



**NIST Internal Report
NIST IR 8504 ipd**

Access Control on NoSQL Databases

Initial Public Draft

Vincent C. Hu

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.IR.8504.ipd>

**NIST Internal Report
NIST IR 8504 ipd**

Access Control on NoSQL Databases

Initial Public Draft

Vincent C. Hu
*Computer Security Division
Information Technology Laboratory*

This publication is available free of charge from:
<https://doi.org/10.6028/NIST.IR.8504.ipd>

January 2024



U.S. Department of Commerce
Gina M. Raimondo, Secretary

National Institute of Standards and Technology
Laurie E. Locascio, NIST Director and Under Secretary of Commerce for Standards and Technology

Certain equipment, instruments, software, or materials, commercial or non-commercial, are identified in this paper in order to specify the experimental procedure adequately. Such identification does not imply recommendation or endorsement of any product or service by NIST, nor does it imply that the materials or equipment identified are necessarily the best available for the purpose.

There may be references in this publication to other publications currently under development by NIST in accordance with its assigned statutory responsibilities. The information in this publication, including concepts and methodologies, may be used by federal agencies even before the completion of such companion publications. Thus, until each publication is completed, current requirements, guidelines, and procedures, where they exist, remain operative. For planning and transition purposes, federal agencies may wish to closely follow the development of these new publications by NIST.

Organizations are encouraged to review all draft publications during public comment periods and provide feedback to NIST. Many NIST cybersecurity publications, other than the ones noted above, are available at <https://csrc.nist.gov/publications>.

NIST Technical Series Policies

[Copyright, Use, and Licensing Statements](#)

[NIST Technical Series Publication Identifier Syntax](#)

Publication History

Approved by the NIST Editorial Review Board on YYYY-MM-DD [Will be added in final publication.]

How to Cite this NIST Technical Series Publication:

Hu VC (2024) Access Control on NoSQL Databases. (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or Internal Report (IR) NIST IR 8504 ipd. <https://doi.org/10.6028/NIST.IR.8504.ipd>

Author ORCID iDs

Vincent C. Hu: 0000-0002-1648-0584

Public Comment Period

January 29, 2024 - March 15, 2024

Submit Comments

ir8504-comments@nist.gov

National Institute of Standards and Technology
Attn: Computer Security Division, Information Technology Laboratory
100 Bureau Drive (Mail Stop 8930) Gaithersburg, MD 20899-8930

All comments are subject to release under the Freedom of Information Act (FOIA).

1 **Abstract**

2 NoSQL database systems and data stores often outperform traditional RDBMS in various
3 aspects, such as data analysis efficiency, system performance, ease of deployment,
4 flexibility/scalability of data management, and users' availability. However, with an increasing
5 number of people storing sensitive data in NoSQL databases, security issues have become
6 critical concerns. NoSQL databases suffer from vulnerabilities, particularly due to the lack of
7 effective support for data protection, including weak authorization mechanisms. As access
8 control is a fundamental data protection requirement of any database management system
9 DBMS, this document focuses on access control on NoSQL database systems.

10 **Keywords**

11 access control; attribute-based access control; authorization; database systems; No-SQL; SQL.

12 **Reports on Computer Systems Technology**

13 The Information Technology Laboratory (ITL) at the National Institute of Standards and
14 Technology (NIST) promotes the U.S. economy and public welfare by providing technical
15 leadership for the Nation's measurement and standards infrastructure. ITL develops tests, test
16 methods, reference data, proof of concept implementations, and technical analyses to advance
17 the development and productive use of information technology. ITL's responsibilities include
18 the development of management, administrative, technical, and physical standards and
19 guidelines for the cost-effective security and privacy of other than national security-related
20 information in federal information systems.

21 **Call for Patent Claims**

22 This public review includes a call for information on essential patent claims (claims whose use
23 would be required for compliance with the guidance or requirements in this Information
24 Technology Laboratory (ITL) draft publication). Such guidance and/or requirements may be
25 directly stated in this ITL Publication or by reference to another publication. This call also
26 includes disclosure, where known, of the existence of pending U.S. or foreign patent
27 applications relating to this ITL draft publication and of any relevant unexpired U.S. or foreign
28 patents.

29 ITL may require from the patent holder, or a party authorized to make assurances on its behalf,
30 in written or electronic form, either:

- 31 a) assurance in the form of a general disclaimer to the effect that such party does not hold
32 and does not currently intend holding any essential patent claim(s); or
- 33 b) assurance that a license to such essential patent claim(s) will be made available to
34 applicants desiring to utilize the license for the purpose of complying with the guidance
35 or requirements in this ITL draft publication either:
 - 36 i. under reasonable terms and conditions that are demonstrably free of any unfair
37 discrimination; or
 - 38 ii. without compensation and under reasonable terms and conditions that are
39 demonstrably free of any unfair discrimination.

40 Such assurance shall indicate that the patent holder (or third party authorized to make
41 assurances on its behalf) will include in any documents transferring ownership of patents
42 subject to the assurance, provisions sufficient to ensure that the commitments in the assurance
43 are binding on the transferee, and that the transferee will similarly include appropriate
44 provisions in the event of future transfers with the goal of binding each successor-in-interest.

45 The assurance shall also indicate that it is intended to be binding on successors-in-interest
46 regardless of whether such provisions are included in the relevant transfer documents.

47 Such statements should be addressed to: ir8504-comments@nist.gov

48

49	Table of Contents	
50	Executive Summary	1
51	1. Introduction	2
52	2. Overview of NoSQL Database Systems	3
53	2.1. Types of NoSQL Databases Systems	3
54	2.1.1. Key-Value Model	3
55	2.1.2. Wide-Column Model	4
56	2.1.3. Document Model.....	4
57	2.1.4. Graph Model.....	5
58	2.2. Features of NoSQL Database Systems	6
59	2.2.1. Flexible Data Model.....	6
60	2.2.2. Horizontal Scalability	6
61	2.2.3. Fast Queries	6
62	2.2.4. Security Deficiencies.....	7
63	3. Access Control Model on NoSQL Database Systems	8
64	3.1. Relationship Structures of NoSQL Models	8
65	3.2. Access Control Rules from NoSQL Relationship Structures	9
66	3.2.1. Key-Value Model	10
67	3.2.2. Wide-Column and Document Models	10
68	3.2.3. Graph Model.....	12
69	3.3. Access Control Model Implementations.....	13
70	4. Considerations of Access Control for NoSQL Systems	15
71	4.1. Fine-Grained Access Control	15
72	4.2. Security.....	15
73	4.3. Query Language	16
74	4.4. Data Consistency	16
75	4.5. Performance.....	17
76	4.6. Audit.....	17
77	4.7. Environment Conditions Control	17
78	4.8. Support for AI.....	18
79	4.9. Combine RDBMS	18
80	5. Conclusion	19
81	References	20
82	Appendix A	22
83	A.1. Federation of Access Control.....	22

84 A.2. Graph NoSQL Implementation for AC Policies..... 23

85 **List of Tables**

86 **Table 1. Comparison of RDBMS and NoSQL features6**

87 **List of Figures**

88 **Fig. 1. Key-value database model3**

89 **Fig. 2. Wide-column database model4**

90 **Fig. 3. Document database model.....5**

91 **Fig. 4. Graph database model5**

92 **Fig. 5. Hierarchical structure relationships of the key-value, wide-column, and document models8**

93 **Fig. 6. Directed graph relationship structure of the graph model9**

94 **Fig. 7. Example of defining subject and object attributes for the key-value model.....10**

95 **Fig. 8. Example of defining subject and object attributes for wide-column and document models11**

96 **Fig. 9. Example of defining subject and object attributes for a graph model.....12**

97 **Fig. 10. Example of defining subject and object attributes for graph models using edges for permitted**

98 **actions to embed an access control policy13**

99 **Fig. 11. Generic resource federation model.....23**

100 **Fig. 12. Resource federation scheme.....23**

101

102 **Acknowledgments**

103 The author would like to express his thanks to Isabel Van Wyk and Jim Foti of NIST for their
104 detailed editorial review of both the public comment version and the final publication.

105 **Executive Summary**

106 NoSQL stands for “not only SQL” or “non-SQL,” which typically refer to any non-relational
107 database that stores data in a format other than relational tables. It shifts away from relational
108 databases (RDBMS) for dealing with enormous and constantly growing data and infrastructure
109 needs and increasingly uses the Web 3.0 framework for big data and real-time web
110 applications, including handling unstructured data like documents, emails, and social media.
111 NoSQL databases are particularly useful for managing unstructured or very large data objects
112 stored across distributed and cooperating devices. Use cases for retrieving such data range
113 from critical scenarios (e.g., storing financial data and healthcare records) to more casual
114 applications (e.g., chat log data, video, and images, readings from smart devices). Major Web
115 2.0 companies have developed or adopted different flavors of NoSQL databases to meet their
116 growing data and infrastructure needs.

117 NoSQL database systems and data stores often outperform traditional RDBMS in various
118 aspects, such as data analysis efficiency, system performance, ease of deployment,
119 flexibility/scalability of data management, and users’ availability. However, with an increasing
120 number of people storing sensitive data in NoSQL databases, security issues have become
121 critical concerns. NoSQL databases suffer from vulnerabilities, particularly due to the lack of
122 effective support for data protection, including weak authorization mechanisms. As access
123 control is a fundamental data protection requirement of any database management system
124 DBMS, this document discusses access control on NoSQL database systems by illustrating the
125 NoSQL database types along with their support for access control models and describing
126 considerations from the perspective of access control.

127

128 1. Introduction

129 NoSQL stands for “not only SQL” or “non-SQL,” which typically refer to any non-relational
130 database that stores data in a format other than relational tables. It emerged in the late 2000s,
131 triggered by the growing popularity of distributed systems, such as cloud computing and big
132 data, which required a shift away from relational databases (RDBMS) toward NoSQL databases.
133 Organizations now deal with enormous and constantly growing data and infrastructure needs
134 or increasingly use the Web 3.0 framework for big data and real-time web applications,
135 including handling unstructured data like documents, emails, and social media [AA, NO].

136 NoSQL databases are particularly useful for managing unstructured or very large data objects
137 stored across distributed and cooperating devices. Use cases for retrieving such data range
138 from critical scenarios (e.g., storing financial data and healthcare records) to more casual
139 applications (e.g., chat log data, video, images, and readings from smart devices). Major Web
140 2.0 companies like Amazon (Dynamo), Google (BigTable), LinkedIn (Voldemort), and Facebook
141 (Cassandra) have developed or adopted different flavors of NoSQL databases to meet their
142 growing data and infrastructure needs. Their success has inspired many of today’s NoSQL
143 applications [AA, CJ].

144 NoSQL database systems and data stores often outperform traditional RDBMS in various
145 aspects, such as data analysis efficiency, system performance, ease of deployment,
146 flexibility/scalability of data management, and users’ availability. However, with an increasing
147 number of people storing sensitive data in NoSQL databases, security issues have become
148 critical concerns. NoSQL databases suffer from vulnerabilities, particularly due to the lack of
149 effective support for data protection, including weak authorization mechanisms. As access
150 control is a fundamental data protection requirement of any database management system
151 (DBMS) [FF], this document discusses access control on NoSQL database systems by illustrating
152 the NoSQL database types along with their support for access control models and describing
153 considerations from the perspective of access control. Note that an access control system may
154 store and manage access control data (e.g., subjects, objects, actions, and attributes) in
155 external systems rather than the NoSQL database itself, which have a wide range of different
156 implementations are not discussed in this document.

157 This document is organized as follows:

- 158 • Section 1 is the introduction,
- 159 • Section 2 provides an overview of NoSQL database systems,
- 160 • Section 3 introduces access control models for NoSQL database systems,
- 161 • Section 4 describes considerations for NoSQL systems from the perspective of access
162 control.
- 163 • Section 5 is the conclusion,
- 164 • References list articles referred by this document,
- 165 • Appendix A provides an example application of Graph model.

166 **2. Overview of NoSQL Database Systems**

167 Applications in web services, e-commerce, mobile computing, and social media require storing
168 and processing vast amounts of structured and unstructured data driven by the significant
169 increase in global datasets known as “big data,” which includes data with high volume, velocity,
170 and variety. However, handling such immense data becomes increasingly complicated in terms
171 of processing and meeting users’ requirements, such as scalability, performance control, high
172 availability, low latency, workload distribution, and managing big data applications [AA].

173 Traditional RDBMSs often fall short when compared to NoSQL database systems, which utilize
174 distributed and collaborative devices to store and retrieve data. NoSQL databases offer the
175 flexibility to rapidly adapt to changes in the software stack and allow data to be distributed
176 across multiple servers and regions in the cloud, scale out instead of scaling up, and intelligently
177 geo-place data to optimize performance [MO].

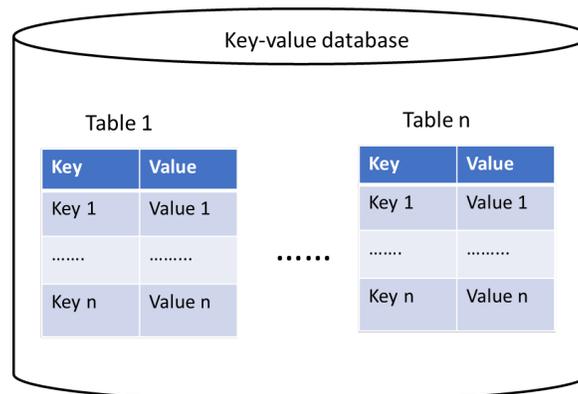
178 **2.1. Types of NoSQL Databases Systems**

179 In general, there are four major types of NoSQL models: key-value, document, wide-column,
180 and graph databases.

181 **2.1.1. Key-Value Model**

182 The key-value model, as shown in Fig. 1, stores data in a schemaless form, making them simple,
183 efficient, and powerful. In this model, each database item contains keys and corresponding
184 values, similar to an RDBMS with only two columns (i.e., key and the value). Data can be
185 efficiently retrieved by using a unique key, which serves as an index like a hash table. The values
186 in a key-value NoSQL database can be simple data types, like strings and numbers, or complex
187 objects [HI]. The key-value model is primarily used for caching, session management, and
188 leaderboard applications [WT].

189



190

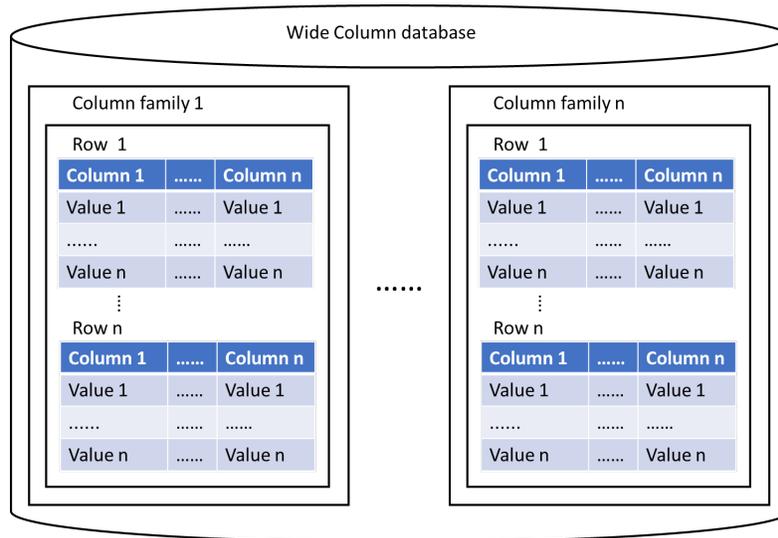
191

Fig. 1. Key-value database model

192 **2.1.2. Wide-Column Model**

193 The wide-column model, as shown in Fig. 2, stores data in tables, rows, and dynamic columns,
194 similar to the key-value model. However, the key in this model is an integration of the row,
195 column, and/or timestamp and refers to one or many columns as a "column family." Each
196 column family is equivalent to a table in an RDBMS, enabling analytics on a small number of
197 columns for data mining and web applications. This design allows for more efficient data
198 reading and retrieval with higher speed compared to traditional RDBMS [AA, HI]. It is well-
199 suited for complex datasets and storing large amounts of data in distributed systems as it
200 enables the easy addition of new columns by creating new files, eliminating the need to rebuild
201 the entire table, as required in RDBMS [AA]. The model is usually optimized for specific use
202 cases (e.g., Facebook's Inbox search feature), allowing it to efficiently handle over 100 million
203 users continuously using the system [OG]. Other popular use cases of the wide-column model
204 include the Internet of Things (IoT), inventory management, and big data processing.

205
206



207
208

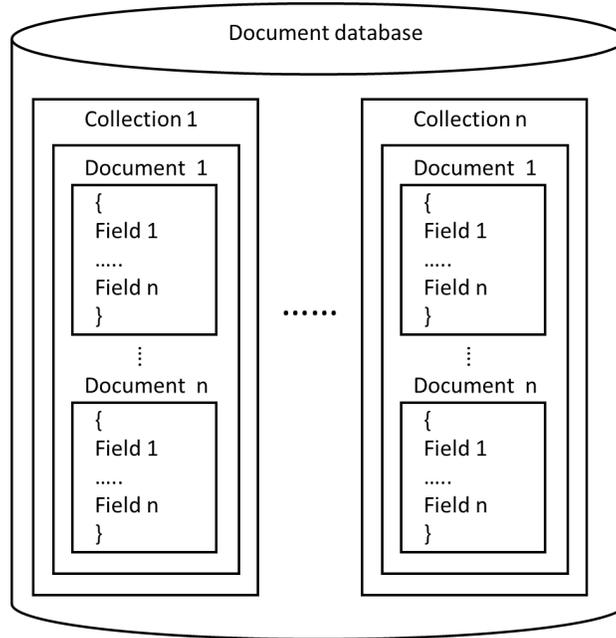
Fig. 2. Wide-column database model

209 **2.1.3. Document Model**

210 The document model, as shown in Fig. 3, stores data in documents, similar to JSON, BSON, and
211 XML documents. In this model, data is organized into collections with a unique key that serves
212 as an index for fast querying. A collection contains documents without a schema. Rather, each
213 document contains pairs of fields and values of various types, such as strings, numbers,
214 Booleans, arrays, and others. The data within documents can include structured data, semi-
215 structured data (e.g., XML files), or unstructured data (e.g., text), which allows for a dynamic
216 structure and the easy modification, addition, or deletion of fields. This flexibility allows
217 documents to be stored and retrieved in a form that closely resembles the data objects used in
218 applications, reducing the need for data translation during application use [HI] and contributing

219 to high performance and horizontal scalability. This model has applications in blog software,
220 content management systems [AA], and product catalogs, big data, and analytics [WT].

221



222

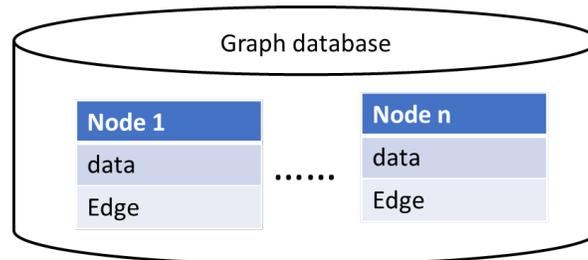
223

Fig. 3. Document database model

224 2.1.4. Graph Model

225 The graph model, as shown in Fig. 4, represents and stores data based on relationships. It is a
226 schemaless model, and data is structured using nodes and edges. Nodes typically store data
227 entities, while edges represent the relationships or links between the data nodes. The model is
228 scalable but complex as it often utilizes shortest path algorithms to optimize data queries for
229 real-time results. In this model, the efficiency of a query depends on the number of
230 relationships among the nodes [HI]. Graph-based databases have various applications, such as
231 recommendation systems, social networking, identity and access management (IAM), and
232 content management [AA WT].

233



234

235

Fig. 4. Graph database model

236 **2.2. Features of NoSQL Database Systems**

237 NoSQL databases offer flexible and simple data models, horizontal scalability, and fast (i.e.,
238 primitive) queries, as well as security deficiencies [MO]. Table 1 compares the features of
239 RDBMS with those of NoSQL databases [AH, SL].

240 **Table 1. Comparison of RDBMS and NoSQL features**

Database Model	RDBMS	NoSQL
Database	Relational	Non-relational
Scheme	Fixed and structured	Dynamic and unstructured
Queries	Complex queries using JOIN	Unsupported
Scalable	Vertical	Horizontal
Properties	Atomicity, consistency, isolation, and durability (ACID)	Consistency (eventually), availability, partial tolerance (CAP)

241 **2.2.1. Flexible Data Model**

242 NoSQL databases provide flexible schemas that enable easy database changes and seamless
243 integration with database applications. For example, in document-based systems, documents
244 do not need to follow the same schema, which allows for faster creation and maintenance of
245 documents with minimal overhead.

246 **2.2.2. Horizontal Scalability**

247 NoSQL databases support horizontal scaling, which means that they handle increased capacity
248 by distributing data across multiple servers, thus avoiding the need to migrate to a larger server
249 when capacity exceeds the requirements of the current server. NoSQL key-value systems have a
250 simple structure with only two columns (i.e., key and the value), enabling horizontal scaling
251 without the need for complex field joins. Similarly, wide-column systems can handle more
252 complex data structures by adding new columns through the creation of a new file.

253 **2.2.3. Fast Queries**

254 NoSQL databases store data in a way that optimizes queries by utilizing simple and efficient
255 (i.e., primitive) query language [OG] without the need for join operations that can degrade
256 query performance in typically normalized DBMS using SQL. For example, in document systems
257 with open formats like XML and JSON, building a document does not require foreign keys. This
258 allows for dynamic relationships between documents, making them independent of each other.
259 Wide-column systems are also designed for efficient data reading and retrieval, resulting in
260 better performance [HI].

261 **2.2.4. Security Deficiencies**

262 Despite their advantages, NoSQL databases lack a mechanism for handling and managing data
263 consistency and maintaining integrity constraints (e.g., using foreign keys). Additionally, they
264 often have limited support for security at the database level [OG].

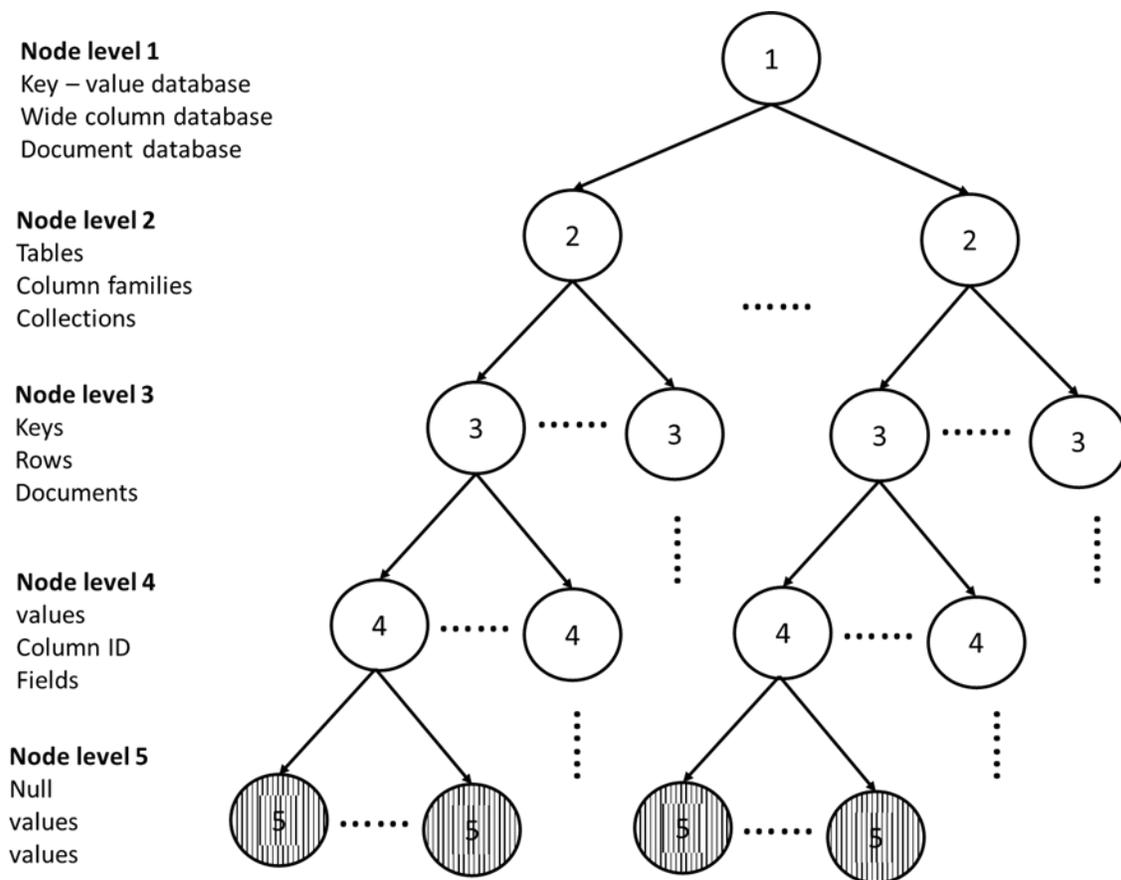
265 3. Access Control Model on NoSQL Database Systems

266 This section illustrates how to apply access control properties on NoSQL models assuming that
267 the access control system stores and manage access control data (e.g., subjects, objects, and
268 attributes) in the NoSQL database.

269 3.1. Relationship Structures of NoSQL Models

270 Schemaless relationships between NoSQL data can be constructed using hierarchical structures.
271 The key-value, wide-column, and document models can be represented in a tree data structure,
272 as shown in Fig. 5, with a maximum of five tree levels if the NoSQL database itself is
273 represented at the first (i.e., root) level.

274



275
276

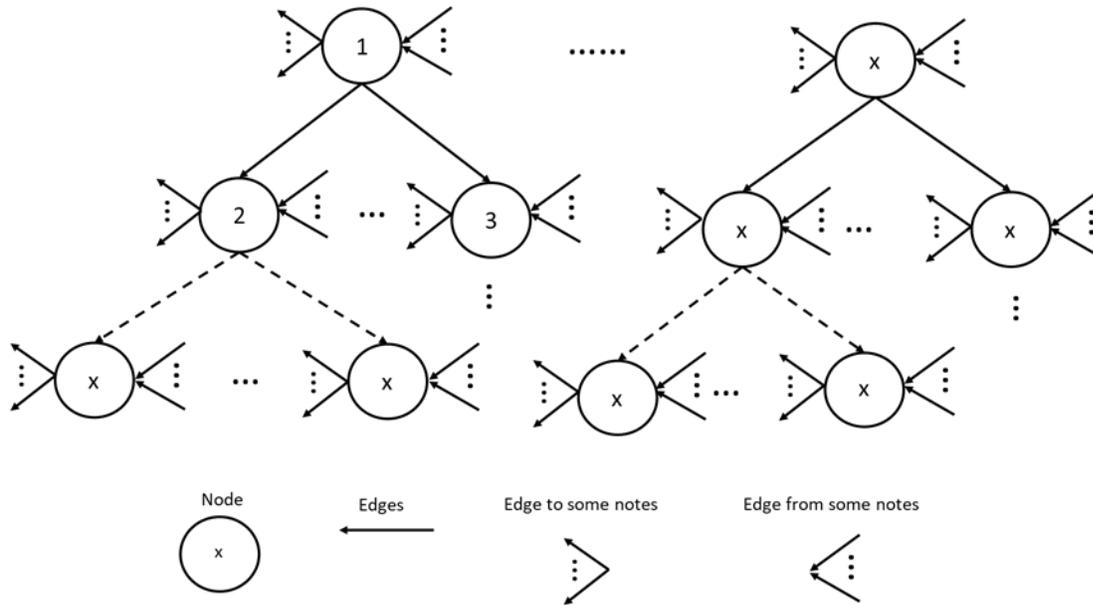
277 **Fig. 5. Hierarchical structure relationships of the key-value, wide-column, and document models**

278 Nodes in the second tree level represent the Tables, Column families, and Collections for key-
279 value, wide-column, and document models, respectively. Nodes in the third tree level represent
280 the data that can be indexed from the nodes in the second tree level, which are Key, Rows, and
281 Documents for the key-value, wide-column, and document models, respectively. Nodes in the
282 fourth tree level represent the data that can be indexed from nodes in the third tree level,

283 which are values, Column IDs, and Fields for the key-value, wide-column, and document
284 models, respectively. Nodes in the fifth tree level represent the data that can be indexed from
285 nodes in the fourth tree level. There is no offspring on this level for the key-value model. It is
286 essential to consider these hierarchical structures when designing access control on NoSQL
287 databases to effectively leverage their schemaless nature and optimize data access.

288 Unlike hierarchical tree structures, the relationship structure of the graph model is represented
289 by a directed graph, as shown in Fig. 6.

290



291
292

293 **Fig. 6. Directed graph relationship structure of the graph model**

294 In this model, each node can have more than one edge, which represent the connections or
295 relationships between nodes in the graph. The directed graph allows for complex and
296 interconnected data relationships, making it ideal for applications in which data connections
297 and dependencies are critical to the analysis and processing of information.

298 **3.2. Access Control Rules from NoSQL Relationship Structures**

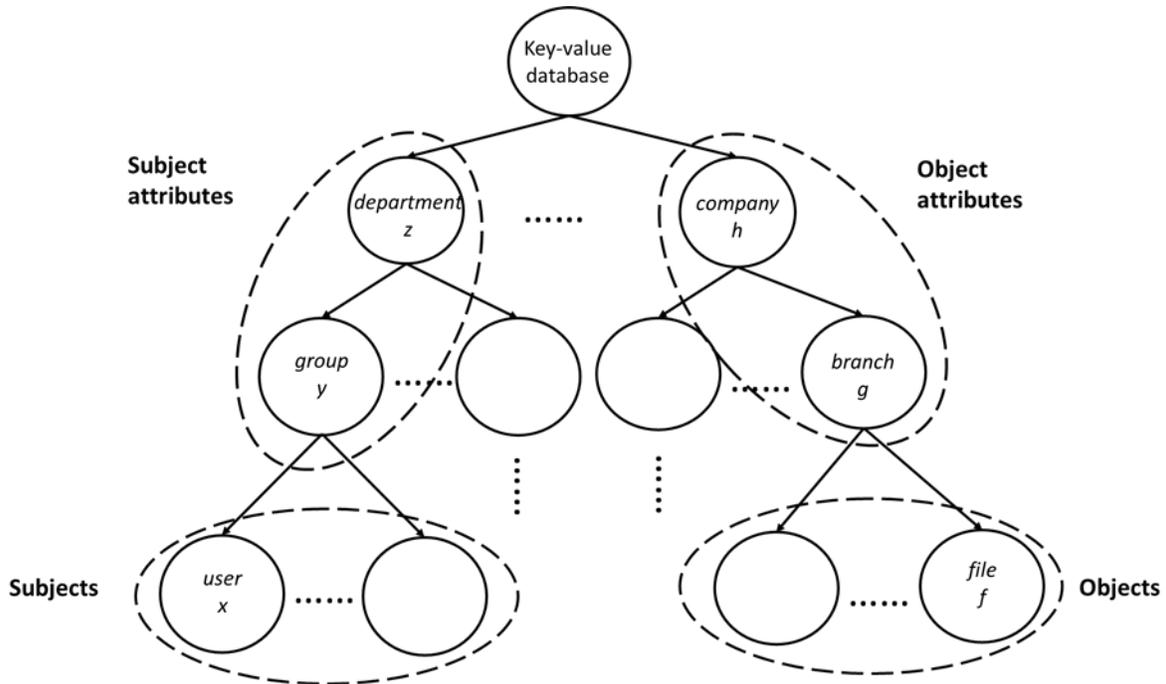
299 Enforcing access control policies — especially for fine-grained access control on schemaless
300 NoSQL databases with data relationships in heterogeneous structures — presents challenges
301 due to the absence of a reference data model and related manipulation language. This
302 exacerbates the enforcement of access control policies on the protected data [CF]. Additionally,
303 managing access control on NoSQL databases is not as straightforward as it is for RDBMSs,
304 making it challenging to query “who can access what” in a straightforward manner. However,
305 any static access control policy model based on subjects and object attributes can still be
306 applied to NoSQL databases. The following sections describe how access control rules can be

307 specified by attributes in the relationship structures for each of the NoSQL models presented in
308 Sec. 3.1.

309 3.2.1. Key-Value Model

310 In the key-value model, a subject's attributes can be specified through the link relationships
311 from level 2 nodes to level 3 nodes, as illustrated in Fig. 7.

312



313

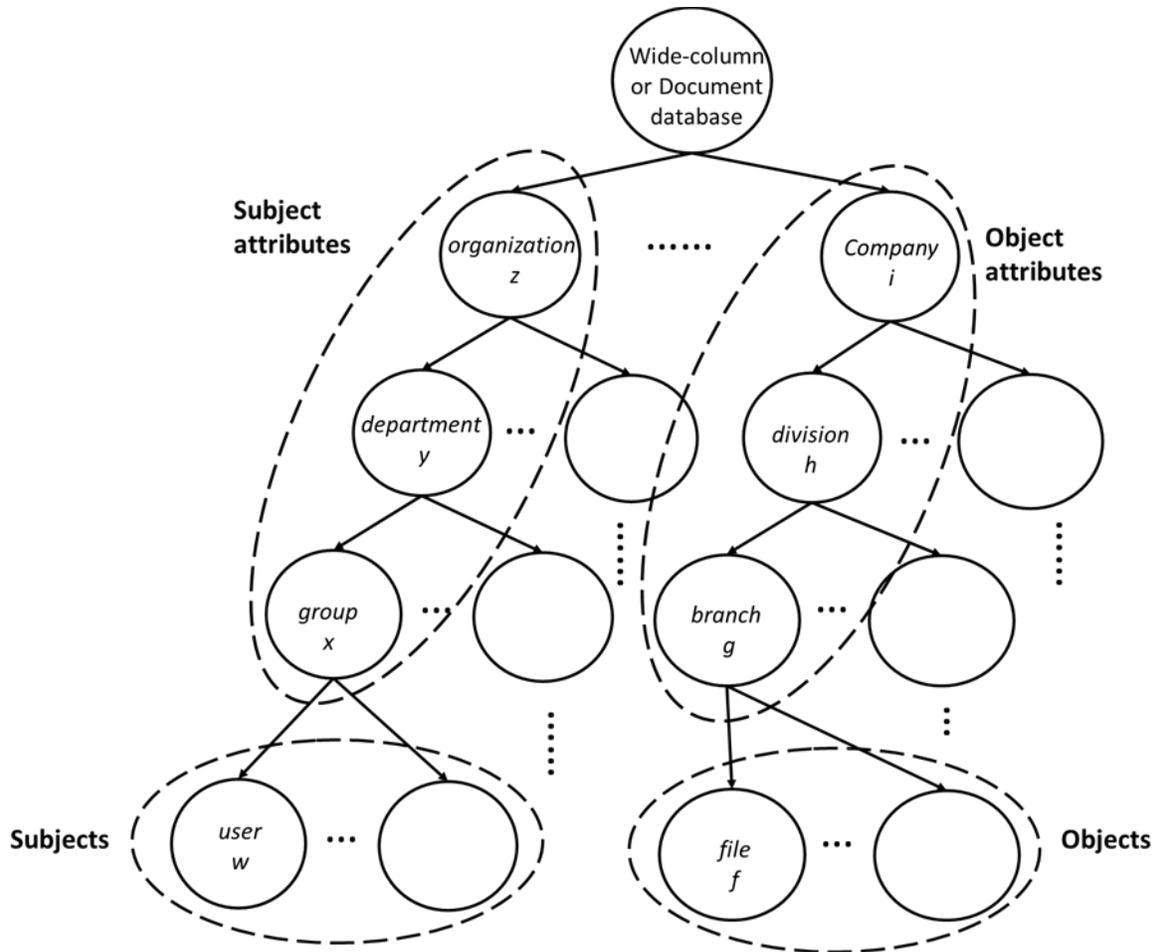
314

Fig. 7. Example of defining subject and object attributes for the key-value model

315 For example, a subject and its attributes can be specified as x (Value *user*), y (Key *group*), and z
316 (Table *department*). Similarly, an object and its attributes can be specified as f (Value *file*), g
317 (Key *branch*), and h (Table *company*). Based on these attributes, an access control policy rule
318 can be created to state that “user x in group y of department z can read file f managed by
319 branch g of company h ” using the relationships of the nodes. It is not necessary to include all
320 levels of nodes except for the leaf to form a policy rule, but at least one node is required for a
321 rule, such as “user x in group y can access objects managed by branch h .”

322 3.2.2. Wide-Column and Document Models

323 In wide-column and document models, the attributes of subjects can be specified through any
324 link from level 2 through level 4 nodes or just one node, as shown in Fig. 8.



325

326

Fig. 8. Example of defining subject and object attributes for wide-column and document models

327

For example, a subject and its attributes can be specified as w (Value *user*), x (Column ID *group*), y (Row *department*), and z (Column family *organization*) for the wide-column model.

328

329

Similarly, a subject and its attributes can be described as w (Value *user*), x (Field *group*), y (Document *department*), and z (Collection *company*) for the document model. An object and its

330

331

attributes can be specified as f (Value *file*), g (Column-ID *branch*), h (Row *division*), and j (Column family *company*) for the wide-column model. For the Document model, an object and

332

333

its attributes can be specified as f (Value *file*), g (Field *branch*), h (Document *division*), and j (Collection *company*). Thus, access control based on attributes for the Wide-column model can

334

335

be represented as follows: “The user w in group x of department y in company z can read file f

336

337

managed by branch g that belongs to division h in company j .” Similar access control rules can

338

339

be specified for the document model. In both models, an access control rule can be specified

340

using only one node (i.e., attribute) without linking to the next level of nodes. This implies that

341

342

the lower-level attributes inherit access permissions from the attributes above them. For

343

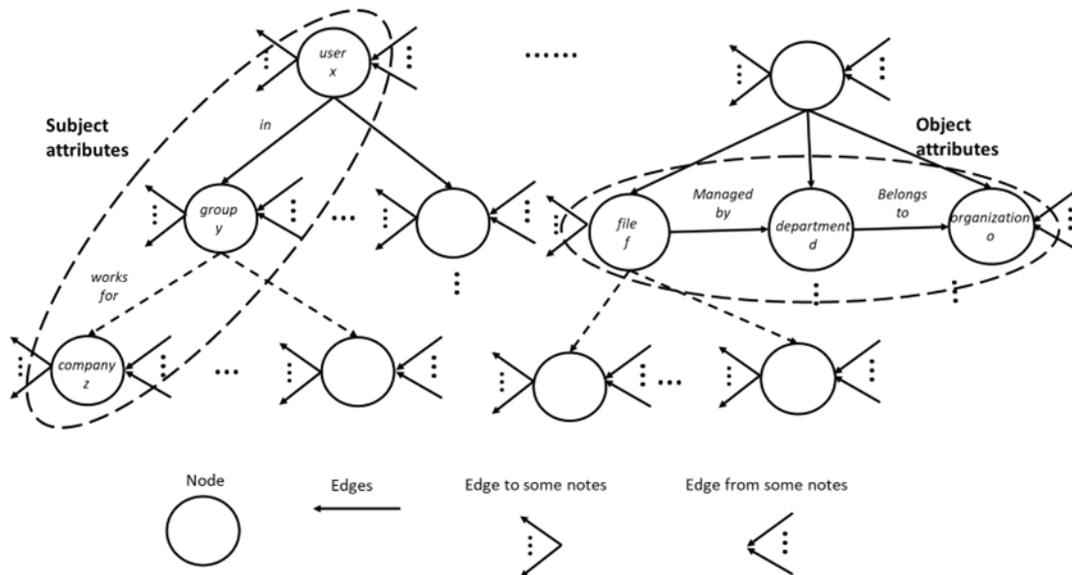
example, “Any users in group x can read objects that belong to division h .”

344 permission} for key-value, wide-column, and document models. Due to the limitations of
345 models based on node hierarchies, accessing data at the finest granularity level of the original
346 data resources becomes a challenge. For example, in wide-column databases, data has to be
347 organized within a collection in Column IDs. In document databases, data is typically stored
348 within Fields. However, there may be situations in which finer-grained access control is
349 required at a level beyond what the existing node structures can support, leading to a need for
350 more attributes in access control rules. In such cases, it becomes necessary to construct tree
351 branches horizontally to accommodate additional attributes. This approach adds complexity to
352 the implementation of the database system and can be challenging to manage effectively. Such
353 a lack of inherent support for fine-grained access control may require developers to find
354 creative solutions to achieve the desired level of access granularity, potentially leading to more
355 complex and less straightforward implementations. In addition to attributes for attribute-based
356 access control, role-based access control (RBAC) can also be implemented by assigning a layer
357 of nodes as roles.

358 3.2.3. Graph Model

359 Attributes in the graph model can be constructed differently since Graph edges do not
360 necessarily have hierarchical relationships between nodes. Instead, edges connected by nodes
361 may form cyclic relationships rather than a tree structure. As shown in Fig. 9, attributes in a
362 graph database can be described through a sequence of edges and/or nodes without
363 hierarchical relationships.

364



365

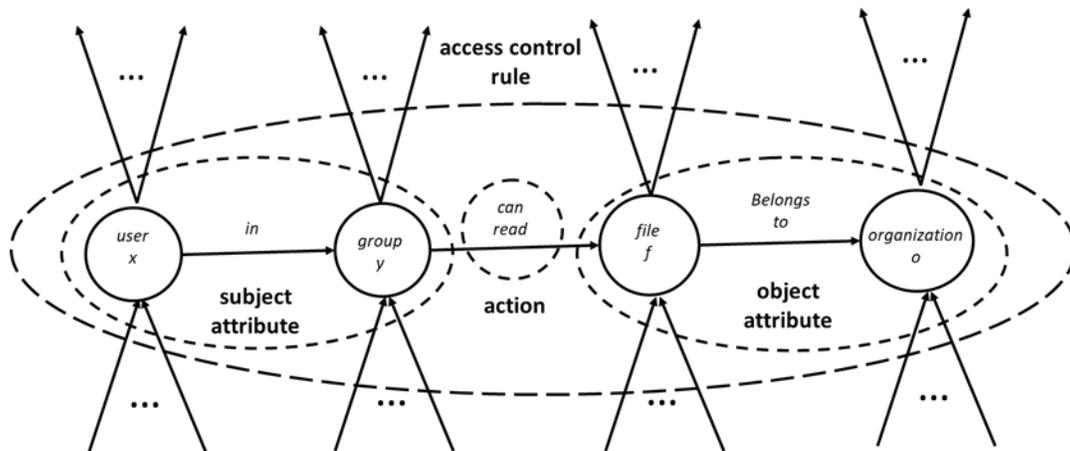
366 **Fig. 9. Example of defining subject and object attributes for a graph model**

367 For example, a subject attribute can be described as user x (node) in (edge) group y (node)
368 works for (edge) company z (node), and an object attribute can be described as file f (node)
369 managed by (edge) department d (node) belongs to (edge) organization o (node). An access

370 control policy based on attributes can then be specified, such as “the user x in group y who
371 works for company z can read file f managed by department d that belongs to organization o ,”
372 or “any users in company z can read objects that belong to organization o ” if only applied to
373 attributes without specific subjects or objects.

374 Edges in the graph model are flexible and can be used to implement either an attribute or a
375 permitted action in an access control rule. If implemented as a permitted action, the access
376 control rule can be directly embedded in the database, treating a sequence of links as the rule.
377 In other words, there is no need to store the access control rule outside of the database. An
378 example graph in Fig. 10 contains the following nodes and edges: user x (node) in (edge
379 directed to group y node) can read (edge directed to file f node).

380



381

382 **Fig. 10. Example of defining subject and object attributes for graph models using edges for permitted actions to**
383 **embed an access control policy**

384 Unlike the key-value, wide-column, and document models, the graph model does not have
385 limitations on the number of edges and nodes connected in a sequence of relationships when
386 constructing attributes for subjects and objects. The graph model allows access to data at the
387 finest possible granularity that matches the original data resource thanks to the unlimited
388 extension of the graph topology. Consequently, when composing access control rules, the
389 graph model offers more flexibility and scalability but also introduces a more complex structure
390 compared to the three other models. Additionally, RBAC can also be applied if any edges or
391 nodes are assigned as roles instead of attributes. This flexibility in representing attributes and
392 roles makes the graph model suitable for various access control scenarios. Appendix A
393 describes an example application of this flexibility for implementations for a federated system
394 environment.

395 **3.3. Access Control Model Implementations**

396 While mandatory access control policies (MAC) (e.g., ABAC [SP 800-162] and RBAC [SF]),
397 discretionary policies (DAC), and various context-based models [NISTIR 7316] can be
398 implemented in NoSQL databases using add-on applications, most current NoSQL

399 implementations enforce authorization at a higher level. Specifically, authorization mechanisms
400 are often implemented at a database level rather than at a more granular level [CF], such as the
401 document collection level or column family level. Some NoSQL systems apply different
402 enforcement mechanisms for different operation modes, such as read-only and read-write
403 permissions being set for users in unshared modes but lacking support for authorization in
404 shared modes [AA]. To support access control models, many research and commercial NoSQL
405 databases propose or implement advanced enforcement mechanisms [CF] that vary from one
406 system to another and may include:

- 407 • Modifying the format (schema) of the NoSQL structures according to the access control
408 model
- 409 • Modifying the query methods according to the access control model [AA]
- 410 • Integrating an access control model in a NoSQL hierarchical level

411 Managing the heterogeneous, non-normalized, schemaless data of NoSQL databases
412 characterized by complex hierarchical structures can make it difficult to handle the dynamic
413 requirement of access control policies, such as conditional control (e.g., a rule regulating “no
414 access to data has a class greater than 10” or “no access for name field is empty”).

415 The graph model’s No-SQL structure of nodes and edges can leverage the pervasive capability
416 of semantic content and the fluency of machine-understandable knowledge of Semantic Web
417 technology. One important application is access control that manages federated resources in
418 federated access control environments, as demonstrated in Appendix A.

419 **4. Considerations of Access Control for NoSQL Systems**

420 NoSQL systems can handle fast-paced agile development, the storage of large volumes of
421 structured and semi-structured data, the requirements for scale-out architecture, and modern
422 application paradigms (e.g., microservices and real-time streaming) [MO]. However, they also
423 face many challenges with regard to access control, such as a lack of proper authentication,
424 encryption, dynamic model support, and fine-grained authorization (as described in Sec. 3.3).
425 Therefore, careful considerations need to be made when implementing access control on
426 NoSQL databases to address these issues effectively.

427 **4.1. Fine-Grained Access Control**

428 Fine-grained access control (FGAC), which determines the scope of data authorized to users,
429 plays a crucial role in access control mechanisms. However, due to the schemaless nature of
430 most NoSQL databases, access control mechanisms are often implemented at a coarse-grained
431 level. Unlike traditional RDBMS, where FGAC allows for the enforcement of access control at
432 the cellular (e.g., row or column) level [FF], granular control is not available in current NoSQL
433 databases. For example, some document NoSQL systems grant access control to the whole
434 database or none [CJ]. This limitation hinders the ability to provide customized data protection
435 levels, which could enhance the usability and expansion of these systems. While additional
436 enforcement mechanisms can be implemented at various levels (e.g., a column family, column,
437 or row), these solutions may not be easily adapted to other NoSQL databases.

438 To address these challenges, various proposed solutions include identifying suitable
439 engineering approaches for encoding policies, defining an enforcement monitor, modifying the
440 format of NoSQL structures, and adapting query methods according to access control models
441 within a NoSQL hierarchical level. Despite ongoing research efforts, the integration of FGAC into
442 NoSQL databases is still a work in progress [AA, FF].

443 **4.2. Security**

444 Due to NoSQL's distributed nature, security becomes a more significant concern compared to
445 center management-oriented RDBMS. From the perspective of data at rest and data in transit
446 (i.e., communications between databases), NoSQL databases generally provide weaker security
447 features when compared to RDBMS.

448 For data at rest, some NoSQL databases may allow users backdoor access to other users' data
449 due to poor logging and log analysis methods [ZA]. In some systems, authorization is applied at
450 a per-database level using a role-based approach, limiting certain roles to specific privileges at
451 different database levels. In contrast, some systems may provide no authorization by default,
452 allowing any action by any user [CJ]. In such cases, an external security enforcement
453 mechanism is essential. For example, some systems utilize metadata and provide management
454 functions based on the database structure implementing access control operations with
455 authorization principles. This allows different applications to implement their own access
456 control.

457 The security of data in transit and communications in horizontally scaled NoSQL databases is a
458 critical issue. Some systems use symmetric key algorithms to encrypted data based on the
459 NoSQL structure (i.e., types). For example, in a key-value system, the value of the key is the
460 encrypted data, including all of the other key values being concatenated and encrypted. When
461 a query is created, the system retrieves data entities and then uses the symmetric key to
462 decrypt the data [ZM]. As a result, the data is protected by encryption even if the key is
463 intercepted while in transit between different NoSQL hosts.

464 Just like RDBMS, NoSQL databases are susceptible to injection attacks, which allow attackers to
465 add malicious data to the NoSQL and cause unavailability and corrupt data. NoSQL injection
466 attacks typically occur when the attack string is parsed, evaluated, or concatenated into a
467 NoSQL API call. Attackers who are familiar with the syntax, data model, and underlying
468 programming language of the target database can design specific exploits, especially in cases
469 where server-side middleware (e.g., JavaScript, PHP) are heavily used to enhance database
470 performance. For example, an internal operator “\$where,” designed to be used like the
471 “where” clause in SQL, can accept sophisticated JavaScript functions to filter data. An attacker
472 can exploit this by passing arbitrary code or commands into the \$where operator as part of the
473 query. The Open Worldwide Application Security Project (OWASP) Test Guide (v4) [OW] plans
474 to include new procedures for testing NoSQL injections to assess NoSQL systems built upon
475 JavaScript and/or PHP engines that may possess similar vulnerabilities [CJ].

476 NoSQL databases are designed to meet the requirements of the analytical world of big data
477 with less emphasis on security during the design stage. Since they do not inherently embed
478 security features in the system itself, developers must impose security mechanisms in the
479 middleware using third-party tools without compromising scalability.

480 **4.3. Query Language**

481 There is currently no standard NoSQL query language in general or for a specific datastore
482 category. Instead, each database adopts its own unique query language. This lack of
483 standardization reduces interoperability among existing systems. For example, it is currently
484 not possible to write even a basic query that can be executed within several different NoSQL
485 systems. Similarly, data portability can be problematic since importing a dataset from one
486 NoSQL database to another often requires preliminary data manipulation activities, even when
487 dealing with the same data type (e.g., JSON objects). The heterogeneity of NoSQL databases
488 and the diversity of their query languages make defining a general FGAC enforcement solution
489 a complex task. Additionally, there is no systematic support for views, as in standard views for
490 RDBMS, which further adds to the challenge of achieving uniformity and standardization across
491 the NoSQL landscape [FF].

492 **4.4. Data Consistency**

493 Data consistency is essential for some access control models, especially dynamic models such as
494 RBAC sections, separation of duty, workflow control, and n-person control [HA]. These models
495 rely on a current and accurate access state maintained by consistent data to make access

496 permission decisions. However, most NoSQL databases are designed using a shared base
497 architecture across distributed commodity servers, which introduces the possibility of inherent
498 data inconsistency among clustering nodes [CJ]. As a result, NoSQL databases do not always
499 guarantee consistent results as not all participating commodity NoSQL servers may be entirely
500 synchronized with other servers that hold the latest information. Additionally, if a single
501 commodity server fails, it can lead to load imbalance among other commodity servers and
502 affect data availability [ZA]. This lack of strong data consistency may explain why NoSQL
503 databases have not been widely adopted to process critical financial transactions. Instead, it is
504 often the responsibility of developers to design applications that can work with the eventual
505 consistency model of NoSQL databases and to carefully weigh the trade-offs between data
506 consistency and performance impacts [CJ].

507 **4.5. Performance**

508 Performance is key in choosing a NoSQL database, particularly when dealing with high volumes
509 of data. However, security — including access control — is often a trade-off that impacts
510 performance. Many NoSQL databases come with default security settings that are either set to
511 none or minimum. Therefore, the most effective way to maintain performance while reducing
512 risk is to deploy these databases in an environment where proper security measures and
513 performance can be implemented and monitored.

514 **4.6. Audit**

515 Most currently distributed monitoring and reporting tools focus on database performance (e.g.,
516 information about the system's running state and connecting clients) with limited support for
517 access auditing [CJ]. Approaches to implementing auditing functions in NoSQL databases may
518 depend on the type of NoSQL database being used, which may have their own specific methods
519 and tools for implementing access auditing and security logging. For example, in wide-column
520 systems, auditing can be enabled on a per-node basis, allowing for maximum audit information
521 by turning on auditing on every node of the system. It is important to consider the specific
522 requirements and capabilities of the NoSQL database being used to ensure effective security
523 logging and monitoring.

524 **4.7. Environment Conditions Control**

525 Properly managing and ensuring situational awareness is crucial for maintaining a secure and
526 efficient NoSQL database environment. The environmental conditions can either be derived
527 from the NoSQL database itself or provided by an independent outside source. If the
528 environment conditions are derived from the NoSQL database itself, considering consistency
529 (Sec. 4.4) and performance (Sec. 4.5) is critical. However, if the environment conditions are
530 provided by an external source, addressing security issues (Sec. 4.2) is essential.

531 **4.8. Support for AI**

532 Formal documents (e.g., laws, statutes, regulations, memorandum, articles, etc.) are typically
533 written in plain natural language (e.g., English) that may contain access control information. An
534 artificial intelligence (AI) natural language processing (NLP) system can be used to render an
535 access control policy by taking natural language documents as input and generating formal
536 access control rules. To achieve this, the natural language contents must first be converted into
537 a formal format. For example, “employees from the product division have to attend a security
538 training course to work on project X” is translated into a formal format for an access control
539 rule: “(user attribute = *employee from product division*) \wedge (operation = *work*) \wedge (resource
540 attribute = *project X*) \wedge (condition = *attended security training*) \rightarrow (permission = *grant*)”
541 connected by Boolean operators.

542 To make this transformation possible, the AI NLP system must:

- 543 a. Recognize if a sentence in the document can form access control rule(s)
- 544 b. Build a dictionary used by the NPL system
- 545 c. Define grammar for formal access control rules¹

546 NoSQLs databases, especially graph systems, can parse and construct hierarchical order
547 relationships between data, which is necessary for the AI analyst work in steps *a* and *b* above.
548 They provide a data structure to support AI NLP systems in rendering access control policy from
549 natural language documents and allow for more efficient and accurate translations of access
550 control information from formal documents to formal access control rules.

551 **4.9. Combine RDBMS**

552 Considering the specific requirements of an application, it may be beneficial to deploy both
553 RDBMS and NoSQL to process different data flows and achieve the optimal combined features
554 of both types of databases. This hybrid approach can help leverage the strengths of each
555 database model while addressing their respective limitations [CJ].

¹ In term of NLP processes, the requirements might go through the following: 1) sentence segmentation: generate sentences (a list of strings); 2) tokenization: generates tokenized sentences (a list of lists of strings); 3) part-of-speech tagging: generate post-tagged sentences (list of lists of tuples); 4) entity recognition: generated chunked sentences (list of trees); and 5) relationship recognition: generate relations (list of tuples).

556 **5. Conclusion**

557 NoSQL database systems offer promising features, such as flexible data models, horizontal
558 scaling, and fast queries that allow for data analysis efficiency, improved system performance,
559 ease of deployment, flexibility/scalability of data management, and users' availability when
560 compared with RDBMS. However, despite these advantages, NoSQL databases lack a
561 mechanism for handling and managing data consistency and maintaining integrity constraints.
562 Additionally, they often have limited support for security at the database level. With an
563 increasing number of people storing sensitive data in NoSQL databases, security issues are
564 becoming critical concerns. Since access control is a fundamental data protection requirement
565 for any database management system, this document discusses access control on NoSQL
566 database systems by illustrating the NoSQL database types along with their support for access
567 control models and describing considerations from the perspective of access control.

568 References

- 569 [AA] Alotaibi AA, Alotaibi RM, Hamza N (2019) Access Control Models in NoSQL Databases:
570 An Overview, *JKAU: Comp. IT. Sci., Vol. 8 No. 1*, pp 1–9. Available at
571 https://marz.kau.edu.sa/Files/320/Researches/72524_45671.pdf
- 572 [AH] Ahmad N (2023) When to Use NoSQL Databases: NoSQL vs. SQL, *ServerWatch*. Available
573 at <https://www.serverwatch.com/guides/when-to-use-nosql>
- 574 [CF] Colombo P, Ferrari E (2021) Evaluating the effects of access control policies within
575 NoSQL, *Future Generation Computer Systems*.
576 <https://doi.org/10.1016/j.future.2020.08.026>
- 577 [CJ] Crawford J (2014) Current Data Security Issues of NoSQL Databases, *Fidelis Cybersecurity*
578 *Volume 17*. Available at [https://silo.tips/download/current-data-security-issues-of-](https://silo.tips/download/current-data-security-issues-of-nosql-databases)
579 [nosql-databases](https://silo.tips/download/current-data-security-issues-of-nosql-databases)
- 580 [FA] Fullstack Academy (2017) RDF Tutorial - An Introduction to the Resource Description
581 Framework. Available at <https://www.youtube.com/watch?v=zeYfT1cNKQg>
- 582 [FF] Colombo P, Ferrari E. (2016) Fine-Grained Access Control Within NoSQL Document-
583 Oriented Datastores, *Data Sci. Eng. 1*, pp 127–138. [https://doi.org/10.1007/s41019-016-](https://doi.org/10.1007/s41019-016-0015-z)
584 [0015-z](https://doi.org/10.1007/s41019-016-0015-z)
- 585 [HI] hiteshreddy2181(2023) Types of NoSQL Databases, *Geeksforgeeks*. Available at
586 <https://www.geeksforgeeks.org/types-of-nosql-databases/>
- 587 [HS] Hu VC, Quirolgico S, Scarfone K (2008) Access Control Policy Composition for Resource
588 Federation Networks Using Semantic Web and Resource Description Framework (RDF),
589 *Proceeding of 2008 International Computer Symposium*, pp 497-502 (v1). Available at
590 https://tsapps.nist.gov/publication/get_pdf.cfm?pub_id=890067
- 591 [MO] MongoDB (2023) What is NoSQL? Available at [https://www.mongodb.com/nosql-](https://www.mongodb.com/nosql-explained)
592 [explained](https://www.mongodb.com/nosql-explained)
- 593 [NISTIR 7316] Hu VC, Ferraiolo DF, Kuhn DR (2006) Assessment of Access Control Systems.
594 (National Institute of Standards and Technology, Gaithersburg, MD), NIST Interagency or
595 Internal Report (IR) 7316. <https://doi.org/10.6028/NIST.IR.7316>
- 596 [NO] NoSQL (2020) Your Ultimate Guide to the Non-Relational Universe! Available at
597 <http://nosql-database.org/>
- 598 [OG] Okman L, Gal-Oz N, Gonen Y, Gudes E, Abramov J (2011) Security
599 Issues_in_NoSQL_Databases, *International Joint Conference of IEEE TrustCom-11/IEEE*
600 *ICES-11/FCST-11*. <https://doi.org/10.1109/TrustCom.2011.70>
- 601 [OW] OWASP (2020) Web Security Testing Guide. Available at [https://owasp.org/www-](https://owasp.org/www-project-web-security-testing-guide)
602 [project-web-security-testing-guide](https://owasp.org/www-project-web-security-testing-guide)
- 603 [SF] Sandhu R, Ferraiolo DF, Kuhn D.R (2000) The NIST Model for Role Based Access Control:
604 Toward a Unified Standard, *5th ACM Workshop Role-Based Access Control*. pp 47–63.
605 <https://doi.org/10.1145/344287.344301>
- 606 [SL] Simplilearn (2003) NoSQL Tutorial For Beginners | A Complete Guide To NoSQL
607 Databases. Available at <https://www.youtube.com/watch?v=aUPVpiYiLcC>
- 608 [SP 800-162] Hu VC, Ferraiolo DF, Kuhn DR, Schnitzer A, Sandlin K, Miller R, Scarfone KA (2014)
609 Guide to Attribute Based Access Control (ABAC) Definition and Considerations. (National
610 Institute of Standards and Technology, Gaithersburg, MD), NIST Special Publication (SP)

- 611 800-162, Includes updates as of August 02, 2019. <https://doi.org/10.6028/NIST.SP.800->
612 [162](https://doi.org/10.6028/NIST.SP.800-162)
- 613 [WT] WeAreTechies (2022) NoSQL database | key value | document | wide column | graph
614 stores. Available at <https://www.youtube.com/watch?v=zG6CHYCx6ag>
- 615 [ZA] Zaki AK (2014) NoSQL DATABASES: NEW MILLENNIUM DATABASE FOR BIG DATA, BIG
616 USERS, CLOUD COMPUTING AND ITS SECURITY CHALLENGES, *International Journal of*
617 *Research in Engineering and Technology*. <https://doi.org/10.15623/ijret.2014.0315080>
- 618 [ZM] Zaki AK, M Indiramma (2015) A Novel Redis Security Extension for NoSQL Database
619 Using Authentication and Encryption, *2015 IEEE International Conference on Electrical,*
620 *Computer and Communication Technologies (ICECCT)*.
621 <https://doi.org/10.1109/ICECCT.2015.7226101>.
622

623 **Appendix A.**

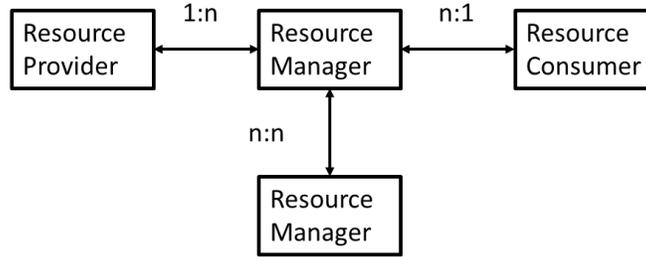
624 **A.1. Federation of Access Control**

625 The availability of pervasive information will be greatly facilitated by the increase of globally
626 distributed and interconnected information services. To support this global architecture,
627 services that reside in groups of local networks (i.e., federations) interact with services that
628 reside in other federations. All member federations form a federated network. To achieve
629 networking in a global-computing framework, it is necessary to facilitate seamless access to
630 federated services through inter-federation resource sharing and inter-trust between limited
631 numbers of participating members of the global federation. However, the management of
632 access control on a multi-organization global environment does not scale well. Since the shared
633 resources of a federation are available both locally and conditionally globally, both the local and
634 global access control policies are integrated under one static access control system so as to not
635 violate the principles of the reference monitor. Therefore, it is challenging to 1) specify access
636 control rules that manage the dynamic trust relations among federated parties, 2) separate
637 local resource access control policy from global (federation) policy and risk the possible leaking
638 of authorization, and 3) share the access control profile among federated members providing
639 similar services.

640 The requirements for the interaction between global (or federation) and local access control
641 policies are complex because most access control mechanisms and models are not flexible
642 enough to arbitrarily combine and compose access control policies. Moreover, federated
643 resources are distributed and shared by interoperating between three services:

- 644 1. Resource providers (RPs) store information for sharing with federated members. The
645 information is managed locally by the resource contributors or administrators of the RP.
646 The availability and integrity of the resource is the central operation goal.
- 647 2. Resource managers (RMs) are responsible for locating the resources in response to the
648 access request from a resource consumer. The security and accessibility of
649 communications between the RPs and their connected resource consumers are the
650 primary concerns of an RM.
- 651 3. Resource consumers (RCs) are client applications that accept user requests for resources
652 and forward those requests to an RM. Ideally, an RP should only have to communicate
653 with one RM because the dissemination of shared resources is achieved by the RM. Only
654 one connection between an RM and an RC is expected as well because the discovery of
655 resource locations should be done by an RM, as illustrated in Fig. 11.

656



657

658

Fig. 11. Generic resource federation model

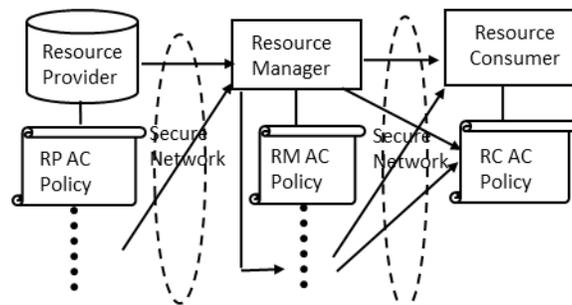
659 In reality, a federated community may be networked in a variety of architectures. The three
660 basic services may be incorporated or simplified such that more than one service is managed or
661 hosted in one physical system. However, these three services and their connections are
662 assumed to be essential for any resource-sharing federation, and the resource-sharing
663 protocols between them are composed by interlacing the following scenarios:

- 664 • Scenario 1: The information request from an RC is sent to an RM and then directly
665 relayed to an RP without passing the request to other RMs or RPs.
- 666 • Scenario 2: A resource query cannot be satisfied by the connected RP, so the RM must
667 collect and consolidate the partial results returned from more than one RP.
- 668 • Scenario 3: An RM does not have a direct or static connection to any RP that can provide
669 the information as requested, so the resource discovery protocols need to be invoked to
670 exchange information with other RMs that may have connections to other RPs with
671 locations for the resources.

672 A.2. Graph NoSQL Implementation for AC Policies

673 To support accessibility and maintain the integrity of resource-sharing, access control policies
674 between the three services are required such that a service has its own policy for the
675 federation. Fig. 12 illustrates a generic scheme for a resource federation network and AC
676 policies associated with each of the services.

677



678

679

Fig. 12. Resource federation scheme

680 In addition to security between services in the lower-level communication mechanism (e.g.,
681 through a PKI infrastructure), support for the federation according to the access control policies

682 posted by the services requires access control functions to be implemented. These functions
683 manage and manipulate the types of enforcement rules listed below. Here, it is assumed that
684 the policy for each service is maintained locally by the administration of the service.

- 685 • RP
 - 686 (p1) share (or conditional) rules
 - 687 (p2) non-share (or conditional) rules
- 688 • RM
 - 689 (m1) list of trusted RPs
 - 690 (m2) list of trusted RMs
 - 691 (m3) credibility rules for m1 and m2 (e.g., RP *A* has more credential than RP *B*)
 - 692 (m4) priority rules for m1 and m2 (e.g., RP *A* can be replaced by RP *B* — *A* “is a
693 replacement” of *B*)
 - 694 (m5) reference rules (information from RP *A* is composed of information from
695 RPs *B*, *C*, and *D* — *A* “should be supplemented by” *B*, *C*, and *D*)
 - 696 (m6) mediation rules (information from RP *A* cannot conflict with information
697 from RP *B*)
- 698 • RC
 - 699 (c1) reference rules (similar to the reference information in RM except at the
700 application level such as logic operations (AND, OR, XOR) between collected
701 information)
 - 702 (c2) mediation rules (similar to the mediation information in RM except at the
703 application level, such as data *a* from RM *X* cannot conflict with data *b* from RM
704 *Y*)
 - 705 (c3) constraint rules (for RMs, such as no information older than 10 days can be
706 trusted)

707 Rules p1, p2, m1, m2, and c3 contain resource availability information, while m3, m4, m5, m6,
708 c1, and c2 contain information for trust management. Each rule is an access control policy
709 assertion enforced upon two of the RPs, RMs, or RCs. Such a formal relationship can be
710 annotated as members of a set that contains the binary relationships that the rule set is
711 enforced upon: Rule $x = \{ \dots (S_x, S_y) \dots \}$, where S_x service is related to service S_y by the
712 enforcement of Rule x . For example, *Credential* = $\{ \dots (S_1, S_2) \dots \}$ says that the resource from RP
713 S_1 has more credentials than RP S_2 and *Replace* = $\{ \dots (S_1, S_2) \dots \}$ says that the resource from RP
714 S_1 should be requested if RP S_2 is not available. Thus, by conventional set operations, an access
715 control trust management policy can be composed and combined through the Boolean or
716 closure properties of the sets of trust management rules. A trust management rule can be
717 expressed by a relationship pair (S_x, S_y) in a set that contains the type of rule such that the pair
718 in the set Rule x , which are subject S_x and predicate Rule x , and object S_y [HS] can be supported

719 by the Graph NoSQL in the form of node S_x edge rule x node S_y for an access control policy rule.
720 For example, *Replace* = {... (S_x , S_y) ...} is translated into S_x and can replace S_y of a triple in a
721 graph NoSQL model.