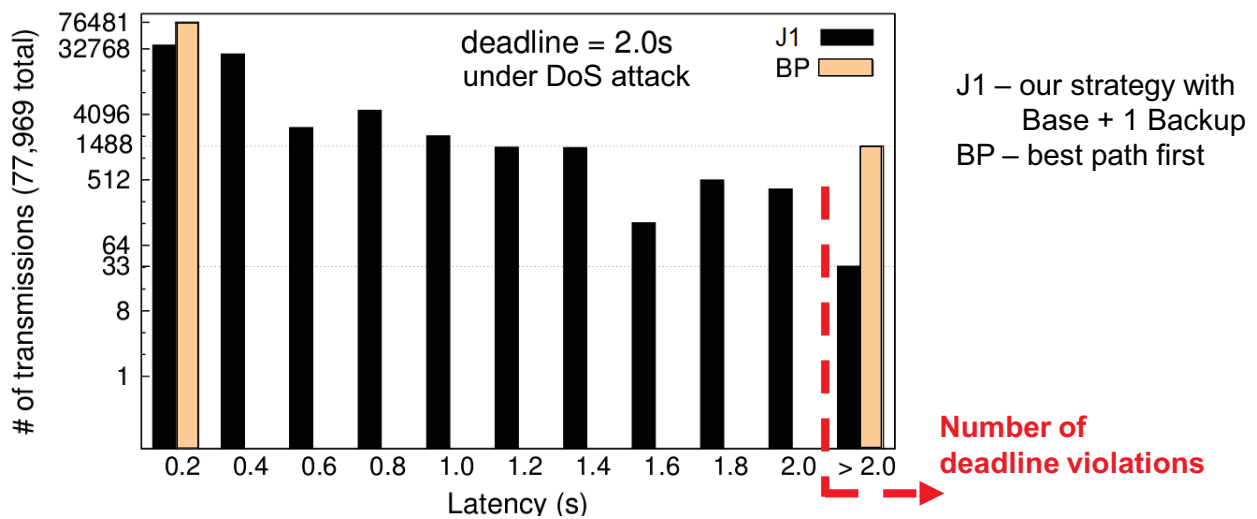
- We can determine the number of tries (a maximum number of retransmissions) to send a message
- **Main Idea:** in each try, use two types of channels to send a message
  - **Base channel:** the worst channel  $c$  that still allows the node to try the maximum number of faster channels if it fails  
*messages don't need to reach their destination fast, they need to arrive on time!*
  - **Backup channel(s):** some other  $B$  channels that can deliver the message on time and that have minimal correlation with the chosen base channel  
*IP networks offers no guarantee so we must take some preventive measures and use channel diversity*

## Example of the Overlay Strategy

- System with a single ISP connecting each MIS
- Message with deadline 180
- Base + 1 backup channel



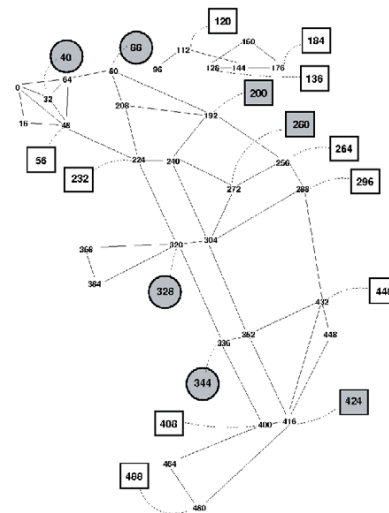
## Why not going through fastest channel first?



- Ends up missing much less deadlines (33 instead of 1488)
- Leaves faster channels for messages with tighter deadlines
- Achieves load balancing across the links

## Evaluation


- Uses a simulation model using the J-Sim network simulation tool
- Topology based on 31 routers and 51 direct channels
- The network is duplicated to simulate multihoming
- Links with 50 ms of latency and bandwidth of 1Gbps
- There are aperiodic and periodic messages with deadlines of 1, 2 and 4 seconds
- Faults are injected following models reported in the literature
  - in each run, 74 faults are injected in each ISP backbone
  - both accidental and DoS are modeled



## Strategies under Evaluation

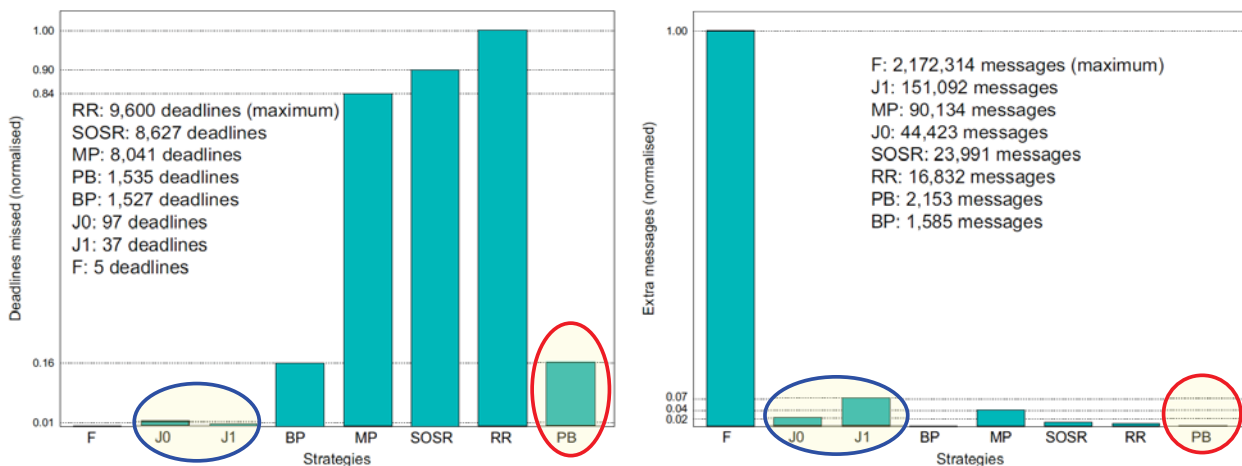
- **J0**: only the base channel
- **J1**: base + 1 backup channel
- **Flooding (F)**: all messages are sent to all channels  
(most reliable solution... uses all redundancy available)
- Multihoming (two channels):
  - **Round-robin (RR)**: direct channels used in round robin fashion
  - **Primary-backup (PB)**: if some direct channel fails, use another
- Overlay (at most one relay node):
  - **Best path (BP)**: always sent through the best non-failed channel
  - **Multi path (MP)**: send through the direct channel and a randomly selected channel (direct or not)
  - **Hybrid (SOSR)**: send first in a random direct channel, then if there is a failure use 4 random channels (direct or not)

## Evaluation: Accidental Faults



- Our strategy misses no deadlines requiring only 4% of the extra messages used in flooding
- Primary-backup (as used today in certain companies) is sufficient to deal with most accidental faults (only misses a few deadlines)

## Evaluation: Malicious Attacks



- Flooding misses 5 deadlines
- Our strategy misses only 37 (or 97) deadlines, using less than 7% of the extra messages used in flooding
- Primary-backup misses a significant number of deadlines (1535)

## Summary

- SIEM systems currently act as the guard of our networks
- The guard is “naked” because only limited security support is available, which will not be able to withstand a serious attack
- SIEM systems can be made secure using a number of principles, including the protection of the communication flows
- An overlay routing strategy was presented for timely message delivery
  - employs one-hop source routing
  - explores redundancy both in time and space
  - delivers messages not fast but just-in-time



Thank you!

Questions?