

- 1 Graphical Abstract
- 2 **Malicious Changeload for the Resilience Evaluation**
- 3 **of Self-adaptive Authorisation Infrastructures**
- 4 Christopher Bailey, Rogério de Lemos

5 Highlights

6 **Malicious Changeload for the Resilience Evaluation**
7 **of Self-adaptive Authorisation Infrastructures**

8 Christopher Bailey, Rogério de Lemos

- 9 • Formal definition of a *malicious changeload*, describing scenarios of
10 abuse in access control that were used in the resilience evaluation of a
11 self-adaptive authorisation infrastructure.
- 12 • Definition of a generic *simulation-based* approach for evaluating the re-
13 silience of self-adaptive authorisation infrastructures under repeatable
14 conditions of system and environmental change.

15 Malicious Changeload for the Resilience Evaluation
16 of Self-adaptive Authorisation Infrastructures

17 Christopher Bailey, Rogério de Lemos*

18 *University of Kent, UK*

19 **Abstract**

20 Self-adaptive systems are able to modify their behaviour and/or structure in
21 response to changes that occur to the system, its environment, or even its
22 goals. In terms of authorisation infrastructures, self-adaptation has shown
23 to be a promising solution for enforcing access control policies and subject
24 access privileges when mitigating insider threat. This paper describes the
25 resilience evaluation of a self-adaptive authorisation infrastructure by sim-
26 ulating a case study related to insider threats. As part of this evaluation,
27 a malicious changeload has been formally defined in order to describe sce-
28 narios of abuse in access control. This malicious changeload was then used
29 to stimulate self-adaptation within a federated authorisation infrastructure.
30 The evaluation confirmed the resilience of a self-adaptive authorisation in-
31 frastructure in handling abuse of access under repeatable conditions by con-
32 sistentlly mitigating abuse under normal and high loads. The evaluation has
33 also shown that self-adaptation had a minimal impact on the authorisation
34 infrastructure, even when adapting authorisation policies while mitigating
35 abuse of access.

36 *Keywords:* self-protecting systems, authorisation infrastructures,
37 changeload, insider threats, autonomic computing, access control

38 **1. Introduction**

39 Self-adaptive systems are able to modify their behaviour and/or struc-
40 ture in response to changes that occur to the system, its environment, or
41 even its goals [26]. A self-adaptive authorisation infrastructure is a self-
42 adaptive system tailored to adapt, at run-time, access control policies and
43 their enforcement [30]. An important aspect when evaluating the resilience

44 of a self-adaptive authorisation infrastructure is to demonstrate its ability to
45 mitigate abuse in access control.

46 In this paper, we present a simulation-based approach for evaluating the
47 resilience of self-adaptive authorisation infrastructures under repeatable con-
48 ditions of system and environmental changes. In our evaluation, in addition
49 to observing performance as a measure of success, we also evaluate the im-
50 pact of self-adaptation as a means to mitigate potential attacks. Although
51 the goal of self-adaptation is to protect dynamically the authorisation infras-
52 tructures from attacks, self-adaptation measures may result in undesirable
53 states, which may include the loss of access to critical resources.

54 For demonstrating the proposed approach, we evaluate the resilience of
55 the Self-adaptive Authorisation Framework (SAAF) [3, 4, 6], whose goal is to
56 make existing authorisation infrastructures self-adaptable. This is achieved
57 by analysing potential attacks once they are detected, and synthesising ap-
58 propriate mitigation actions depending on the operating conditions of the
59 infrastructure. In terms of SAAF, this would also comprise the generation
60 and deployment of new access control policies at run-time, without any hu-
61 man interference. The premise is that, an organisation can benefit from the
62 properties of dynamic access control without the need to adopt new access
63 control models.

64 A common way for evaluating self-adaptive systems is through case stud-
65 ies [19]. They are used to represent environment and system changes, and
66 these are expected to stimulate self-adaptation, thus providing the basis for
67 evaluating the impact of adaptation. An advantage of using case studies is
68 that changes can be repeated to stimulate self-adaptation scenarios, thus al-
69 lowing the evaluation of impact from adaptation in a more consistent way.
70 For the evaluation of a self-adaptive authorisation infrastructure, we use a
71 fictitious case study describing a set of insider attacks within a federated
72 environment. The resilience evaluation of SAAF in a more realistic attack
73 scenario was performed using an ethical on-line game to gather insights on
74 how SAAF would react towards real malicious behaviour, and how malicious
75 users would behave in the presence of self-adaptation [8]. The motivation for
76 the study being reported in this paper is to evaluate whether the proposed so-
77 lution affects the performance of the overall authorisation infrastructure, and
78 to evaluate the effectiveness of the self-adaptive solution to handle malicious
79 behaviour.

80 This paper provides two key contributions:

- 81 • the definition of a generic approach for evaluating the resilience of
82 self-adaptive authorisation infrastructures. This is demonstrated by
83 deploying SAAF within a federated environment, thus showing how
84 SAAF handles and mitigates malicious behaviour, in the form of in-
85 sider threats, given the existence of non-cooperating third party organ-
86 isations;
- 87 • the definition of *malicious changeload* that drives stimulation of adap-
88 tation in response to malicious behaviour (i.e., abuse of access control).
89 The usefulness of malicious changeload in providing systematic means
90 for specifying repeatable behaviour is demonstrated by evaluating the
91 resilience of SAAF under various operational conditions.

92 The definition of malicious changeload in the context of self-protecting
93 systems extends that of changeload, defined for the resilience evaluation of
94 architectural-based self-adaptive systems [12], and which considered faults
95 as the only undesirable type of change. In this paper, malicious changeload
96 considers the abuse of access control in federated authorisation infrastruc-
97 tures. Regarding the resilience evaluation of self-protecting systems, to the
98 best of our knowledge this is the first work that uses the notion malicious
99 changeload from resilience benchmarking. This is an important concept if
100 repeatable conditions of system and environmental changes are necessary in
101 the benchmarking of self-protecting systems.

102 The rest of this paper is structured as follows. In Section 2, we present
103 some basic concepts related to self-adaptive authorisation infrastructures and
104 insider threats. Section 3 positions a motivating case study used as a basis for
105 the evaluation. Section 4 specifies the malicious changeload derived from the
106 case study. Section 5 describes a set of experiments and results that observes
107 the run-time stimulation of malicious changeload, and adaptation of a target
108 system. Section 6 reflects on the outcome of the experiments, along with
109 the benefits and challenges of self-adaptive authorisation. Related work is
110 presented in Section 7. Finally, in Section 8, a summary of the paper is
111 provided in addition to some insights regarding future work.

112 2. Background

113 In this background section, we briefly describe the Self-adaptive Autho-
114 risation Framework (SAAF) to be used in the resilience evaluation of self-
115 adaptive authorisation infrastructures, we provide some insight to insider

116 threats that are representative of malicious behaviour and defined as part of
117 malicious changeload, and then we provide a brief introduction to resilience
118 benchmark.

119 *2.1. Self-adaptive Authorisation Framework*

120 The basis our work is the Self-adaptive Authorisation Framework
121 (SAAF), whose goal is to make existing authorisation infrastructures self-
122 adaptable [3, 4, 6]. The motivation is that, authorisation infrastructures
123 maintained by organisations can benefit from the properties of dynamic ac-
124 cess control without the need to adopt new access control models. SAAF is
125 based on the MAPE-K feedback loop [24], which monitors the distributed
126 services of an authorisation infrastructure, and builds a model of the whole
127 system at run-time (i.e., deployed access control rules, assigned subject privi-
128 leges, and protected resources). Malicious user behaviour observed by SAAF
129 is mitigated through the generation and deployment of access control policies
130 at run-time, preventing any identified abuse from continuing. Adaptation at
131 the model level ensures that abuse can no longer continue. In addition, model
132 transformation supports the generation of access control policies from an ab-
133 stract access model. This enables the generation and deployment of policies
134 that are specific to different implementations of access control.

135 Figure 1 presents a conceptual design of SAAF in which the services of the
136 authorisation infrastructure, and their interactions, are represented as solid
137 lines, while the autonomic controller and its interactions are represented as
138 dash lines. The role of the SAAF controller is to monitor and adapt the
139 services of the authorisation infrastructure. The interactions between the
140 services are annotated with a sequence of events, and as it can be observed,
141 the autonomic controller does not affect the sequence of events related to
142 the functionality of the identity and authorisation services. The autonomic
143 controller affects only the properties associated these services, respectively,
144 identities and policies, as a means to mitigate malicious behaviour. A major
145 challenge while implementing SAAF is that no single service provides a com-
146 plete view of access control in terms of what users own in access rights, what
147 access control rules exist, and finally, how users are utilising access rights.

148 Figure 2 presents a detailed design of the SAAF controller. The SAAF
149 controller comprises the analysis, which generates new access control mod-
150 els, and the plan, which selects the most appropriate access control model
151 amongst the valid ones. For each identified attack, the SAAF controller se-
152 lects a subset of solutions applicable to a particular attack, which depends




Figure 1: SAAF conceptual design

153 on the type of attack and current access control model. At deployment-time,
 154 the SAAF controller is loaded with a set of predefined solutions that mitigate
 155 malicious events. The solutions match a finite set of actions that can be per-
 156 formed within the application domain, and are parametric in order to tailor
 157 the solutions to specific cases of insider attacks. Given a detected attack, a
 158 solution is selected from the following alternatives: 1) increasing, limiting or
 159 removing access rights owned by an individual, 2) increasing or limiting the
 160 scope of access defined by access control rules, 3) warning the individual(s)
 161 of their behaviour, and 4) monitoring the behaviour further. Associated with
 162 each solution there is a potential impact. Depending on the type of action
 163 invoked, it can cause either negligible or severe consequences to the system
 164 (which may be warranted given the severity of the attack detected). Which
 165 solution is selected depends on how severe the SAAF controller deems the
 166 identified malicious behaviour to be. For example, what is the number of
 167 non malicious users impacted negatively by the solution (thus losing access
 168 to resources). High severity may be justified for cases when many users are
 169 identified as being malicious in relation to specific resources or roles. In
 170 this cases, changing access control rules provides a more effective means to
 171 responding to attacks.

172 In its current form, SAAF ensures that whatever adaptations take place
 173 will not break conformance to the service’s implemented access control
 174 methodology - in our case Attribute Based Access Control (ABAC), nor
 175 conflict with application domain requirements (e.g., ensure access to busi-
 176 ness critical systems). To implement ABAC, we provide an identity service

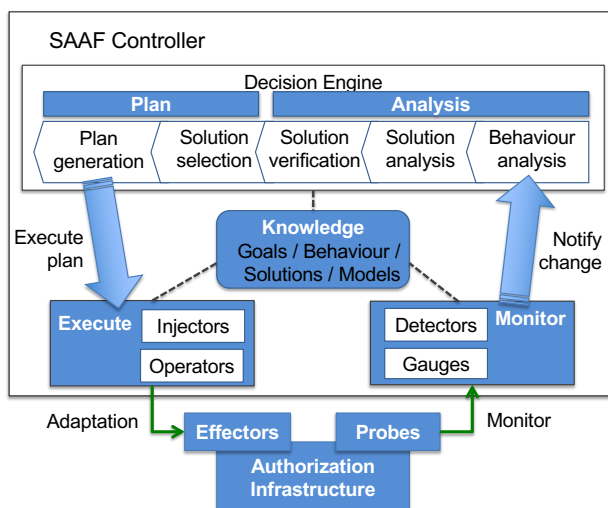


Figure 2: SAAF controller design

177 referred as LDAP [25], which is a directory service commonly used to hold
 178 information (including user roles) about users within an organisation. To
 179 generate ABAC access control decisions, based on roles owned by users, we
 180 use a standalone service authorisation service, known as PERMIS [16].

181 2.2. Insider Threats

182 Insider threat refers to an organisation’s risk of attack by their own users
 183 or employees [13]. This is particularly relevant to access control, where the
 184 active management of authorisation has the potential to mitigate and prevent
 185 users from abusing their own access rights to carry out attacks.

186 A common characteristic of insider threat is that malicious insiders use
 187 their knowledge of their organisation’s systems, and their assigned access
 188 rights, to conduct attacks. This places a malicious insider in a fortuitous
 189 position, whereby the insider (as an authorised user) can cause far greater
 190 damage than an external attacker, simply due to their access rights [14].
 191 Such form of attack is representative of the attacks that many organisations
 192 consider to be most vulnerable from: the abuse of privileged access rights by
 193 the employees of an organisation [34]. Unless additional measures are put
 194 into place, malicious insiders can abuse existing security measures, where
 195 current approaches fail to robustly adapt and respond to the unpredictable
 196 nature of users. Whilst there are a number of novel techniques that enable

197 the detection of insider threat [22, 32, 38], there is little research that uses
198 such techniques within an automated setting, like our proposed approach
199 that relies on self-adaptation.

200 *2.3. Resilience Benchmarking*

201 A benchmark is a standard procedure that allows characterising and com-
202 paring systems or components according to specific characteristics (e.g., per-
203 formance, dependability) [23]. Previous work on computer benchmarking can
204 be divided in three main areas: performance benchmarking [20], dependabil-
205 ity benchmarking [23], and security benchmarking [27].

206 In the context of self-protecting systems, resilience benchmarking rep-
207 represents a step further since it needs to consider system and environment
208 dynamics, although it is bound to encompass techniques from these previ-
209 ous efforts due to its inherent relation to performance, dependability and
210 security. Given an application domain (that specifies the target systems
211 and its environment), a resilience benchmark should provide generic ways
212 for characterising system behaviour in the presence of changes, allowing to
213 compare similar systems quantitatively. If a system is effective and efficient
214 in accommodating or adjusting to changes, avoiding successful attacks as
215 much as possible and operating as close as possible to its defined goals, it
216 is reasonable to consider the system to be resilient. This capability can be
217 benchmarked by submitting the system to various types of changes, time and
218 resources dedicated to mitigate them, as well as, the impact of this process
219 in the fulfilment of the system goals. As the changes affecting the system
220 may lead to the degradation of its performance, without leading necessarily
221 to security breaches, we need to assess variations in the properties of inter-
222 est when the system is under varying environmental conditions, in order to
223 characterise its behaviour from a resilience perspective.

224 **3. Case Study: LGZLogistics**

225 Given the challenges in obtaining detailed data on actual cases of insider
226 threats, this fictitious case study draws upon several historical cases dis-
227 cussed in the CERT guide to insider threat [13]. In this paper, we consider
228 data theft attacks that are performed to a fictitious logistics company, called
229 *LGZLogistics*, representing a service provider and identity provider within a
230 federated authorisation infrastructure. Malicious behaviour is conducted by
231 disgruntled employees of the logistics company, as well as employees of an

232 external *Trusted Business Partner* (TBP) [13]. The role of a TBP is key to
 233 this case study, as it is representative of the relationship a service provider
 234 organisation has with an identity provider organisation (e.g., *LGZLogistics*
 235 trusts the TBP to provide IT help desk services).

236 The case study focuses on two areas of insider threat that organisations
 237 are highly vulnerable to: the abuse of user access rights by employees of the
 238 organisation, and the abuse of access rights by TBP organisations [34].

239 3.1. Context and Architecture

240 *LGZLogistics* portrays a small to medium sized company of 1000 employ-
 241 ees, ten of which are IT staff that support and administer a set of protected
 242 resources. These resources are protected via an instantiated Attribute Based
 243 Access Control (ABAC) model, in the form of subject attribute assignments
 244 within identity services, and an access control policy within an authorisation
 245 service.




Figure 3: *LGZLogistics* authorisation infrastructure architecture

246 *LGZLogistics* maintains a SimpleSAML.php [37] identity service **lgzIS** to
 247 authenticate its subjects (employees), and issue access rights (as credentials).
 248 The organisation also maintains a PERMIS standalone authorisation service
 249 **as** [36], to authorise subject access to its resources. These resources include
 250 an employee database **empDB**, and a bespoke logistics tool **lgT**. The employ-
 251 ee database contains personnel information about the logistic company’s
 252 employees, which is required for general IT help desk enquires.

253 *LGZLogistics* uses the authorisation service **as** to authorise access for its
 254 own subjects, as well as subjects from a second offshore contractor organ-
 255 isation (a TBP). *LGZLogistics* trusts the contractor organisation to issue
 256 access rights to their subjects, as part of a business contract for providing
 257 IT help desk services. As such, the contractor organisation also operates
 258 a SimpleSAMLphp identity service **conIS** that manages its own employees’
 259 access rights to the requesting service providers (i.e., *LGZLogistics*). As part
 260 of their contract, subjects from the contractor organisation are permitted ac-

261 cess to `empDB` to facilitate help desk duties. Access for subjects from either
262 identity service is obtained as follows:

- 263 1. A subject attempts to perform an action on a resource;
- 264 2. The resource enacts a policy enforcement point (PEP) that requires the subject to
265 authenticate with their identity service (i.e., `lgzIS` or `conIS`);
- 266 3. Upon authentication, a short term credential is released to the resource's PEP,
267 denoting a signed set of subject attributes (e.g., a SAML assertion [33]);
- 268 4. The PEP forwards the subject's issued credential to the authorisation service `as`,
269 which validates the contents of the credential to ensure attributes released have
270 been issued by a trusted identity provider;
- 271 5. If valid, the attributes are used to request access via the authorisation service `as`;
272 along with the resource, and action to be performed.
- 273 6. Lastly, the authorisation service `as` decides whether to grant access in accordance
274 to its authorisation policy.

275 3.2. Access Control Model

276 *LGZLogistics* employ an ABAC methodology to protect its resources. As
277 such, an instantiation of ABAC considers the subjects of *LGZLogistics* and
278 the subjects of the contractor organisation. Each set of subjects have a
279 permissible scope of access rights that can be assigned to them.

280 Figure 4 defines access in the form of a class diagram. There are five
281 'permisRole type' attributes [15] (specific to the PERMIS standalone au-
282 thorisation service) with corresponding values. Subjects are assigned these
283 attributes, which can then be used to invoke permissions.

284 In addition to the subject attribute assignments and attribute
285 permission assignments, *LGZLogistics* define a set of valid at-
286 tribute assignment rules (within its authorisation policy). Figure 5
287 specifies what attributes an identity provider is trusted to issue
288 on behalf of its employees. For example, *LGZLogistics* identity
289 provider `lgzIS` is trusted to assign attributes $\langle \text{permisRole}, \text{SysAdmin} \rangle$,
290 $\langle \text{permisRole}, \text{SysAnalyst} \rangle$, and $\langle \text{permisRole}, \text{Staff} \rangle$ to its employees. The
291 contractor organisation identity provider `conIS` can only assign attributes
292 $\langle \text{permisRole}, \text{ContractorSupervisor} \rangle$ and $\langle \text{permisRole}, \text{Contractor} \rangle$.

293 3.3. Subject Behaviour

294 This section identifies typical subject behaviour for the day to day oper-
295 ations of *LGZLogistics*, as well as a malicious behaviour scenario.

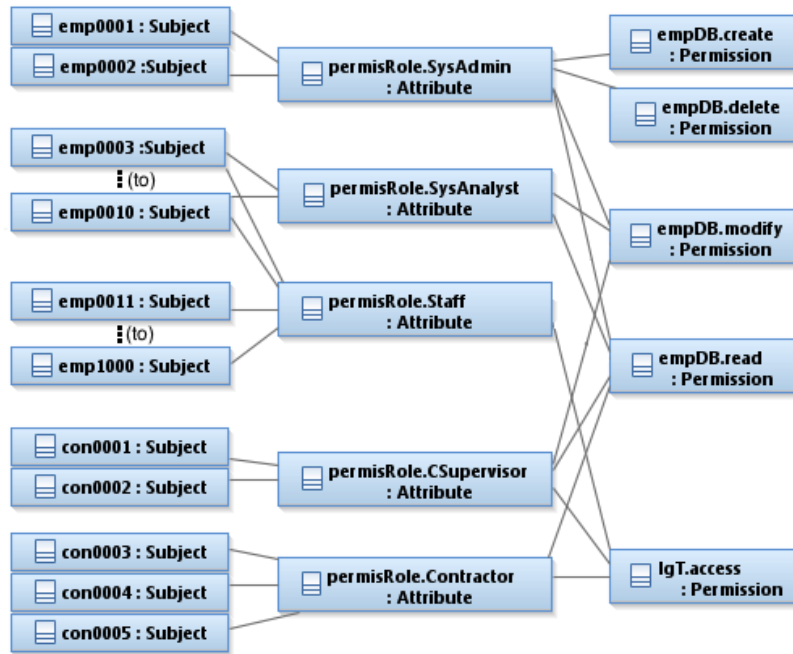


Figure 4: *LGZLogistics* subject attribute permission assignments

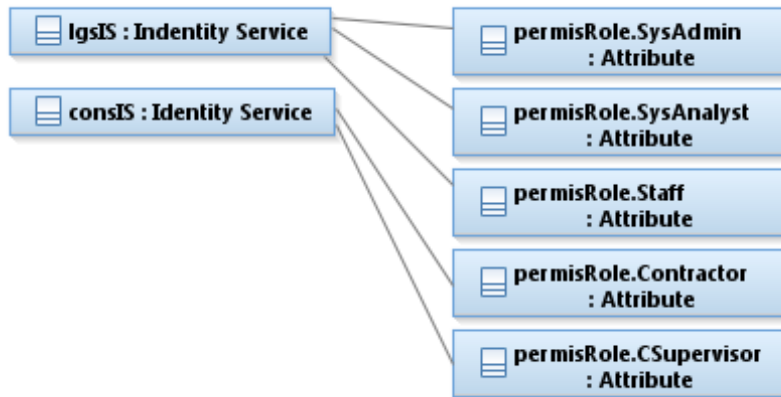


Figure 5: *LGZLogistics* valid attribute assignments

296 *3.3.1. Typical Behaviour*

297 The following describes a base line of subject behaviour, detailing the
 298 average usage of the authorisation infrastructure likely to occur in the day

299 to day operations of *LGZLogistics*:

- 300 • Each staff member requests ‘access’ to the lgT resource on average two times per
301 day;
- 302 • Contractors receive on average fifty calls per day, each call requiring one ‘read’
303 access to empDB;
- 304 • On average, one in five calls require access to ‘modify’ the empDB, which can only
305 be performed by a contractor supervisor, systems analyst, or system administrator;
- 306 • On average, each system analyst performs ten ‘read’ requests, and five ‘modify’
307 requests per day to the empDB;
- 308 • A system admin performs on average one ‘read’, ‘modify’, ‘delete’, and ‘create’
309 request per day to the empDB.

310 3.3.2. *Malicious Behaviour Scenario*

311 The logistics company is victim of an insider attack, largely as a result
312 of a catalyst event [32]. The catalyst event refers to a notification to several
313 key IT workers that they have been selected for job redundancy ¹.

314 A systems analyst that has been selected for redundancy is unhappy about
315 the decision, and attempts to damage the company in three ways. The first is
316 to attack the empDB resource by randomly corrupting employee records, in-
317 voking the permission ‘modify’ empDB. The second is an attempt to disrupt
318 access to the lgT resource, essentially flooding the resource by initiating nu-
319 merous authorised sessions. The final attempt is socially motivated, whereby
320 the analyst, who works closely with employees of the contractor organisation,
321 informs them that *LGZLogistics* is going to cancel their contract to cut costs.

322 A contractor supervisor, now fearing job redundancy, decides to steal
323 data from the empDB resource. The supervisor has links with the internet
324 underground [13], and is aware of anonymous buyers looking for data fit for
325 identify theft. By persuading his peers, three other contractors decide to
326 collaborate in stealing employee information from the empDB, to sell it to
327 the internet underground.

328 4. Specification of Malicious Changeload

329 In this section, we define the changeload related specifically to malicious
330 behaviour in the context of authorisation infrastructures. Essentially, it ap-

¹Instead of generalising an attack as being “harmful”, the labelling of an attack in the context of a case study is fundamental for specifying a meaningful malicious changeload.

331 plies Cámara et al.’s definitions of a changeload model [12] (which is spe-
332 cific to architectural-based self-adaptation) to authorisation infrastructures.
333 Cámara et al.’s changeload model was chosen in order to concretely define
334 the scope of change within an authorisation infrastructure.

335 Cámara et al. formulated their changeload model primarily to classify
336 system and environment changes that stimulate adaptation [12]. They have
337 defined *changeload* as a set of change scenarios that demonstrates changes,
338 which are: valid within a conventional operational profile, invalid thus stim-
339 ulating adaptation, or as the result of adaptation.

340 A *malicious changeload*, in the context of authorisation infrastructures,
341 drives stimulation of adaptation in response to the abuse of access control
342 (i.e., places a system into a non-conventional operational state). It is consid-
343 ered that both environment and system stimulation are capable in generat-
344 ing non-conventional operational states (and are often a by-product of each
345 other), whereby environment change leads to system change.

346 4.1. System and Environment Models

347 For the specification of system and the environment properties, we need
348 to define, respectively, the *system model* and *environment model*. These
349 models enable the specification of system properties that describe an autho-
350 risation infrastructure’s run-time parameters and workload, and environment
351 properties that characterise the operational conditions imposed on an autho-
352 risation infrastructure. The properties contained in both system and envi-
353 ronment models are dependent on a given deployment of an authorisation
354 infrastructure and its protected resources.

355 The *LGZLogistics* authorisation infrastructure is formally defined in
356 terms of an architecture model (Figure 6). For SAAF, on the other hand,
357 the *access control model* provides the relations between components of an
358 architectural model (i.e., how a subject of an identity service component
359 can access a resource component). Despite this, the use of an architectural
360 model is beneficial for defining properties of a system and the environment.
361 It enables the specification of system properties that describe an autho-
362 risation infrastructure’s run-time parameters and workload, and properties of
363 the environment that characterises the operational conditions imposed on
364 an authorisation infrastructure. These properties are said to be contained
365 within a system model and environment model, derived from the architecture
366 model.

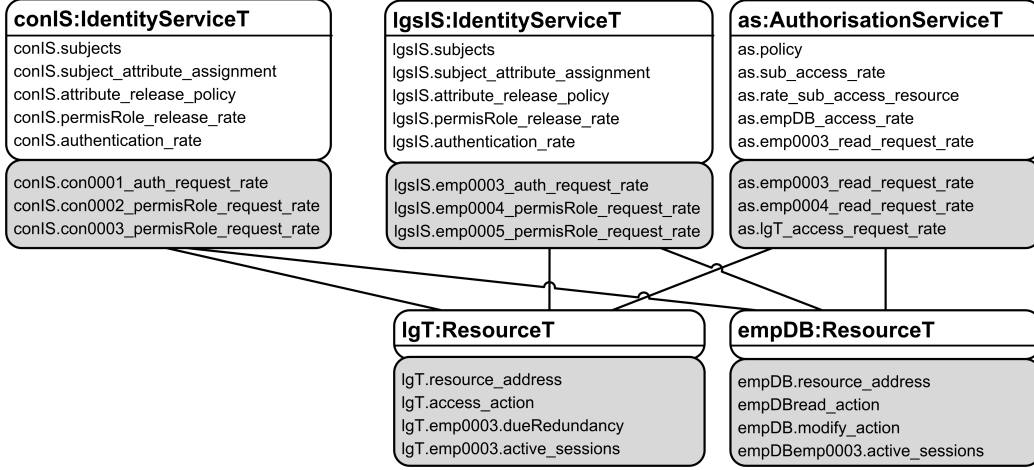


Figure 6: Example architecture model for an authorisation infrastructure.

367 **Example 1.** Figure 6 displays an architecture model of the LGZL-
 368 ogistics authorisation infrastructure, where the set of architectural types
 369 is $\mathcal{T} = \{IdentityServiceT, AuthorisationServiceT, ResourceT\}$. Examples of
 370 system properties include $\Gamma_{sys}(AuthorisationServiceT) = \{policy, sub_access_rate\}$.
 371 Examples of environment properties (displayed inside the grey boxes) include
 372 $\Gamma_{env}(AuthorisationServiceT) = \{sub_access_req_rate, lgT_access_request_rate\}$,
 373 and $\Gamma_{env}(ResourceT) = \{activeSessions, latency\}$.

374 4.2. System and Environment State

375 The *system state* captures the value of system properties and the execu-
 376 tion of services at a given moment in time. For example, the authentication
 377 of subjects, the release of subject attributes, the validation of subject at-
 378 tributes, and the authorisation of subject access.

379 **Example 2.** In this example, two system attributes are identified that denote
 380 execution of the authorisation infrastructure: (i) rate of attribute releases (of
 381 any kind) from the identity service *lgzIS*, and (ii) rate of successful read re-
 382 quests per interval to *empDB*. For typical execution of the identity service
 383 *lgzIS*, there is a constant throughput of one attribute release per minute. For
 384 typical execution of the authorisation service *as*, there is a constant through-
 385 put of three successful access decisions to *empDB*.

- 386 • $A = \langle lgzIS.sub_attr_release_rate, as.empDB_read_rate \rangle$

- 387 • $B = \langle \text{constant_function}, \text{constant_function} \rangle$
- 388 • $V_A = \langle 1, 3 \rangle$
- 389 • $V_B = \langle \theta_{lgzIS}(t) = 1/\text{min}, \theta_{as}(t) = 3/\text{min} \rangle$

390 The *environment state* captures operational conditions of external sys-
 391 tems and users that interact with the authorisation infrastructure. This in-
 392 cludes conditions, such as, the rate of access requests by subjects, or number
 393 of active sessions in resources. Building a perception of environment state is
 394 essential to identifying states that exhibit malicious behaviour (e.g., subjects
 395 exhibiting excessive deviation from normal activity).

396 **Example 3.** *In this example, three environment attributes are identified that*
 397 *denote operational conditions: (i) the number of active sessions in empDB,*
 398 *(ii) the rate of authentication requests made by all subjects against the iden-*
 399 *tity service lgzIS, and (iii) the rate of access requests to access the resource*
 400 *empDB to authorisation service as. The values associated with these opera-*
 401 *tional conditions are, respectively, five active sessions, a throughput of one*
 402 *authentication requests per minute, and a throughput of three access requests*
 403 *per minute.*

- 404 • $A = \langle \text{empDB.active_sessions}, \text{lgzIS.sub_authentications_req_rate},$
 405 $\text{as.empDB_read_req_rate} \rangle$
- 406 • $B = \langle \text{constant_function}, \text{constant_function}, \text{constant_function} \rangle$
- 407 • $V_A = \langle 5, 1, 3 \rangle$
- 408 • $V_B = \langle \theta_{\text{empDB}}(t) = 5, \theta_{\text{lgzIS}}(t) = 1/\text{min}, \theta_{\text{as}}(t) = 3/\text{min} \rangle$

409 4.3. Operational Profiles

410 An authorisation infrastructure can be in one of two types of states, a
 411 conventional operational state, or non-conventional operational state [12]. A
 412 *conventional operational state* refers to a state where there is no ongoing
 413 abuse of access rights.

414 **Example 4.** *A conventional operational profile is described as a set of states*
 415 *that does not contain any known patterns of abuse of access (i.e., violations).*

$$416 \text{COP} = \{s \in S \mid s \models \neg(\text{empDBViolation})\}$$

417 *A violation describes a predicate, that if true, denotes malicious behaviour*
 418 *within the environment state.*

419 A *non-conventional operational state* refers to a state where there is on-
420 going abuse of access rights.

421 **Example 5.** A *non-conventional operational profile* is described as a set of
422 states that contain one or more occurrences of malicious behaviours. In this
423 case, a violation (*empDBViolation*) denotes a specific violation in access to
424 the *empDB*.

$$425 \quad NCOP_{empDBViolation} = \{s \in S \mid s \models empDBViolation\}$$

426 The violation *empDBViolation* is focused on determining if any particular
427 subject is requesting access to the *empDB* resource in a rapid manner. A
428 subject that requests access to *empDB* at a rate (*subAccessReqRate_{empDB}*)
429 greater than a maximum prescribed rate (*maxSubAccessReqRate_{empDB}*) is
430 considered to be malicious.

$$431 \quad empDBViolation = subAccessReqRate_{empDB} > maxSubAccessReqRate_{empDB}$$

432 4.4. Change Types and Changes

433 Change types affect either identity services or authorisation services,
434 which are characterised as part of an authorisation infrastructure, or its envi-
435 ronment, consisting mainly of protected resources. *Change types* are defined
436 as a vector of ‘attributes’² that describe a change and the dynamics of a
437 change.

438 In application to authorisation infrastructures, a change type describes an
439 observable event within identity services, authorisation services, or protected
440 resources. Essentially, the observation of such change will have a consequence
441 on properties contained within the system and environment model.

442 **Example 6.** In the following, several examples of low level environment
443 change types are exemplified, depicting the process of a subject requesting
444 access to a resource. The instantiation of these change types will have a
445 consequence on one or more environment properties.

²Note that the domain of authorisation infrastructures refer to ‘attributes’ as a piece of information that expresses something about the subject or the current conditions within an accessed resource. This is not to be confused with attributes of a formal model of change (i.e., changeload).

446 (i) **Authentication request type** captures (within an identity service)
447 the identity service receiving a request for authentication of a user.

$$\text{auth_request_type} = \langle \text{identity_service}, \\ \langle \text{authRequest}(\text{username}, \text{password}) \rangle, \\ \langle \text{event} \rangle \rangle$$

448 (ii) **Attribute release request type** captures a request received by the
449 identity service made by a service provider for a subject's identity at-
450 tributes.

$$\text{attr_release_request_type} = \langle \text{identity_service}, \\ \langle \text{attrRequest}(\text{identity}, \\ \langle i\text{Attribute_type}_1, \dots, i\text{Attribute_type}_n \rangle, \text{target}) \rangle, \\ \langle \text{event} \rangle \rangle \\ \text{identity} = \langle \text{identity_type}, \text{identity_value} \rangle$$

451 It describes the request of a service provider (target) for a set of identity
452 attributes (*iAttribute_type*) that have been issued to a subject (*identity*).
453 The set of attributes requested can be a null set, therefore requesting all
454 releasable attribute types for the subject identity. Note that an identity is
455 referred to by a type of identifier and a value. For example, *identity_type*
456 may be an LDAP distinguished name.

457 (iii) **Credential validation request type** is the receipt of a credential
458 validation request within an authorisation service.

$$\text{cred_validation_request_type} = \\ \langle \text{auth_service}, \\ \langle \text{valRequest}(\text{identity}, \text{issuer}, \langle i\text{Condition}_1, \dots, i\text{Condition}_n \rangle, \\ \langle i\text{Attribute}_1, \dots, i\text{Attribute}_n \rangle) \rangle, \\ \langle \text{event} \rangle \rangle$$

459 It contains attributes issued by a given identity provider (*issuer*) for a
460 requesting subject, detailing a request to validate a subject's attributes.
461 A set of conditions specified by the issuer can also be contained, whereby
462 a condition refers to a type / value tuple, such as a single use declara-
463 tion, or validity time. A credential validation request can either push
464 the subject's known attributes, or (given a null set) require the autho-
465 risation service to pull the subject's known attributes from the subject's
466 identity provider. In the latter case, the authorisation service invokes
467 an attribute release request.

468 (iv) **Access request type** is the request, received by an authorisation ser-
469 vice, and made by a resource on behalf of a subject.

$$\text{access_request_type} = \langle \text{auth_service}, \\ \langle \text{accessRequest}(\langle i\text{Attribute}_1, \dots, i\text{Attribute}_n \rangle, \text{resource}, \\ \text{action}, \langle r\text{Attribute}_1, \dots, r\text{Attribute}_n \rangle, \text{identity}) \rangle, \\ \langle \text{event} \rangle \rangle$$
$$i\text{Attribute} = \langle i\text{Attribute_type}, i\text{Attribute_value} \rangle$$
$$r\text{Attribute} = \langle r\text{Attribute_type}, r\text{Attribute_value} \rangle$$

470 The request contains 1) the subject's identity attributes ($i\text{Attribute}$), 2)
471 the resource and action to be carried out by the subject, 3) a set of
472 resource environment attributes ($r\text{Attribute}$) provided by the resource
473 (e.g., $\langle \text{timeOfDay_type}, 11\text{am} \rangle$), and 4) the requesting subject's iden-
474 tity.

475 (v) **Resource action step type** models an action that has occurred within
476 any protected resource. The type is generic as resources are generally
477 unique to the organisation and their purpose, unlike with an authorisa-
478 tion service type that exists to fulfil access control requirements.

$$\text{resource_action_step_type} = \langle \text{resource}, \\ \langle r\text{Attribute} \rangle, \\ \langle \text{step_function} \rangle \rangle$$

479 The type identifies an attribute modification by means of a step function.
480 The attribute modified ($r\text{Attribute}$) is a tuple of type / value, and can
481 represent anything modelled within the resource type, be it generic or
482 specific. For example, this type could be instantiated to increase the total
483 amount of bandwidth consumed by a subject, within a given session.

484 **Example 7.** In the following, several examples of system change types are
485 described, conveying the system's response to a subject requesting access.

486 (i) **Authentication decision type** captures the consequence (within an
487 identity service) of an authentication request being responded to.

$$\text{auth_decision_type} = \langle \text{identity_service}, \\ \langle \text{authDecision}(\text{auth_request}) \rangle, \\ \langle \text{event} \rangle \rangle$$

488 (ii) **Attribute release type** is the consequence of an attribute release re-
489 quest, within an identity service.

$$\begin{aligned} \mathit{attr_release_type} &= \langle \mathit{identity_service}, \\ &\quad \langle \mathit{attrRelease}(\mathit{attr_release_request}), \\ &\quad \langle \mathit{event} \rangle \rangle \\ \mathit{attrRelease}(\mathit{attr_release_request}) &= \langle \mathit{issuer}, \mathit{identity}, \\ &\quad \langle \mathit{iCondition}_1, \dots, \mathit{iCondition}_n \rangle \\ &\quad \langle \mathit{iAttribute}_1, \dots, \mathit{iAttribute}_n \rangle \rangle \end{aligned}$$

490 It details the releasable identity attributes ($\mathit{iAttribute}$) as a tuple stating
 491 the type of identity attribute and its value. Identity attributes are re-
 492 leased along with the issuer of the attributes (i.e., an ID of the identity
 493 provider or individual whom assigned these attributes), the identity of
 494 the subject (i.e., a persistent ID), and a set of conditions. Conditions
 495 are a type value tuple, detailing the use of the released attributes. For
 496 example, a condition may assert the released attributes may only be used
 497 once, or can only be used in a given time frame.

498 (iii) **Credential validation type** is the consequence of a credential vali-
 499 dation request, within an authorisation service.

$$\begin{aligned} \mathit{cred_validation_type} &= \langle \mathit{auth_service}, \\ &\quad \langle \mathit{valCredentials}(\mathit{cred_validation_request}), \\ &\quad \langle \mathit{event} \rangle \rangle \\ \mathit{valCredentials}(\mathit{cred_validation_request}) &= \\ &\quad \langle \mathit{viAttribute}_1, \dots, \mathit{viAttribute}_n \rangle \end{aligned}$$

500 It returns valid attributes ($\mathit{viAttribute}_i$) for a subject if the provided
 501 $\mathit{iAttributes}$ conform to the authorisation service's credential validation
 502 policy. These are effectively the same as identity attributes, however,
 503 they are referred to as valid because an authorisation service has checked
 504 that the identity service is trusted to issue them.

505 (iv) **Access decision type** is the consequence of an access request, provid-
 506 ing a decision based on the attributes within an access request, and an
 507 authorisation service's access control policy.

$$\begin{aligned} \mathit{access_decision_type} &= \langle \mathit{auth_service}, \langle \mathit{accessDecision}(\mathit{access_request}), \\ &\quad \langle \mathit{event} \rangle \rangle \\ \mathit{accessDecision}(\mathit{access_request}) &= \mathit{decision} \end{aligned}$$

508 A *change* is an instantiation of a change type. Once enacted, the percep-
 509 tion of state (either system or environment) has changed.

510 **Example 8.** *In this example, the environment change types, provided in Ex-*
511 *ample 6, are instantiated into actual changes relevant to the LGZLogistics*
512 *case study.*

513 (i) **Authentication request** change defines the attributes received as part
514 of a request for authentication within identity provider *lgzIS*.

```
emp0003_auth_request = (auth_request_type, lgzIS,  
                        <authRequest(emp0003, password)>,  
                        <event>, 1373463234, 0)
```

515 (ii) **Attribute release request** change could be requested by a resource's
516 policy enforcement point at the time of authentication. However, it is
517 also used by authorisation services as part of a credential validation
518 request (depending on its configuration). The request states a set of
519 identity attribute types (i.e., attribute types that can exist with an iden-
520 tity, such as e-mail), and an identity. The identity, shown as a set
521 of numerical and alphabetical characters, is a privacy protected persis-
522 tent id (PID), however, equally could denote a non-privacy protected
523 identifier (e.g., an e-mail address).

```
emp0003_permisRole_request =  
    (attr_release_request_type, lgzIS,  
     <attrRequest(<pid, brx915810faa4910>,  
                 <permisRole>, empDB)>,  
     <event>, 1373463236, 0)
```

524 (iii) **Credential validation request**

```
emp0003_cred_validation_request =  
    (cred_validation_request_type, as,  
     <valRequest(<PID, brx915810faa4910>, lgzIS,  
               <<notOnOrBefore, 1373462240>,  
               <notOnOrAfter, 1373473240>)>,  
               <<permisRole, SysAnalyst>>>)>,  
     <event>, 1373463240, 0)
```

525 *This change portrays a credential validation request being received in*
526 *authorisation service as. A privacy protected identity is provided,*

527 along with the authenticating identity service (*lgzIS*), in order for
 528 the authorisation service *as* to validate the subject's released attribute
 529 $\langle \text{permisRole}, \text{SysAnalyst} \rangle$. The two conditions (*notOnOrBefore*,
 530 *notOnOrAfter*) state the validity of the released attribute.

531 (iv) **Access request** change captures the subject *emp0003*, with
 532 assigned identity PID: *bxu915810faa4910* and attribute
 533 $\langle \text{permisRole}, \text{SysAnalyst} \rangle$, requesting to execute 'read' action on
 534 the resource *empDB* (Payroll service). The change is observed as the
 535 receipt of request within the authorisation service *as*.

emp0003_empDB_request =
 (*access_request_type*, *as*,
 $\langle \text{accessRequest}(\langle \langle \text{permisRole}, \text{SysAnalyst} \rangle \rangle, \text{empDB},$
 read, $\langle \text{NULL} \rangle$,
 $\langle \text{pid}, \text{bxu915810faa4910} \rangle \rangle$),
 $\langle \text{event} \rangle, 1373463245, 0$)

536 (v) **Resource action step** is a change that increments the total bandwidth
 537 the subject *emp0003* has used within an active session, specifically, to
 538 the *empDB* resource. This change indicates a step change to the at-
 539 tribute *activeSessions[emp0003].bandwidth*, which contains the subject's
 540 current used bandwidth for their active session; the change increases
 541 200kb bandwidth to 800kb.

incr_emp0003_empDB_bandwidth =
 (*resource_action_step_type*, *empDB*,
 $\langle 200\text{kb} \rangle$,
 $\langle \theta_{\text{empDB}}(t) = \text{activeSessions}[\text{emp0003}].\text{bandwidth} + 600\text{kb} \rangle$,
 1373465245, 0)

542 **Example 9.** This example provides instantiations of the system change
 543 types identified in Example 7. An instantiation of a change type is defined
 544 as (*change_type*, *srcinst*, V_A , V_B , *time*, *duration*).

545 (i) **Authentication decision** change indicates the subject *emp0003* au-
 546 thenticating themselves against the identity service *lgzIS*, which is clas-
 547 sified by an event. The change is coupled with the attributes of the re-
 548 quest, in order to provide the decision. The decision generates a grant
 549 and the generation of a new session for the subject within the *lgzIS*
 550 identity service.

```

emp0003_auth = (auth_decision_type, is,
                <authDecision(emp0003_auth_request)>),
                <event>,
                1373463235, 0)
authDecision(emp0003_auth_request) = success

```

551 (ii) **Attribute release** change indicates a change observed at the *lgzIS*
552 identity service, where a resource *empDB* has requested the attribute
553 release of attribute type ‘*permisRole*’ of the subject *emp0003*. Identity
554 service *lgzIS* releases a tuple of attributes that match the request
555 from the resource for the required subject. In this case, it releases
556 $\langle \text{permisRole}, \text{SysAnalyst} \rangle$. The time indicates the time and date of
557 the attribute release, and as this is not associated to any session, the
558 duration is instant (0).

```

emp0003_permisRole_release =
    (attr_release_type, is,
     <attrRelease(emp0003_permisRole_request)>),
     <event>,
     1373463239, 0)
attrRelease(emp0003_permisRole_request) =
    <<permisRole, SysAnalyst>>

```

559 (iii) **Credential validation** change indicates a change observed within the
560 authorisation service *as*, being a credential validation. Credential val-
561 idation either validates the provided attributes in the request, or pulls
562 the subject’s attributes from the identity provider. In this example, the
563 authorisation service has validated the pushed attributes, asserting that
564 $\langle \text{permisRole}, \text{SysAnalyst} \rangle$ is valid.

```

emp0003_cred_validation =
    (cred_validation_type, as,
     <valCredentials(emp0003_cred_validation_request)>),
     <event>,
     1373463240, 0)
valCredentials(emp0003_cred_validation_request) =
    <<permisRole, SysAnalyst>>

```

565 (iv) **Access decision** change indicates the authorisation service *as* receiv-
566 ing a request and generating an authorisation decision based on the

567 *attributes of the request. The authorisation service has granted the re-*
 568 *quest. The change is instant and is only relevant for the specific request,*
 569 *therefore there is no duration.*

$$\begin{aligned} emp0003_empDB_grant &= (access_decision_type, as, \\ &\quad (accessDecision(emp0003_empDB_request)), \\ &\quad (event), \\ &\quad 1373463245, 0) \\ accessDecision(emp0003_empDB_request) &= permit \end{aligned}$$

570 4.5. Scenarios and Changeload

571 A *scenario* encompasses a set of changes over time, in light of a set of
 572 system goals, and a given state. It is used to formally describe malicious
 573 behaviour over time, such as a progression of violations. Key to a scenario
 574 is the definition of goals that should be fulfilled as the system undergoes
 575 change. In relation to detecting and mitigating malicious behaviour, a goal
 576 may refer to an error margin in detecting attacks, maximum response time
 577 to resolving attacks, and impact of attacks before required policy changes.

578 A *base scenario* defines a state that conforms to a system's conven-
 579 tional operational profile, i.e., a state where no known malicious behaviour
 580 is present. Such an assumption requires that only malicious behaviour in-
 581 tended to be stimulated against the base scenario can be evaluated, as it is
 582 not possible to rule out the existence of unknown malicious behaviour. The
 583 *LGZLogistics* case study has several valid base scenarios. For example, a
 584 base scenario could describe the typical workload during a normal business
 585 day within *LGZLogistics*. This includes a typical definition of criteria and
 586 assignment of access. Alternatively, it could represent the initial deployment
 587 of its authorisation infrastructure (i.e., no workload).

588 **Example 10.** *A base scenario for the LGZLogistics case study portrays the*
 589 *authorisation infrastructure, and its expected system and environment proper-*
 590 *ties, for a typical work day. For simplicity, only the system and environment*
 591 *properties relating to subject access are described. DailyAccess captures a*
 592 *base scenario of a typical system state relating to access decisions, and a*
 593 *typical environment state relating to access requests.*

$$594 \quad DailyAccess_{BaseScenario} = (SysAccess_{state}^t, EnvReq_{state}^t, G_f, \emptyset)$$

595 Each element of the base scenario tuple is expressed below. The system
 596 state combines properties that indicate the run-time parameters of services
 597 (e.g., authorisation policies), as well as system workload properties (e.g., rate
 598 of permitted access for a given subject). Both the system and environment
 599 states are defined in conformance to LGZLogistic's access control model (Sec-
 600 tion 3.2), and its definition of typical subject behaviour (Section 3.3).

$$\begin{aligned}
 SysAccess_{state}^t = & \\
 & \langle \langle as.policy, lgzIS.emp0003.attr, conIS.con0003.attr, \\
 & \quad as.emp0003.empDB.read, as.con0003.empDB.read \rangle, \\
 & \langle constant_function, constant_function, constant_function, \\
 & \quad constant_function, constant_function \rangle, \\
 & \langle AP_1, \{Staff, SysAnalyst\}, \{Contractor\}, 0.6, 1.25 \rangle, \\
 & \langle \theta_{as.policy}(t) = AP_1, \theta_{as}(emp0003.permisRole) = \{Staff, SysAnalyst\}, \\
 & \quad \theta_{as}(con0003.permisRole) = \{Contractor\}, \theta_{as}(t) = 0.6/min, \\
 & \quad \theta_{as}(t) = 1.25/min \rangle \\
 & \rangle
 \end{aligned}$$

601 $SysAccess_{state}^t$ defines the state of access control, including policies and
 602 attribute assignments. For example, subject emp0003 from identity service
 603 lgzIS, is assigned attributes $\langle permisRole, \{Staff, SysAnalyst\} \rangle$. AP_1
 604 denotes a PERMIS authorisation policy that implements the valid attribute as-
 605 signment rules in Figure 5, and attribute permission assignments in Figure 4.
 606 Note, the system state defined is not exhaustive, rather it focuses only on:
 607 system properties that define the current state of access; provides an exam-
 608 ple of attribute assignment to a subject from each identity provider; and an
 609 example rate of permitted access to the empDB resource.

$$\begin{aligned}
 EnvReqs_{state}^t = & \\
 & \langle \langle as.SysAdmin.empDB.Read, as.SysAdmin.empDB.Modify, \\
 & \quad as.SysAdmin.empDB.Create, as.SysAdmin.empDB.Delete, \\
 & \quad as.SysAnalyst.empDB.Read, as.SysAnalyst.empDB.Modify, \\
 & \quad as.ContractorSupervisor.empDB.Read, \\
 & \quad as.ContractorSupervisor.empDB.Modify, \\
 & \quad as.Contractor.empDB.Read, as.Staff.logisticsTool.Access \rangle, \\
 & \langle constant_function, constant_function, constant_function \\
 & \quad constant_function, constant_function, constant_function \\
 & \quad constant_function, constant_function, constant_function \\
 & \quad constant_function \rangle, \\
 & \langle 0.8, 0.4, 0.2, 0.2, 4.8, 1.6, 6.7, 6.7, 3.8, 20.0 \rangle, \\
 & \langle \theta_{as}(t) = 0.8/min, \theta_{as}(t) = 0.4/min, \theta_{as}(t) = 0.2/min, \theta_{as}(t) = 0.2/min, \\
 & \quad \theta_{as}(t) = 4.8/min, \theta_{as}(t) = 1.6/min, \theta_{as}(t) = 6.7/min, \\
 & \quad \theta_{as}(t) = 6.7/min, \theta_{as}(t) = 3.8/min, \theta_{as}(t) = 20.0/min \rangle \\
 & \rangle
 \end{aligned}$$

610 $EnvReq_{state}^t$ defines the state of the environment with regards to sub-
611 jects requesting access. The environment properties identified in the state
612 condition refer to collective behaviour per attribute per permission. For
613 example, for subjects requesting access to ‘read’ `empDB`, with attribute
614 $\langle permisRole, SysAdmin \rangle$, a rate of 0.8 requests per minute is observed. As
615 there are two subjects with this attribute (Figure 4), it is assumed that each
616 subject has an average rate of 0.4 requests per minute (i.e., one request every
617 150 seconds).

618 The fixed goals (G_f) define the conditions that must be maintained within
619 the authorisation infrastructure, regardless of change. Ultimately, a goal re-
620 quires a system to be brought out of a non-conventional operational state
621 (once identified). However, goals also focus on a wider scope of conditions
622 that attempt to ensure that only necessary adaptations are taken, once in a
623 non-conventional state. The following describes a set of goals relevant to the
624 LGZLogistics case study:

- 625 • Probability of 99% that all instances of known violation types are de-
626 tected;
- 627 • Probability of 90% that violations are mitigated through subject adap-
628 tation;
- 629 • Probability of 99% that all adaptations performed exhibit a lower cost
630 than current and unmitigated violations, to the organisation.

631 Probabilities cited are pseudo values that indicate LGZLogistics requirements
632 for mitigation. However, an accurate probability can only be achieved through
633 rigorous benchmarking of the scenario in an off-line environment [10]. In
634 any case, probabilities defined are specific to the deployment environment,
635 and configuration of the authorisation infrastructure.

636 Cámara et al. state that a *change scenario* represents a set of changes
637 applied to a base scenario [12]. As such, a change scenario instantiates a set
638 of changes within the authorisation infrastructure when it is in a particular
639 state. Through the application of change scenarios, it is expected to bring
640 the authorisation infrastructure into an non-conventional state, where the
641 authorisation infrastructure’s fixed goals can be evaluated.

642 **Example 11.** *The following sequence of changes describes subject emp0003*
643 *accessing the empDB resource.*

$c_1 = (\text{access_request_type}, \text{as}, \langle \text{request}(\langle \langle \text{permisRole}, \text{SysAnalyst} \rangle \rangle), \text{empDB},$
 $\text{read}, \langle \text{NULL} \rangle, \text{pid} = \text{bxu915810faa4910} \rangle \rangle, \langle \text{event} \rangle, 5, 0)$
 $c_2 = (\text{access_request_type}, \text{as}, \langle \text{request}(\langle \langle \text{permisRole}, \text{SysAnalyst} \rangle \rangle), \text{empDB},$
 $\text{read}, \langle \text{NULL} \rangle, \text{pid} = \text{bxu915810faa4910} \rangle \rangle, \langle \text{event} \rangle, 10, 0)$
 $c_3 = (\text{access_request_type}, \text{as}, \langle \text{request}(\langle \langle \text{permisRole}, \text{SysAnalyst} \rangle \rangle), \text{empDB},$
 $\text{read}, \langle \text{NULL} \rangle, \text{pid} = \text{bxu915810faa4910} \rangle \rangle, \langle \text{event} \rangle, 15, 0)$

644 c_1 describes a single access request for emp0003, identified
645 by privacy protected id (PID) bxu915810faa4910, using attribute
646 $\langle \text{permisRole}, \text{SysAnalyst} \rangle$, to access empDB. Thereafter, at 5 second
647 intervals, new changes are instantiated, whereby the request for access to
648 empDB is repeated. As a result, the subject affects a number of environment
649 properties associated to the system, namely, the subject's rate of access to
650 empDB.

651 The sequence of changes describes a rapid rate of access to the empDB
652 resource (labelled as $C_{\text{rapidAccess}}$). With the sequence of changes defined, it is
653 applied to the base scenario with the following notation:

$$\begin{aligned}
654 \quad & \text{ExpectedToRapidUsageChangeScenario} = \\
655 \quad & (\text{SysAccess}_{\text{state}}^t, \text{EnvReq}_{\text{state}}^t, G_f, C_{\text{rapidAccess}}) \\
656 \quad & C_{\text{rapidAccess}} = \{c_1, c_2, c_3\}
\end{aligned}$$

657 Cámara et al. formulated their changeload model primarily to classify
658 system and environment change that stimulates adaptation. As such, a ma-
659 licious changeload, in the context of authorisation infrastructures, is one that
660 drives stimulation of adaptation in response to the abuse of access control
661 (i.e., places a system into a non-conventional operational state). It is consid-
662 ered that both environment and system stimulation are capable in generat-
663 ing non-conventional operational states (and are often a by-product of each
664 other), whereby environment change leads to system change.

665 4.6. Violations

666 A set of violations are defined as the upper bounds of abnormal behaviour,
667 based on the normal behaviour described in the *LGZLogistics* case study. It is
668 assumed that historical data of subject behaviour (if present), coupled with
669 an expert approach, is used to define relevant violations. With reference
670 to the Self-Adaptive Authorisation Framework (SAAF), each violation is
671 defined as a trigger rule (with an associated cost).

672 The following violations detail patterns of access against *LGZLogistic's*
673 resources, regarding short term and long term rates of access invoking certain

674 permissions. For each violation a maximum rate of access is defined, whereby
 675 a short term rate refers to subject access within a minute interval, and a long
 676 term rate refers to subject access within a 10 minute interval (to simulate a
 677 scaled measure of prolonged change).

678 For example, violation `empDBShortRead` and violation `empDBLongRead`
 679 classifies malicious behaviour as any subject successfully requesting access to
 680 invoke the ‘read’ action on `empDB`, at a greater rate than a max allowable.
 681 A constraint is applied to the violation, whereby this violation only applies
 682 to subjects whom do not have the attribute $\langle permisRole, SysAdmin \rangle$.

```
empDBShortRead =
  (subAccessReqRateempDB.read >
   MaxAccessReqShortRateempDB.read) ∧
  (subAttribute <>  $\langle permisRole, SysAdmin \rangle$ )
```

```
empDBLongRead =
  (subAccessReqRateempDB.read >
   MaxAccessReqLongRateempDB.read) ∧
  (subAttribute <>  $\langle permisRole, SysAdmin \rangle$ )
```

683 For violations `empDBShortModify` and `empDBLongModify`, malicious be-
 684 haviour is classified in terms of any subject successfully requesting access to
 685 invoke the ‘modify’ action on `empDB`, at a greater rate than a max allow-
 686 able. As with the aforementioned violations, a constraint is applied meaning
 687 the violation is only applicable to subjects who do not have the attribute
 688 $\langle permisRole, SysAdmin \rangle$.

```
empDBShortModify =
  (subAccessReqRateempDB.modify >
   MaxAccessReqShortRateempDB.modify) ∧
  (subAttribute <>  $\langle permisRole, SysAdmin \rangle$ )
```

```
empDBLongModify =
  (subAccessReqRateempDB.modify >
   MaxAccessReqLongRateempDB.modify) ∧
  (subAttribute <>  $\langle permisRole, SysAdmin \rangle$ )
```

689 Violation `empDBShortDelete` classifies malicious behaviour in a subject
 690 rapidly gaining access to delete entries within the `empDB` resource.

```
empDBShortDelete =
  (subAccessReqRateempDB.delete >
   MaxAccessReqShortRateempDB.delete)
```

691 Violation `lgTShortAccess` classifies malicious behaviour of subjects
692 rapidly accessing the `lgT` (logistic tool) resource.

$$\text{lgTShortAccess} = (\text{subAccessReqRate}_{\text{lgT}} > \text{MaxAccessReqFastRate}_{\text{lgT}})$$

693 Violation `empDBTransaction` is slightly different, whereby it classifies a
694 transaction of non-conventional change. This type of violation denotes a
695 pattern whereby a rate of transactional requests are compared against a
696 maximum rate. The violation requires an environment property that mea-
697 sures the rate of access requests, by a subject, in performing a read action
698 succeeded by a modify action against `empDB`. Basically, it aims to identify
699 subjects who rapidly read and write to the `empDB` resource.

$$\text{empDBTransaction} = (\text{subAccessReqRate}_{\text{empDB.readModifyTransaction}} > \text{MaxAccessReqLongRate}_{\text{empDB.readModifyTransaction}})$$

700 A final violation, albeit by contrast does not capture subject activity
701 directly, is `dueRedundancy`. This violation is a consequence of a change
702 made within the `empDB` resource, indicating that a subject has been marked
703 for job redundancy. A subject facing the prospect of redundancy is seen
704 as a potential risk, and as such, a violation is used to increase the impact a
705 subject has on an organisation. This is viewed as a motivator for adaptation,
706 as when combined with previously identified violations, the subject's activity
707 may now warrant adaptation.

$$\text{dueRedundancy} = (\text{subDueRedundancy} == \text{true})$$

708 4.7. Identifying Change Types and Change

709 To stimulate violations within the context of the *LGZLogistics* case study,
710 it is necessary to identify properties of interest and the change types that
711 will impact such properties. For this specific case study, only environment
712 properties are considered. These are properties that concern subject activity
713 that cannot be directly controlled (e.g., a subject's rate of access requests).

714 *4.7.1. Environment Properties*

715 For each violation, and for each subject, there exists a set of environment
716 properties that measure the extent of change in the environment. Many envi-
717 ronment properties represent composite properties of subject-related changes
718 over time. In reference to SAAF, these properties are dynamically created
719 as mutable attributes within SAAF’s behaviour model, and updated through
720 the observation of environment change via probes.

721 For example, `empDBShortRead` asserts that if a subject’s access rate
722 in requesting a read on `empDB` (who is not a *SysAdmin*) goes be-
723 yond a maximum number of requests within a minute interval, a vi-
724 olation has occurred. To measure against this violation, an environ-
725 ment property of `as.subject.AccessReqRateempDB.read` is required (e.g.,
726 `as.emp0003.AccessReqRateempDB.read`).

727 *4.7.2. Change Types and Changes*

728 Once environment properties are identified it is necessary to select rele-
729 vant change types (and changes) that result in a non-conventional operational
730 state. For example, a non-conventional operational profile that contains
731 the violation `empDBShortRead` is realised through a succession of changes,
732 whereby a single subject successfully requests access to ‘read’ `empDB`. The
733 violation occurs when a subject, e.g., `emp0003`, has performed a number
734 of **Access request** change type, and is permitted by an **Access decision**
735 change type.

736 The **Access request** change type is the result of a number of sequential
737 changes, such as the subject first authenticating with their identity provider,
738 requesting a release of attributes as credentials, and validation of attributes.
739 In this instance, these changes need to be realised before a subject performs
740 an **Access request** change type.

741 All but one violation described for the *LGZLogistics* case study is trig-
742 gered by an **Access request** change type. The violation `dueRedundancy`
743 is triggered by a **Resource action step** change, whereby an environment
744 property for a given subject indicates a subject is due for job redundancy
745 (e.g., `empDB.emp0003.isSetForRedundancy`).

746 *4.8. Malicious Changeload*

747 Using the *LGZLogistics* malicious behaviour scenario, the following set
748 of change scenarios are defined. Together they represent the malicious
749 changeload for the case study. There are seven change scenarios defined

750 within the malicious changeload, representative of the case study’s malicious
 751 behaviour scenario, and each change scenario is applicable to the base sce-
 752 nario.

753 The first change scenario (*setSubjectRedundanciesChangeScenario*) consid-
 754 ers a set of resource changes relevant to **empDB**, where changes identify four
 755 system analysts are to be made redundant (**dueRedundancy**).

$$\begin{aligned}
 & \textit{setSubjectRedundanciesChangeScenario} = \\
 & \quad (\textit{SysAccess}_{state}^t, \textit{EnvReq}_{state}^t, G_f, C_{\textit{setRedundancies}}) \\
 C_{\textit{setRedundancies}} = \{ & \\
 & \mathbf{c}_1 = (\textit{resource_action_step_type}, \textit{empDB}, \\
 & \quad \langle \textit{emp0003.Redundancy} \rangle, \\
 & \quad \langle \textit{emp0003.Redundancy} = \textit{true} \rangle, 0, 0) \\
 & \mathbf{c}_2 = (\textit{resource_action_step_type}, \textit{empDB}, \\
 & \quad \langle \textit{emp0004.Redundancy} \rangle, \\
 & \quad \langle \textit{emp0004.Redundancy} = \textit{true} \rangle, 0, 0) \\
 & \mathbf{c}_3 = (\textit{resource_action_step_type}, \textit{empDB}, \\
 & \quad \langle \textit{emp0005.Redundancy} \rangle, \\
 & \quad \langle \textit{emp0005.Redundancy} = \textit{true} \rangle, 0, 0) \\
 & \mathbf{c}_4 = (\textit{resource_action_step_type}, \textit{empDB}, \\
 & \quad \langle \textit{emp0006.Redundancy} \rangle, \\
 & \quad \langle \textit{emp0006.Redundancy} = \textit{true} \rangle, 0, 0) \}
 \end{aligned}$$

756 The second scenario (*emp0003ReadModifyChangeScenario*) describes a
 757 malicious change scenario resulting in violations **empDBLongReadModify**,
 758 **empDBLongRead**, and **empDBLongModify**, whereby subject emp0003 persis-
 759 tently reads and modifies records in the **empDB** resource, every four seconds.
 760 An arbitrary function δ is defined in order to calculate the time at which a
 761 change is executed within the change scenario (for each scenario, we assume
 762 a different function δ). For the following change scenario, δ is defined as
 763 $\delta(n) = \frac{1}{2}(4n - (-1)^n + 1)$, where n refers to the n th change in the change
 764 scenario.

$$\begin{aligned}
 & \textit{emp0003ReadModifyChangeScenario} = \\
 & \quad (\textit{SysAccess}_{state}^t, \textit{EnvReq}_{state}^t, G_f, C_{\textit{maliciousTransactions}}) \\
 C_{\textit{maliciousTransactions}} = \{ & \\
 & \mathbf{c}_1 = (\textit{access_request_type}, \textit{as}, \\
 & \quad \langle \textit{request}(\langle \langle \textit{permisRole}, \textit{SysAnalyst} \rangle \rangle), \textit{empDB}, \\
 & \quad \textit{read}, \langle \textit{NULL} \rangle, \textit{pid} = \textit{emp0003} \rangle, \langle \textit{event} \rangle, 3, 0) \\
 & \mathbf{c}_2 = (\textit{access_request_type}, \textit{as}, \\
 & \quad \langle \textit{request}(\langle \langle \textit{permisRole}, \textit{SysAnalyst} \rangle \rangle), \textit{empDB}, \\
 & \quad \textit{modify}, \langle \textit{NULL} \rangle, \textit{pid} = \textit{emp0003} \rangle, \langle \textit{event} \rangle, 4, 0) \\
 & \dots
 \end{aligned}$$

$$\begin{aligned} \mathbf{c}_n &= (\text{access_request_type}, \text{as}, \\ &\quad \langle \text{request}(\langle \langle \text{permisRole}, \text{SysAnalyst} \rangle \rangle), \text{empDB}, \\ &\quad \text{read}, \langle \text{NULL} \rangle, \text{pid} = \text{emp0003} \rangle, \langle \text{event} \rangle, \delta(n), 0) \\ \mathbf{c}_{n+1} &= (\text{access_request_type}, \text{as}, \\ &\quad \langle \text{request}(\langle \langle \text{permisRole}, \text{SysAnalyst} \rangle \rangle), \text{empDB}, \\ &\quad \text{modify}, \langle \text{NULL} \rangle, \text{pid} = \text{emp0003} \rangle, \\ &\quad \langle \text{event} \rangle, \delta(n + 1), 0) \end{aligned}$$

765 The third scenario (*con0002FastReadChangeScenario*) describes a malicious
766 change scenario by subject con0002, resulting in violations **empDBShortRead**
767 and **empDBLongRead**. In this scenario, a contractor supervisor persistently
768 accesses the **empDB** resource at a rate of 2 seconds, in order to obtain em-
769 ployee data. The scenario begins 150 seconds relative to the start of mali-
770 cious changeload. The function δ denotes the progression in time (seconds)
771 between changes, defined as $\delta(n) = 2n$.

$$\begin{aligned} \text{con0002FastReadChangeScenario} &= \\ &(\text{SysAccess}_{state}^t, \text{EnvReq}_{state}^t, G_f, C_{\text{con0002FastRead}}) \\ C_{\text{con0003FastRead}} &= \{ \\ \mathbf{c}_1 &= (\text{access_request_type}, \text{as}, \\ &\quad \langle \text{request}(\langle \langle \text{permisRole}, \text{ContractorSupervisor} \rangle \rangle), \\ &\quad \text{empDB}, \text{read}, \langle \text{NULL} \rangle, \text{pid} = \text{emp0003} \rangle, \\ &\quad \langle \text{event} \rangle, 150, 0) \\ &\dots \\ \mathbf{c}_n &= (\text{access_request_type}, \text{as}, \\ &\quad \langle \text{request}(\langle \langle \text{permisRole}, \text{ContractorSupervisor} \rangle \rangle), \\ &\quad \text{empDB}, \text{modify}, \langle \text{NULL} \rangle, \text{pid} = \text{emp0003} \rangle, \\ &\quad \langle \text{event} \rangle, \delta(n), 0) \end{aligned}$$

772 For contractors con0003, con0004, and con0005, similar change scenarios
773 exist based on *con0002FastReadChangeScenario*. However, the scenarios are
774 introduced in stages of 30 second intervals (i.e., con0003 begins at 3 minutes
775 from the start of the malicious changeload, con0004 begins at 3.5minutes,
776 etc.). The rate of changes is defined as $\delta(n) = 2.5n$, where subjects utilise
777 their $\langle \text{permisRole}, \text{Contractor} \rangle$ attribute, in accessing **empDB**.

778 Lastly, *emp0003FastAccessChangeScenario* describes a further malicious
779 change scenario by emp0003, resulting in violation **lgTShortAccess**. In this
780 scenario, emp0003 persistently requests access every 370ms and gains access
781 to **lgT** resource (logisticsTool resource), to disrupt the performance of the
782 resource. The scenario begins 900 seconds relative to the start of malicious
783 changeload. δ denotes a function whereby progression in time (seconds) is
784 defined as $\delta(n) = 0.37n$.


```

emp0003FastAccessChangeScenario =
  (SysAccesststate, EnvReqststate, Gf, Cemp0003FastAccess)
Cemp0003FastAccess = {
  c1 = (access_request_type, as,
        <request(<<permisRole, Staff>>), lgT,
        read, <NULL>, pid = emp0003),
        <event>, 900, 0)
    . . .
  cn = (access_request_type, as,
        <request(<<permisRole, SysAnalyst>>), lgT,
        read, <NULL>, pid = emp0003),
        <event>, δ(n), 0)
}

```

785 This malicious changeload (consisting of the seven change scenarios) con-
786 cisely describes the *LGZLogistics* malicious behaviour scenario. It is the
787 intention that the changeload can be repeated under various operational
788 conditions, and also used to compare future approaches to self-adaptive au-
789 thorisation. As such, it can be exploited to execute simulation of changes
790 within an authorisation infrastructure, in order to evaluate the impact of
791 violations, and trigger self-adaptive responses. However, one limitation is
792 that no parser currently exists to execute a defined changeload. Therefore, a
793 changeload can only be viewed as a model of change, which must be manually
794 transformed into an executable script (e.g., JMeter simulation scripts).

795 5. Experiments

796 The *LGZLogistics* case study is simulated within a live self-adaptive au-
797 thorisation infrastructure. This self-adaptive authorisation infrastructure is
798 instantiated across four individual machines. Two machines running **Debian**
799 **Linux** (512MB of memory) are deployed hosting an LDAP directory and an
800 installation of SimpleSAMLphp [37]. These are configured to operate as
801 the **lgzIS** and **conIS** identity services, respectively. A single machine running
802 **Ubuntu Linux** (2048MB of memory) is deployed hosting an installation of
803 the PERMIS standalone service, instantiating authorisation service **as**, and
804 a prototype of the SAAF controller. Lastly, a single ‘client’ machine run-
805 ning **Windows** (2048MB) is deployed to simulate activity between subjects
806 accessing a resource, and communicating with services of the authorisation
807 infrastructure.

808 The rest of this section details a brief overview of the deployment of
809 the prototype of the SAAF controller, a description of how the malicious

810 changeload is simulated within the environment, the execution of experi-
 811 ments, and lastly, a summary of results.

812 5.1. Deploying the SAAF Controller

813 For the *LGZLogistics* case study, SAAF is deployed as a federated autho-
 814 risation infrastructure in order to facilitate the adaptation process.

815 Figure 7 portrays *LGZLogistic's* federated authorisation infrastructure,
 816 based on the architectural model described in Section 2.1. Here, the in-
 817 frastructure is distributed across multiple management domains (identity
 818 provider and service provider domains). *LGZLogistics* operates a service
 819 provider domain (to handle authorisation and provision access to resources),
 820 and their own identity provider domain (to handle identity management of
 821 their own employees). In addition, the *contractor* organisation is said to op-
 822 erate their own identity provider domain (to handle identity management of
 823 their own employees).

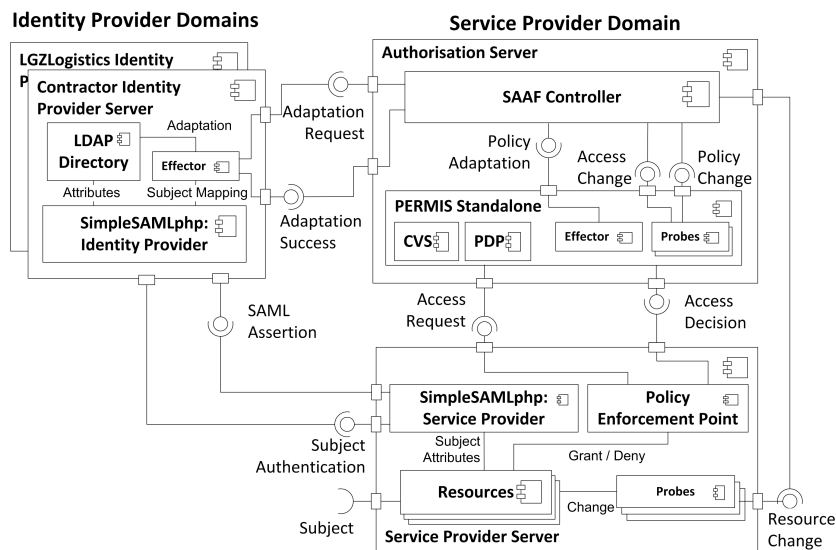


Figure 7: Self-adaptive federated authorisation infrastructure

824 SimpleSAMLphp [37] is used as the enabling technology to facilitate
 825 communication between these management domains. It provides a layer
 826 of control over ‘what’ information can be released or requested (in regards
 827 to subjects), and how subjects can be authenticated. Deployments of Sim-
 828 pleSAMLphp are capable of exchanging information via signed or unsigned

829 SAML assertions [33], such as messages containing a set of subject attributes
830 and the subject's unique identifier.

831 A prototype of the SAAF controller is deployed within *LGZLogistics* ser-
832 vice provider domain, whereby it is expected to manage authorisation as-
833 sets across both management domains. However, self-adaptation over multi-
834 ple management domains is a challenging and non-trivial problem. Identity
835 providers often do not release uniquely identifiable personal information to
836 service providers, and use transient (TID) or persistent (PID) IDs to allow
837 service providers to identify subjects. In addition, identity providers may not
838 be as forthcoming to accepting adaptations from an organisation outside of
839 their management domain, meaning the SAAF controller can only 'request'
840 adaptation.

841 A solution to enabling adaptation across multiple management domains
842 is the deployment of an effector managed by the identity provider domain [5].
843 Here an effector can map a service provider's view of a subject (i.e., from a
844 subject TID / PID to subject LDAP entry), and govern which adaptations
845 to perform. Instantiations of this effector are deployed on each of the iden-
846 tity services (*subject identity adaptation*), as well as an effector capable of
847 deploying and activating policies within the PERMIS authorisation service
848 (*policy adaptation*).

849 A resource probe is deployed on the *empDB* resource to observe changes
850 to the state of an employee's job redundancy property (*resource change*). In
851 addition, a probe is deployed on the PERMIS authorisation service to detect
852 *access change* and *policy change*. A probe is not deployed on the contractor's
853 identity service (*conIS*) simulating a limitation in federated authorisation in-
854 frastructures, where third party organisations may prevent immediate access
855 to their subject's attributes (*subject change*). This limits the SAAF proto-
856 type's view of the state of access, whereby the SAAF prototype must infer
857 its perception of subjects from the observation of access requests (via the
858 authorisation service *as*).

859 The SAAF controller is configured to detect and mitigate the set of viola-
860 tions described in Section 4.6. Here, a solution policy exists containing a set
861 of solutions applicable to mitigating instances of these violations. Each solu-
862 tion contains a weighting of cost to the deploying organisation (e.g., the cost
863 in removing subject access, or removing the trust in an identity provider). A
864 minimum subject impact weighting, on a scale of 0 to 1, is also defined, which
865 is used to constrain a subset of solutions relevant to resolving differing scales
866 of malicious subject behaviour. These weightings are used as part of solution

867 analysis and solution selection. The following solutions are configured for
868 this deployment:

- 869 • `S0 - noAdaptation` default solution for when all other solutions cause greater im-
870 pact over an observed behaviour;
- 871 • `S1 - removeSubjectAttribute` removal of an individual abused attribute from a
872 subject (i.e., the cause of a violation);
- 873 • `S2 - removeAllSubjectAttributes` removal of all attributes from a subject, typ-
874 ical for when subjects are persistently abusing access;
- 875 • `S3 - removeAttributeAssignment` removal of trust in an identity provider in is-
876 suing valid attributes (policy change);
- 877 • `S4 - removeAllAttributeAssignments` removal of all trust in an identity provider
878 in issuing valid attributes (policy change);
- 879 • `S5 - deactivatePolicy` removal of all access to all of *LGZLogistic's* resources.

880 A limitation in this deployment is the inability to use the integrated
881 `rbacDSML` tool³, preventing solution verification from taking place. This is
882 due to the fact that our deployment operates within a federated environment
883 that conforms to ABAC, which `rbacDSML` is unable to accommodate for. It
884 was decided that evaluating SAAF within a federated environment provided
885 greater contributions as opposed to enhancing `rbacDSML` to operate within
886 a federated ABAC environment [6]. As a consequence, we assume all adap-
887 tation solutions result in acceptable implementations of the access control
888 model.

889 5.2. Executing *LGZLogistics* Changeload

890 The execution of the *LGZLogistics* malicious changeload (Section 4.8) is
891 achieved through enacting environment change via a number of protocols:

- 892 • LDAP binds [25] - for authenticating subjects within identity providers;
- 893 • SAML assertions [33] - for requesting and deliverance of released at-
894 tributes as signed credentials (to and from identity provider services);

³`rbacDSML` allows the verification of RBAC access control models, enabling organisa-
tions to manage their access control models with greater accuracy, efficiency, and assur-
ance [29].

895 • SOAP messages [9] - for credential validation requests, credential val-
 896 idation responses, access requests, and access decisions (to and from
 897 protected resources and authorisation services).

898 An installation of the JMeter testing tool [2] (deployed on the Windows
 899 ‘client’ machine) automates each of the change types applicable to an autho-
 900 risation infrastructure, using the aforementioned protocols. Here, subjects
 901 are simulated in authenticating, requesting, and obtaining access to protected
 902 resources.

903 Using the experimentation profile proposed by Cámara et al. [12], the
 904 malicious changeload is executed across multiple runs as part of four experi-
 905 ments. The two experiments are designed to evaluate the SAAF prototype.
 906 Exp1 and Exp2 evaluate the prototype in mitigating malicious changeload
 907 under normal and high loads, respectively.

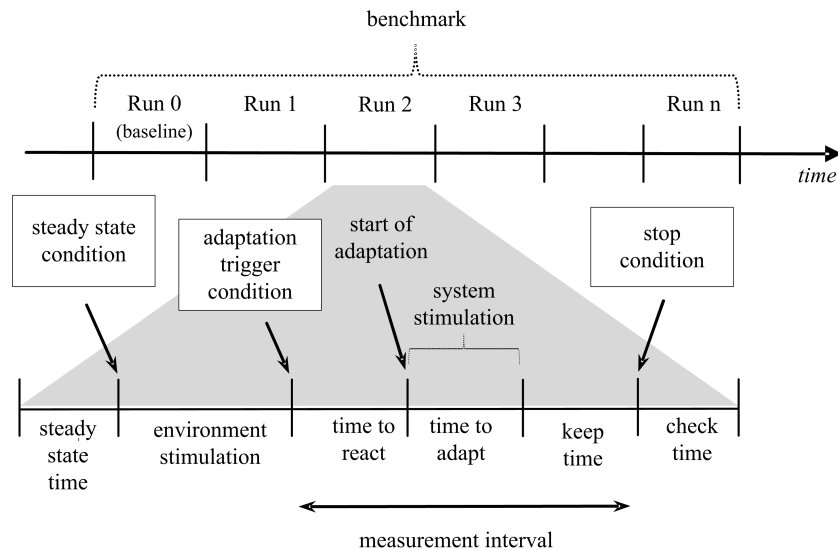


Figure 8: Executing changeload experimentation profile [12]

908 Each experiment is executed six times (referred to as ‘runs’), where each
 909 run follows the set of stages stated in Figure 8. The stages of a run are:

- 910 1. steady state time - the realisation of a base scenario for the *LGZLogis-*
 911 *tics* case study that portrays the authorisation infrastructure and its
 912 expected system properties, and environment properties, for a typical

- 913 work day. Steady state time is maintained for a period of 30 minutes in
914 order to ensure the controller and authorisation infrastructure is eval-
915 uated in a warmed up state. During this time, the baseline scenario is
916 simulated within the authorisation infrastructure;
- 917 2. environment stimulation - the execution of the malicious changeload.
918 From this period on there is a set of staggered violations in which sev-
919 eral periods of ‘time to react’ and ‘time to adapt’ overlap environment
920 stimulation. This is necessary in order to evaluate the prototype’s abil-
921 ity to detect and mitigate multiple attacks that have been conducted
922 collaboratively;
 - 923 3. time to react - the detection of malicious behaviour and decision to act;
 - 924 4. time to adapt - time it takes to perform adaptations;
 - 925 5. keep time - time needed to observe system recovery post adaptation.
926 In this post adaptation the baseline scenario resumes and no further
927 adaptation takes place.
 - 928 6. check time - the post analysis of each run. It remains the same for each
929 run within an experiment.

930 At the end of each run the system and environment states are reset before
931 performing the next run. For consistency, each run (and experiment) observes
932 the same steady state, malicious changeload, and keep / check time.

933 *5.3. Experiments Execution*

934 Experiments **Exp1** and **Exp2** demonstrate adaptation under increasing
935 loads on the controller (in terms of processing environment change).

936 *5.3.1. Baseline execution*

937 Baseline runs identify the impact of the malicious changeload whereby no
938 adaptation takes place. During these runs, the prototype controller is active,
939 yet limited to only detecting the number and types of violations that have
940 occurred.

941 Figure 9 (i) and (ii) describe the rate of access of key subjects within the
942 *LGZLogistics* authorisation infrastructure (taken at minute intervals). Note
943 that ‘all.Staff’ indicates an aggregate rate for all subjects with attribute
944 $\langle permisRole, staff \rangle$, whereas, all others represent the access requests of an

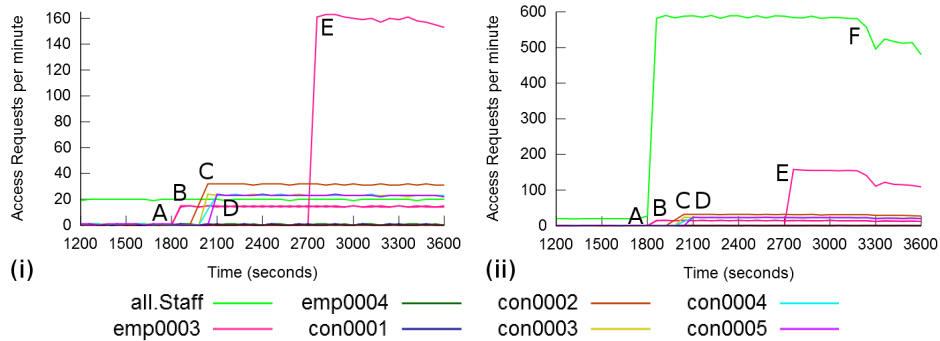


Figure 9: (i) Baseline Exp1 normal load, (ii) baseline Exp2 high load

945 individual. Figure 9 (i) depicts execution of malicious changeload under
 946 normal load, simulated as the continuation of the base scenario through-
 947 out environment stimulation. Figure 9 (ii) depicts execution of malicious
 948 changeload under high load, simulated as an increase to the base scenario’s
 949 ‘staff’ rate of access, from 20req/min to 600req/min.

950 The normal load baseline (i) is representative of a baseline run for Exp1
 951 and Exp2 since the experiments undergo the same steady state and malicious
 952 changeload scenarios for their corresponding runs.

953 Comparing the baseline runs portrayed minimal difference in violations
 954 observed. Point A indicates the start of the malicious changeload (1800
 955 seconds into the run), where the *setSubjectRedundancies* change scenario is
 956 executed, sending the controller several resource change events. It also indi-
 957 cates the start of *emp0003ReadModify* change scenario, at point B, where a
 958 system analyst begins to persistently read and modify records in empDB at
 959 a rate of 15req/min. At point C a contractor supervisor (*con0002FastRead*)
 960 begins to persistently read the empDB resource at a rate of 33req/min. This
 961 is followed by D, where three contractors also begin malicious behaviour,
 962 exhibiting a slightly lower request rate of 24req/min. Lastly, at point E,
 963 *emp0003FastAccess* change scenario is stimulated, representing a system an-
 964 alyst attempting to disrupt the performance of resource lgT, accessing at a
 965 rate of 160req/min.

966 The only exception between the two baselines is indicated at point F
 967 (high load baseline). As a result of the client machine being pushed to its
 968 limits (overloaded by *emp0003FastAccess*), a slowdown in load occurred after
 969 3000 seconds into the run. Whilst this presented an anomaly to the baseline,

970 adaptation runs were not impacted, as shown in Figure 10.

971 Regarding the detection of malicious behaviour, the controller detected
 972 275 violations in normal load (i), and 260 violations in high load (ii). These
 973 were confined to six types of violations: `dueRedundancy`, `empDBTransaction`,
 974 `empDBShortRead`, `empLongRead`, `lgTShortAccess`, `empLongModify`. The
 975 high load baseline had fewer detections due to the slowdown of client re-
 976 quests at 3000 seconds into the run.

977 5.3.2. Adaptation execution

978 Experiments `Exp1` and `Exp2` undergo the same malicious changeload as
 979 the baseline execution, albeit against a normal and high load, respectively.
 980 The experiments and results to be discussed in the following take as a basis
 981 Figure 10 and Table 1.




Figure 10: (i) `Exp1` normal load, (ii) `Exp2` high load

982 Both experiments resulted in the consistent identification and selection of
 983 solutions to violations, where 14 attack steps were identified and responded
 984 to. Table 1 details these attack steps, where each *step* describes:

- 985 • A malicious *subject* who has carried out a violation;
- 986 • The calculated *impact* of the *subject* to the organisation;
- 987 • The *violation* observed;
- 988 • A set of *identified solutions* in which to mitigate the violation;
- 989 • The *selected solution* used to mitigate the violation;

Step	Subject	Impact	Violation	Identified Solutions	Selected Solutions	R_{Time} (Avg, Std)	A_{Time} (Avg, Std)	Result
1	emp03	0.07	dueRedundancy	S0	S0	4.6, 0.5	N/A	N/A
2	emp04	0.07	dueRedundancy	S0	S0	2, 0	N/A	N/A
3	emp05	0.07	dueRedundancy	S0	S0	2, 0	N/A	N/A
4	emp06	0.07	dueRedundancy	S0	S0	1.8, 0.4	N/A	N/A
5	emp03	0.4	empDBTransaction	S1	S1	291.6, 59.6	182.2, 48.6	1
6	con02	0.07	empShortRead	S0	S0	1.4, 0.5	N/A	N/A
7	con02	0.27	empShortRead	S1	S1	186.6, 77.5	153.4, 22	1
8	con03	0.07	empShortRead	S0	S0	2.4, 1.1	N/A	N/A
9	con04	0.07	empShortRead	S0	S0	1.6, 0.5	N/A	N/A
10	con03	0.27	empShortRead	S1	S1	94.8, 33.8	165.2, 63.2	1
11	con05	0.07	empShortRead	S0	S0	4.6, 2.8	N/A	N/A
12	con04	0.27	empShortRead	S1	S1	139.8, 23.1	146.2, 38.9	1
13	con05	0.27	empShortRead	S1	S1	70.2, 9.7	120.4, 27.7	1
14	emp03	0.8	lgTShortAccess	S2,S3,S4,S5	S2	297.8, 35.8	189.4, 63.8	1

Table 1: Exp1: Adaptation with normal load (time in milliseconds)

- 990 • The response time (R_{Time}) in which to select the solution;
- 991 • The time to carry out an adaptation A_{Time} ;
- 992 • The result as to whether a solution was successful (1), failed (0), or not
993 applicable (i.e., no adaptation performed).

994 Reviewing Table 1, steps 1 to 4 identify a resource change event (at
995 point A), which triggered the violation `dueRedundancy`. For each subject, an
996 impact level was calculated based on past behaviour observed. The controller
997 calculated a low impact for these subjects (0.07), as none of these subjects
998 have any previous violations. However, the result of this means that the
999 controller is less tolerable to future violations.

1000 Step 5 portrays the controller’s first adaptation in response to subject
1001 `emp0003` triggering a second violation (`empDBTransaction` at point B). As a
1002 result, the subject’s impact level was recalculated from 0.07 to 0.4. Solution
1003 S1 was identified (as being within scope of the subject’s level of impact),
1004 and realised in the form of an adapted ABAC model (whereby the subject’s
1005 $\langle permisRole, SysAnalyst \rangle$ attribute is removed). The adapted ABAC model
1006 was then assessed by the controller’s planning stage, ensuring the identified
1007 solution does not cause a greater cost to the organisation over observed vi-
1008 olations. As the solution only impacts the malicious subject, and no other
1009 solution is applicable to the impact of the subject, solution S1 is selected.
1010 Solution S1 is enacted as a SOAP message [5] sent to the subject’s rele-
1011 vant identity provider (lgzIS) effector. The effector accepts the request and

1012 removes *emp0003*'s *SysAnalyst* attribute. The consequence of this adapta-
1013 tion is that *emp0003* is no longer able to gain future access to *empDB*, as
1014 the employee lacks the necessary access rights.

1015 In steps 6 to 13, the contractor supervisor (*con0002*) and three other
1016 contractors are detected at points C and D respectively, triggering violation
1017 *empDBShortRead*. Detection results in a similar process to that of the first
1018 attack by *emp0003* (step 5), where each subject's impact is recalculated
1019 and appropriate solutions are identified and enacted. Eventually, each con-
1020 tractor's relevant attributes in accessing the *empDB* resource are identified
1021 and removed via a SOAP message sent to the contractor's identity provider
1022 (*conIS*) effector.

1023 In the final step, *emp0003* rapidly accesses the *lgT* resource, this time
1024 using their remaining attribute $\langle permisRole, Staff \rangle$. This triggered the
1025 violation *lgTShortAccess*, whereby the controller calculates the subject's
1026 impact level as 0.8. Several solutions are now applicable given this impact
1027 weighting, including solutions that result in policy adaptation. However, as
1028 the subject has been identified as the source of two previous violations, but
1029 is the only subject that has abused their *permisRole* attribute of *SysAnalyst*
1030 and *Staff*, solution S2 is enacted. This results in the complete removal of
1031 access for subject *emp0003*.

1032 As a result of subject identity adaptation, malicious behaviour by subjects
1033 were mitigated given the abuse of access rights. Moreover, subject identity
1034 adaptation ensured that non-malicious subjects remained able to request and
1035 gain access to protected resources, as evident by the rate of access maintained
1036 for *all.staff*, as shown in Figure 10.

1037 5.4. Summary of Results

1038 In the experiments performed, adaptation resulted in preventing the de-
1039 tected malicious subject(s) from gaining further access. This was achieved
1040 through removing the abused access right (assigned attribute), removing all
1041 of a subject's access rights at identity provider level, or removing varying
1042 degrees in trust of the contractor identity provider.

1043 In *Exp1* and *Exp2*, no impact was made to authorisation services in terms
1044 of the service being able to perform its duties. This reflects the fundamental
1045 design of SAAF, which promotes separation of concerns between adaptation
1046 and authorisation. In these experiments, the impact on identity provider
1047 services was negligible. There was no observable rise in latency in subject

1048 authentication and attribute release as a result of identity provider adapta-
1049 tion. However, malicious subjects were impacted at identity provider level,
1050 in terms of attribute removal, yet this was the desired consequence of adap-
1051 tation.

1052 Further experiments have been performed using different malicious
1053 changeloads that trigger other self-adaptive strategies for mitigating mali-
1054 cious behaviour [7], and all of them have demonstrated the effectiveness of
1055 the SAAF prototype in handling the violations identified in the malicious
1056 changeloads.

1057 *5.4.1. Subject Identity versus Policy Adaptation*

1058 The experiments portray scenarios that exemplify subject identity adap-
1059 tation and policy adaptation. A scenario where a controller is capable of
1060 performing subject identity adaptation against all identity providers, and a
1061 scenario where the controller is limited in performing subject identity adap-
1062 tation (requiring policy adaptation).

1063 Subject identity adaptation is seen as the economical choice, whereby
1064 malicious behaviour can be mitigated with no impact to non-malicious sub-
1065 jects. When subject identity adaptation was possible, the malicious subjects'
1066 behaviour was mitigated almost immediately, preventing future violations.
1067 However, where subject identity adaptation was not possible, subjects were
1068 capable of repeating violations until the controller identified that the cost of
1069 unresolved violations warranted policy adaptation.

1070 Policy adaptation has far greater consequence in comparison to subject
1071 identity adaptation, which is calculated (in part) by the number of non-
1072 malicious subjects that will lose access to resources as a result of an adapta-
1073 tion.

1074 Regardless of type, each adaptation results in a concrete adjustments to
1075 the authorisation infrastructure. Adaptations ultimately control the outcome
1076 of future access decisions, and whether or not subjects can be authorised in
1077 accessing resources.

1078 *5.4.2. Performance*

1079 Whilst benchmarking the performance is not the main objective for this
1080 paper, the performance observed in the experiments requires some explana-
1081 tion. Of particular interest is the performance of different types of adapta-
1082 tion. Performance is directly related to the number of violations the controller
1083 can identify, the size of its access control model, the number of previously

1084 identified violations, the number of solutions applicable to an identified vi-
1085 olation, and the type of adaptation to be performed. For each experiment,
1086 these factors remained persistent, relative to the given experiment step.

1087 A concern was the high standard deviation observed (max 99ms) between
1088 experiment runs for some adaptations, specifically in regards to the time it
1089 took the controller to react and decide upon solutions. Steady state time was
1090 used to place the controller in a warmed up state. However, due to a mix of
1091 factors the standard deviation failed to improve. Some of these factors in-
1092 clude network fluctuation between communication of the prototype controller
1093 and its effectors, the triggering of Java garbage collection and Java's code
1094 optimisation, and that the controller prototype is yet to be optimised. To
1095 compensate for this, further experiment runs are required, but were limited
1096 to 6 runs per experiment (due to the hour long run-time of each run).

1097 **6. Discussion**

1098 The *LGZLogistics* case study has provided a scenario for demonstrating
1099 and evaluating the detection and mitigation a malicious changeload. This
1100 has been achieve through self-adaptation in the context of a federated envi-
1101 ronment consisting of multiple management domains. Probes and effectors
1102 have shown to facilitate automated adaptation across these management do-
1103 mains, where there is arguably a greater need for automation given the fact
1104 that federations contain large and unknown user bases.

1105 *6.1. Evaluation Approach*

1106 The experiment was designed to demonstrate the resilience of the SAAF
1107 prototype in mitigating malicious behaviour under repeatable conditions.
1108 This required simulating a predefined malicious changeload for triggering self-
1109 adaptation, and capture the mitigating responses from the SAAF prototype.

1110 The evaluation of SAAF prototype was performed using a simulation
1111 of a large scale deployment, akin to a small to medium sized organisation.
1112 This was critical to providing evidence of the SAAF prototype's feasibility in
1113 operating within the real world, and that the prototype would consistently
1114 mitigate violations in a resilient manner. As such, it was observed that the
1115 SAAF prototype was capable of mitigating violations when operating under
1116 high loads, and when faced with non-cooperating management domains.

1117 *6.2. Detection and Adaptation*

1118 The goal of this paper was not to improve existing techniques for detec-
1119 tion malicious behaviour, rather, to demonstrate a new process for handling
1120 insider threat. With that said, detection within the SAAF prototype is worth
1121 discussing. The SAAF prototype uses detectors to identify known types of
1122 attacks, typically focused on thresholds, which is a common approach in
1123 detection of malicious behaviour [17, 41].

1124 Adopting a threshold approach for detecting malicious behaviour has the
1125 advantage of clearly identifying extremes in user behaviour. These rules are
1126 then incorporated into rules defined by experts, knowledgeable in the differ-
1127 ences between normal and abnormal behaviour. However, if a rule is incorrect
1128 or inappropriate for the current state of the system, there is the potential
1129 for many false positives. Clearly a challenge for SAAF is to employ detec-
1130 tion techniques that can evolve and accommodate such legitimate changes in
1131 behaviour.

1132 In SAAF, the decision whether to adapt uses cost sensitive modelling [39]
1133 to assess subject impact and impact of solutions. This approach has al-
1134 lowed the aggregation of multiple violations before enacting the appropriate
1135 solution. Multiple occurrences of violations arguably strengthens the per-
1136 ception in the subject being malicious⁴, as well as adjudicating the extent of
1137 appropriate adaptation. Lastly, through this approach, the organisation in
1138 which the SAAF prototype is being deployed has the ability to fine tune the
1139 enactment of the alternative solutions, in terms of their costs.

1140 In the experiments discussed, the SAAF prototype considers the metric
1141 of rate of access requests as the primary environment property in identifying
1142 malicious behaviour. Whilst using this metric has shown to be successful
1143 in identifying attacks, for it to be efficient, the level of access control must
1144 be fine grained. In addition, a subject's ability to access a resource should
1145 be determined by short term (or one time use) credentials issued by their
1146 identity provider.

1147 The experiments demonstrated the selection and escalation of solutions
1148 in response to detected violations. Whilst this was successful and ultimately
1149 viewed as enacting 'appropriate' solutions to violations, the cost sensitive
1150 modelling approach employed has several limitations. One of the limitation

⁴One exception to this is if the behaviour rules specified are incorrect, which is addressed as part of SAAF's limitations.

1151 is the fact the approach relies upon weighting solutions by a perceived cost
1152 of negative impact to an organisation, which is then compared to a perceived
1153 cost of subject activity. Although not observed within the experiment it-
1154 self, there is potential for multiple solutions in conjunction with observed
1155 behaviour to present identical costs (i.e., benefits) to an organisation.

1156 One issue that was not addressed in the experiments is the presence of
1157 bottlenecks. Given this implementation of SAAF is a prototype, a notable
1158 deficiency in its design is its inability to consider multiple violations during
1159 a single iteration of its feedback loop. If violations are detected during the
1160 prototype’s current analysis of behaviour, multiple violations are queued,
1161 analysed, planned and executed in a sequential manner. This may result in
1162 an increase of the response times when mitigating behaviour identified in the
1163 aforementioned manner, due to failed or redundant adaptations if a previous
1164 adaptation has already resolved the violation.

1165 *6.3. Threats to Validity*

1166 This paper has presented the resilience evaluation of the Self-adaptive
1167 Authorisation Framework (SAAF) prototype in handling and mitigating ma-
1168 licious behaviour by simulating a case study related to insider threats. How-
1169 ever, there are assumptions that may affect the validity of the results.

1170 Regarding internal validity, there is the nature of using case studies and
1171 simulation. Specifically, simulation presents a certain amount of bias whereby
1172 the violations performed are known, and the prototype controller can be con-
1173 figured in an optimum way to best handle such violations. This is the case
1174 in the deployment of the SAAF prototype in which the specification of the
1175 malicious changeload and the development of SAAF was done by the same
1176 person. Therefore, the simulation approach can only be seen as a means to
1177 demonstrate the prototype’s resilience in handling known violations. What
1178 we have not evaluated is how the prototype handles unknown malicious be-
1179 haviour, and in particular, unpredictable changes that might violate known
1180 behavioural patterns. In our understanding, the key threat to validity are
1181 the unknown malicious behaviour, and not so much whether a third party
1182 should be responsible for the specification of either SAAF or the malicious
1183 changeload. Since unpredictable changes are challenging to simulate, it is
1184 vital to evaluate SAAF in a live environment in which real users carry out
1185 malicious behaviour. This has been done [8], and that study complements the
1186 outcomes of this paper that uses case studies, and simulation of a malicious
1187 changeload.

1188 Another threat to internal validity is related to the parameters adopted
1189 for the definition of scenarios, malicious changeloads and violations (see sec-
1190 tions 4.5 and 4.6). In this paper, these were selected in order to stimulate
1191 the SAAF prototype under normal and high load conditions. However, since
1192 there are no clear indicators how these should be selected because it depends
1193 on several factors, such as, application domain and deployment environment,
1194 the key criterion adopted for the selection of these parameters was to max-
1195 imise stress conditions when evaluating the resilience of the SAAF prototype.

1196 Regarding external validity, although we have identified the mali-
1197 cious changeload for a specific case study, the general model of malicious
1198 changeload, briefly presented in Section 4, could be easily instantiated into
1199 different case studies. However, the major challenge would be related to the
1200 deployment of SAAF prototype on a new environment, and this could in-
1201 troduce new vulnerabilities at the level of probes and effectors, for example.
1202 Moreover, changeload does not consider potential attacks to the infrastruc-
1203 ture in which the SAAF prototype is deployed. Since the defined changeload
1204 is restricted to activities related to access control policies and their enforce-
1205 ment, direct attacks to the deployment infrastructure were not considered
1206 because they are outside the scope of the study. For these changes to be
1207 considered, a new changeload definition is required.

1208 While several validity threats exist, the results obtained have shown the
1209 value of using a simulation of a malicious changeload when evaluating the
1210 resilience of self-adaptive authorisation infrastructures.

1211 7. Related Work

1212 In this section, we present related work on self-protecting systems, and
1213 on how resilience benchmarking provides a practical way for characterising
1214 and comparing self-adaptive authorisation infrastructures.

1215 Self-protecting systems are a specialisation of self-adaptive systems with a
1216 goal to mitigating malicious behaviour, and as such, the SAAF prototype can
1217 be considered a self-protecting system. In the following, we discuss few of the
1218 works that have demonstrated self-protection within the context of mitigating
1219 insider attacks. In particular, we discuss two self-protection approaches based
1220 on the state of access control, and one approach based on the state system
1221 architecture.

1222 One of the approaches to self-protection via access control is Securi-
1223 TAS [35]. SecuriTAS is a tool that enables dynamic decisions in awarding

1224 access, which is based on a perceived state of the system and its environ-
1225 ment. SecuriTAS is similar to dynamic access control approaches, such as
1226 RADac [28], in that it has a notion of risk (threat) to resources, and changes
1227 in threat leads to a change in access control decisions. However, it furthers
1228 the concepts in RADac to include the notion of utility. The main difference
1229 between SecuriTAS and SAAF, is that SecuriTAS authorisation infrastruc-
1230 ture incorporates self-adaptation by design, and it is based on its own be-
1231 spoke access control model. SAAF, on the other hand, is a framework that
1232 describes how existing access control models, like ABAC, and legacy autho-
1233 risation infrastructures, like PERMIS [15, 16], can be made self-adaptive, so
1234 that it can be configured to actively mitigate insider threat. With that said,
1235 both approaches demonstrate an authorisation infrastructure’s resilience in
1236 mitigating insider attacks, by ensuring that authorisation remains relevant
1237 to system and environment states (and preventing continuation of attacks by
1238 adaptation of security controls).

1239 In contrast to self-protection via access control, architectural-based self-
1240 protection (ABSP) [43] presents a general solution to detection and miti-
1241 gation of security threats via run-time structural adaptation. Rather than
1242 reason at the contextual layer of ‘access control’, ABSP uses an architectural
1243 model of the running system to identify the extent of impact of identified
1244 attacks. Once attacks or security threats have been assessed, a self-adaptive
1245 architectural manager (Rainbow [19]) is used to perform adaptations to
1246 mitigate the attack. ABSP shares a number of similarities with intrusion
1247 response and prevention systems, particularly, with the scope of adaptations
1248 that ABSP can perform (e.g., structural adaptation against network devices
1249 and connections). However, because ABSP maintains a notion of ‘self’, it is
1250 able to reason about the impact of adaptations and provide assurance over
1251 adaptation before adapting its target system.

1252 Another similar example of self-protection is one proposed by Morin et
1253 al., which takes a novel approach in managing access control, through the use
1254 of architectural adaptation [31]. When access control policies are changed,
1255 the architectural model is updated, resulting in the running system being
1256 reconfigured through adaptation. Morin et al.’s approach shows the effec-
1257 tive deployment of access control across an entire system, where unlike a top
1258 down approach proposed by XACML, there is no centralised point of failure.
1259 A limitation in this approach is that this form of architectural adaptation is
1260 expensive, requiring all resources that need access control to be engineered
1261 in a particular manner, lowering the usefulness of the approach in industry.

1262 In addition, it lacks the ability to automatically evolve and reflect changes
1263 to access control, once malicious behaviour has occurred. However, the tech-
1264 nique poses a novel and viable means of realising a change to access control,
1265 once such a change has been formulated.

1266 Although there has been several publications on self-adaptive security [40,
1267 42], there are few contributions on how these systems should mitigate in a
1268 dynamic way cyber attacks, and in particular, insider threats. However, some
1269 of the existing work complements the work of SAAF by providing means
1270 in order to manage access control policies [21], and looking into selection of
1271 policies based on risk [18]. On other hand, there is similar work to ours, whose
1272 goal is to enabling legacy systems to incorporate dynamic security policies,
1273 which in the case of Al-Ali et al. is by adapting the system architecture [1].

1274 Compared to other established benchmarks (see Section 2.3), a resilience
1275 benchmark can be specified following the same generic principles of bench-
1276 marking, but changeload needs to be revised and expanded, and should
1277 support a risk-based approach for evaluating by comparison the adaptation
1278 mechanisms of a self-adaptive software system [12]. An example of such
1279 changeload is the one applied to the resilience evaluation of an architectural-
1280 based self-adaptive system [11]. In this paper, we have tailored the original
1281 changeload model [12] to the context of self-adaptive authorisation infras-
1282 tructures and focusing on insider threats. There is no other work, to be
1283 best of our knowledge, that has attempted to evaluate the resilience of self-
1284 protecting system in a systematic way, not even providing a generic frame-
1285 work that could be instantiated into a wide range of systems.

1286 8. Conclusions

1287 This paper presented an approach for the resilience evaluation of self-
1288 adaptive authorisation infrastructures. For demonstrating the overall ap-
1289 proach for handling and mitigating malicious behaviour, we have defined a fic-
1290 titious case study of insider threat, defined a repeatable malicious changeload,
1291 and deployed a Self-Adaptive Authorisation Framework (SAAF) prototype.
1292 Changes on the operational conditions included: changes to the run-time load
1293 of the authorisation infrastructure, changes to the autonomic controller, and
1294 changes in the availability of probes and effectors.

1295 The evaluation demonstrated that SAAF was resilient in handling abuse
1296 in access control under repeatable conditions, and that SAAF consistently
1297 mitigated abuse under normal and high loads. In addition, when faced with

1298 restrictions in enacting adaptation, SAAF was able to escalate its selection of
1299 policy adaptation in order to overcome failures in subject identity adaptation.
1300 It was shown that, whilst subject identity adaptation created minimal impact
1301 on non-malicious subjects, it was authorisation policy adaptation that was
1302 effective in halting abuse of access. Finally, the experiments based on SAAF
1303 prototype have also demonstrated its effectiveness in mitigating abuse of
1304 access control in federated environments.

1305 Several limitations could be associated with the proposed approach. In
1306 the security domain, case studies, compared with real systems, are not able
1307 to achieve the same level of coverage regarding the range of attacks the sys-
1308 tem may encounter. On the other hand, the use of case studies and malicious
1309 changeload are fundamental for obtaining the repetitive conditions necessary
1310 for evaluating the resilience of a self-adaptive system. However, in order to
1311 balance out this limitation, we have demonstrated the resilience of SAAF in
1312 the context of a honeypot based deployment that made use of a game [8]. An-
1313 other limitation when evaluating the resilience of self-adaptive authorisation
1314 infrastructures using use cases is the inability of dealing with a wide range
1315 of changes that are representative of unexpected subject behaviour. This
1316 may include, for example, subjects changing their behaviour when reacting
1317 to self-adaptation. This has confirmed that simulation does not consider the
1318 run-time consequences of mitigation based on self-adaptation.

1319 As future work, the plan is to use mutation when specifying a malicious
1320 changeload in order to uncover vulnerabilities that are particular to certain
1321 deployments of self-adaptive authorisation infrastructures.

1322 **References**

- 1323 [1] Al-Ali, R., Hnetyuka, P., Havlik, J., Krivka, V., Heinrich, R., Seifer-
1324 mann, S., Walter, M., Juan-Verdejo, A., 2019. Dynamic security rules
1325 for legacy systems, in: Proceedings of the 13th European Conference
1326 on Software Architecture - Volume 2, ACM, New York, NY, USA.
1327 pp. 277–284. URL: <http://doi.acm.org/10.1145/3344948.3344974>,
1328 doi:10.1145/3344948.3344974.
- 1329 [2] Apache Software Foundation, n.d. JMeter [Online]. Available from:
1330 <http://jmeter.apache.org/> [Accessed 30 September 2018].
- 1331 [3] Bailey, C., Chadwick, D.W., de Lemos, R., 2011. Self-adaptive autho-
1332 rization framework for policy based RBAC/ABAC models, in: Proceed-

- 1333 ings of the 2011 IEEE Ninth International Conference on Dependable,
1334 Autonomic and Secure Computing, IEEE Computer Society, Washing-
1335 ton, DC, USA. pp. 37–44. URL: [http://dx.doi.org/10.1109/DASC.](http://dx.doi.org/10.1109/DASC.2011.31)
1336 2011.31, doi:10.1109/DASC.2011.31.
- 1337 [4] Bailey, C., Chadwick, D.W., de Lemos, R., 2014a. Self-adaptive feder-
1338 ated authorization infrastructures. *Journal of Computer and System Sci-*
1339 *ences* 80, 935–952. URL: [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0022000014000154)
1340 [article/pii/S0022000014000154](http://www.sciencedirect.com/science/article/pii/S0022000014000154), doi:[http://dx.doi.org/10.1016/](http://dx.doi.org/10.1016/j.jcss.2014.02.003)
1341 [j.jcss.2014.02.003](http://dx.doi.org/10.1016/j.jcss.2014.02.003).
- 1342 [5] Bailey, C., Chadwick, D.W., de Lemos, R., Siu, K.W.S., 2013. Enabling
1343 the autonomic management of federated identity providers, in: *Proceed-*
1344 *ings of the 7th IFIP WG 6.6 International Conference on Autonomous*
1345 *Infrastructure, Management, and Security: Emerging Management*
1346 *Mechanisms for the Future Internet - Volume 7943*, Springer-Verlag,
1347 Berlin, Heidelberg. pp. 100–111. URL: [http://dx.doi.org/10.1007/](http://dx.doi.org/10.1007/978-3-642-38998-6_14)
1348 [978-3-642-38998-6_14](http://dx.doi.org/10.1007/978-3-642-38998-6_14), doi:10.1007/978-3-642-38998-6_14.
- 1349 [6] Bailey, C., Montrieux, L., de Lemos, R., Yu, Y., Wermelinger, M., 2014b.
1350 Run-time generation, transformation, and verification of access control
1351 models for self-protection, in: *Proceedings of the 9th International Sym-*
1352 *posium on Software Engineering for Adaptive and Self-Managing Sys-*
1353 *tems*, ACM, New York, NY, USA. pp. 135–144. URL: [http://doi.](http://doi.acm.org/10.1145/2593929.2593945)
1354 [acm.org/10.1145/2593929.2593945](http://doi.acm.org/10.1145/2593929.2593945), doi:10.1145/2593929.2593945.
- 1355 [7] Bailey, C.M., 2015. *Self-adaptive Authorisation Infrastructures*. Ph.D.
1356 thesis. University of Kent,.
- 1357 [8] Bailey, C.M., de Lemos, R., 2018. Evaluating self-adaptive authorisa-
1358 tion infrastructures through gamification, in: *48th Annual IEEE/IFIP*
1359 *International Conference on Dependable Systems and Networks, DSN*
1360 *2018, Luxembourg City, Luxembourg, June 25-28, 2018*, pp. 502–513.
1361 URL: <https://doi.org/10.1109/DSN.2018.00058>, doi:10.1109/DSN.
1362 2018.00058.
- 1363 [9] Box, D., Ehnebuske, D., Kakivaya, G., Layman, A., Mendelsohn, N.,
1364 Nielsen, H.F., Thatte, S., Winer, D., 2000. Simple object access protocol
1365 (SOAP) 1.1.

- 1366 [10] Cámara, J., Correia, P., de Lemos, R., Vieira, M., 2014. Empirical
1367 resilience evaluation of an architecture-based self-adaptive software sys-
1368 tem, in: Proceedings of the 10th International ACM Sigsoft Confer-
1369 ence on Quality of Software Architectures, ACM, New York, NY, USA.
1370 pp. 63–72. URL: <http://doi.acm.org/10.1145/2602576.2602577>,
1371 doi:10.1145/2602576.2602577.
- 1372 [11] Cámara, J., de Lemos, R., Laranjeiro, N., Ventura, R., Vieira, M.,
1373 2017. Robustness-driven resilience evaluation of self-adaptive soft-
1374 ware systems. IEEE Trans. Dependable Secur. Comput. 14, 50–64.
1375 URL: <https://doi.org/10.1109/TDSC.2015.2429128>, doi:10.1109/
1376 TDSC.2015.2429128.
- 1377 [12] Cámara, J., de Lemos, R., Vieira, M., Almeida, R., Ventura, R.,
1378 2013. Architecture-based resilience evaluation for self-adaptive sys-
1379 tems. Computing 95, 689–722. URL: [https://doi.org/10.1007/
1380 s00607-013-0311-7](https://doi.org/10.1007/s00607-013-0311-7), doi:10.1007/s00607-013-0311-7.
- 1381 [13] Cappelli, D.M., Moore, A.P., Trzeciak, R.F., 2012. The CERT Guide to
1382 Insider Threats: How to Prevent, Detect, and Respond to Information
1383 Technology Crimes. 1st ed., Addison-Wesley Professional.
- 1384 [14] Caputo, D., Maloof, M., Stephens, G., 2009. Detecting insider theft
1385 of trade secrets. IEEE Security and Privacy 7, 14–21. URL: [http:
1386 //dx.doi.org/10.1109/MSP.2009.110](http://dx.doi.org/10.1109/MSP.2009.110), doi:10.1109/MSP.2009.110.
- 1387 [15] Chadwick, D.W., Otenko, A., 2002. The PERMIS X.509 role based
1388 privilege management infrastructure, in: Proceedings of the Seventh
1389 ACM Symposium on Access Control Models and Technologies, ACM,
1390 New York, NY, USA. pp. 135–140. URL: [http://doi.acm.org/10.
1391 1145/507711.507732](http://doi.acm.org/10.1145/507711.507732), doi:10.1145/507711.507732.
- 1392 [16] Chadwick, D.W., Zhao, G., Otenko, S., Laborde, R., Su, L., Nguyen,
1393 T.A., 2008. PERMIS: A modular authorization infrastructure. Concurr.
1394 Comput. : Pract. Exper. 20, 1341–1357. URL: [http://dx.doi.org/10.
1395 1002/cpe.v20:11](http://dx.doi.org/10.1002/cpe.v20:11), doi:10.1002/cpe.v20:11.
- 1396 [17] Chung, C.Y., Gertz, M., Levitt, K., 2000. DEMIDS: A misuse de-
1397 tection system for database systems, in: van Biene-Hershey, M.E.,
1398 Strous, L. (Eds.), Integrity and Internal Control Information Systems.

- 1399 Kluwer Academic Publishers, Norwell, MA, USA, pp. 159–178. URL:
1400 <http://dl.acm.org/citation.cfm?id=342030.342078>.
- 1401 [18] Díaz-López, D., Dólera-Tormo, G., Gómez-Mármol, F., Martínez-Pérez,
1402 G., 2016. Dynamic counter-measures for risk-based access control
1403 systems: An evolutive approach. *Future Generation Computer Sys-*
1404 *tems* 55, 321 – 335. URL: [http://www.sciencedirect.com/science/](http://www.sciencedirect.com/science/article/pii/S0167739X14002052)
1405 [article/pii/S0167739X14002052](http://www.sciencedirect.com/science/article/pii/S0167739X14002052), doi:[https://doi.org/10.1016/j.](https://doi.org/10.1016/j.future.2014.10.012)
1406 [future.2014.10.012](https://doi.org/10.1016/j.future.2014.10.012).
- 1407 [19] Garlan, D., Cheng, S.W., Huang, A.C., Schmerl, B., Steenkiste, P., 2004.
1408 *Rainbow: Architecture-based self-adaptation with reusable infrastruc-*
1409 *ture*. *Computer* 37, 46–54. URL: [http://dx.doi.org/10.1109/MC.](http://dx.doi.org/10.1109/MC.2004.175)
1410 [2004.175](http://dx.doi.org/10.1109/MC.2004.175), doi:[10.1109/MC.2004.175](https://doi.org/10.1109/MC.2004.175).
- 1411 [20] Gray, J., 1992. *Benchmark Handbook: For Database and Transaction*
1412 *Processing Systems*. Morgan Kaufmann Publishers Inc., San Francisco,
1413 CA, USA.
- 1414 [21] Hummer, M., Kunz, M., Netter, M., Fuchs, L., Pernul, G., 2016.
1415 Adaptive identity and access management—contextual data based
1416 policies. *EURASIP Journal on Information Security* 2016, 19.
1417 URL: <https://doi.org/10.1186/s13635-016-0043-2>, doi:[10.1186/](https://doi.org/10.1186/s13635-016-0043-2)
1418 [s13635-016-0043-2](https://doi.org/10.1186/s13635-016-0043-2).
- 1419 [22] IBM, n.d. *IBM Security Intelligence with Big Data* [Online]. Available
1420 from: <http://www-03.ibm.com/security/solution/intelligence-big-data/>
1421 [Accessed 30 September 2018].
- 1422 [23] Kanoun, K., Spainhower, L., 2008. *Dependability Benchmarking for*
1423 *Computer Systems*. Wiley-IEEE Computer Society Pr.
- 1424 [24] Kephart, J.O., Chess, D.M., 2003. The vision of autonomic computing.
1425 *Computer* 36, 41–50. URL: [http://dx.doi.org/10.1109/MC.2003.](http://dx.doi.org/10.1109/MC.2003.1160055)
1426 [1160055](http://dx.doi.org/10.1109/MC.2003.1160055), doi:[10.1109/MC.2003.1160055](https://doi.org/10.1109/MC.2003.1160055).
- 1427 [25] Koutsonikola, V., Vakali, A., 2004. LDAP: Framework, practices, and
1428 trends. *IEEE Internet Computing* 8, 66–72. URL: [http://dx.doi.org/](http://dx.doi.org/10.1109/MIC.2004.44)
1429 [10.1109/MIC.2004.44](http://dx.doi.org/10.1109/MIC.2004.44), doi:[10.1109/MIC.2004.44](https://doi.org/10.1109/MIC.2004.44).

- 1430 [26] de Lemos, R., et al, 2013. Software engineering for self-adaptive
1431 systems: A second research roadmap, in: de Lemos, R., Giese, H.,
1432 Müller, H., Shaw, M. (Eds.), *Software Engineering for Self-Adaptive*
1433 *Systems II*. Springer Berlin Heidelberg. volume 7475 of *Lecture Notes*
1434 *in Computer Science*, pp. 1–32. URL: [http://dx.doi.org/10.1007/](http://dx.doi.org/10.1007/978-3-642-35813-5_1)
1435 [978-3-642-35813-5_1](http://dx.doi.org/10.1007/978-3-642-35813-5_1), doi:10.1007/978-3-642-35813-5_1.
- 1436 [27] Madeira, H., 2005. Towards a security benchmark for database manage-
1437 ment systems, in: *Proceedings of the 2005 International Conference on*
1438 *Dependable Systems and Networks*, IEEE Computer Society, Washing-
1439 ton, DC, USA. pp. 592–601. URL: [http://dx.doi.org/10.1109/DSN.](http://dx.doi.org/10.1109/DSN.2005.93)
1440 [2005.93](http://dx.doi.org/10.1109/DSN.2005.93), doi:10.1109/DSN.2005.93.
- 1441 [28] McGraw, R., 2009. Risk-Adaptable Access Control (RADac). Technical
1442 Report. National Institute of Standards and Technology (NIST).
- 1443 [29] Montrieux, L., 2013. Model-Based Analysis of Role-Based Access Con-
1444 trol. Ph.D. thesis. The Open University.
- 1445 [30] Montrieux, L., de Lemos, R., Bailey, C., 2019. Challenges in engi-
1446 neering self-adaptive authorisation infrastructures, in: Yu, Y., Ban-
1447 dara, A., Honiden, S., Hu, Z., Tamai, T., Muller, H., Mylopoulos, J.,
1448 Nuseibeh, B. (Eds.), *Engineering Adaptive Software Systems: Com-*
1449 *munications of NII Shonan Meetings*. Springer Singapore, Singapore,
1450 pp. 57–94. URL: https://doi.org/10.1007/978-981-13-2185-6_3,
1451 doi:10.1007/978-981-13-2185-6_3.
- 1452 [31] Morin, B., Mouelhi, T., Fleurey, F., Le Traon, Y., Barais, O.,
1453 Jézéquel, J.M., 2010. Security-driven model-based dynamic adap-
1454 tation, in: *Proceedings of the IEEE/ACM International Conference*
1455 *on Automated Software Engineering*, ACM, New York, NY, USA.
1456 pp. 205–214. URL: <http://doi.acm.org/10.1145/1858996.1859040>,
1457 doi:10.1145/1858996.1859040.
- 1458 [32] Nurse, J.R., Buckley, O., Legg, P.A., Goldsmith, M., Creese, S., Wright,
1459 G.R., Whitty, M., 2014. Understanding insider threat: A frame-
1460 work for characterising attacks, in: *Workshop on Research for Insider*
1461 *Threat (WRIT) held as part of the IEEE Computer Society Secu-*
1462 *rity and Privacy Workshops (SPW14)*, in conjunction with the IEEE

- 1463 Symposium on Security and Privacy (SP), IEEE. pp. 214–228. URL:
1464 <http://www.sei.cmu.edu/community/writ2014/>.
- 1465 [33] OASIS, 2005. Security Assertion Markup Language (SAML) Version
1466 2.0. Technical Report. OASIS.
- 1467 [34] Oltsik, J., 2013. The 2013 Vormetric insider
1468 threat report [Online]. Available from:
1469 [http://www.vormetric.com/sites/default/files/vormetric-insider-threat-](http://www.vormetric.com/sites/default/files/vormetric-insider-threat-report-oct-2013.pdf)
1470 [report-oct-2013.pdf](http://www.vormetric.com/sites/default/files/vormetric-insider-threat-report-oct-2013.pdf) [Accessed 30 September 2018].
- 1471 [35] Pasquale, L., Menghi, C., Salehie, M., Cavallaro, L., Omoronyia, I., Nu-
1472 seibeh, B., 2012. Securitas: A tool for engineering adaptive security,
1473 in: Proceedings of the ACM SIGSOFT 20th International Symposium
1474 on the Foundations of Software Engineering, ACM, New York, NY,
1475 USA. pp. 19:1–19:4. URL: [http://doi.acm.org/10.1145/2393596.](http://doi.acm.org/10.1145/2393596.2393618)
1476 [2393618](http://doi.acm.org/10.1145/2393596.2393618), doi:10.1145/2393596.2393618.
- 1477 [36] PERMIS Standalone Authorisation Server, n.d. [Online]. Available
1478 from: <http://sec.cs.kent.ac.uk/permis/> [Accessed 30 September 2018].
- 1479 [37] SimpleSAMLphp, n.d. [Online]. Available from:
1480 <http://simplesamlphp.org/> [Accessed 30 September 2018].
- 1481 [38] Spitzner, L., 2003. Honeypots: Catching the insider threat, in: Proceed-
1482 ings of the 19th Annual Computer Security Applications Conference,
1483 IEEE. pp. 170–179.
- 1484 [39] Strasburg, C., Stakhanova, N., Basu, S., Wong, J.S., 2009. A framework
1485 for cost sensitive assessment of intrusion response selection, in: Proceed-
1486 ings of the 2009 33rd Annual IEEE International Computer Software and
1487 Applications Conference - Volume 01, IEEE Computer Society, Wash-
1488 ington, DC, USA. pp. 355–360. URL: [http://dx.doi.org/10.1109/](http://dx.doi.org/10.1109/COMPSAC.2009.54)
1489 [COMPSAC.2009.54](http://dx.doi.org/10.1109/COMPSAC.2009.54), doi:10.1109/COMPSAC.2009.54.
- 1490 [40] Tziakouris, G., Bahsoon, R., Babar, M.A., 2018. A survey on self-
1491 adaptive security for large-scale open environments. *ACM Comput.*
1492 *Surv.* 51, 100:1–100:42. URL: <http://doi.acm.org/10.1145/3234148>,
1493 doi:10.1145/3234148.

- 1494 [41] Wang, K., Stolfo, S., 2004. Anomalous payload-based network in-
1495 trusion detection, in: Jonsson, E., Valdes, A., Almgren, M. (Eds.),
1496 Recent Advances in Intrusion Detection. Springer Berlin Heidel-
1497 berg. volume 3224 of *Lecture Notes in Computer Science*, pp. 203–
1498 222. URL: http://dx.doi.org/10.1007/978-3-540-30143-1_11,
1499 doi:10.1007/978-3-540-30143-1_11.
- 1500 [42] Yuan, E., Esfahani, N., Malek, S., 2014. A systematic survey of self-
1501 protecting software systems. *ACM Trans. Auton. Adapt. Syst.* 8, 17:1–
1502 17:41. URL: <http://doi.acm.org/10.1145/2555611>, doi:10.1145/
1503 2555611.
- 1504 [43] Yuan, E., Malek, S., Schmerl, B., Garlan, D., Gennari, J., 2013.
1505 Architecture-based self-protecting software systems, in: Proceedings of
1506 the 9th international ACM Sigsoft conference on Quality of software
1507 architectures, ACM. pp. 33–42.