

# 1 Big Ten Academic Alliance Provisioning

## 2 Cookbook

- 3 • [1. Introduction](#)
- 4 • [2. Problem Statement](#)
- 5 • [3. Basic Concepts](#)
- 6 • [4. Identity Provisioning](#)
- 7 • [5. Identity Lifecycle](#)
- 8 • [6. Passwords, Multi-Factor Authentication, and Provisioning](#)
- 9 • [7. Service Provisioning](#)
- 10 • [8. Target Directory Provisioning](#)
- 11 • [9. Authorization](#)
- 12 • [10. Assuring Provisioned Authorization Is Correct](#)
- 13 • [11. Product Lifecycle](#)

## 14 1. Introduction

15 Provisioning is a core function of any identity and access management system. Whether you're  
16 provisioning identities for new students, MFA devices for staff members, e-mail boxes for  
17 faculty, or secure cloud storage for a research group, there's lots to think about. Each kind of  
18 item that you provision has its own set of considerations and pitfalls. It's important to have  
19 standard practices and policies around your provisioning logic so that things flow smoothly and  
20 everyone knows what to expect.

21 Some provisioning practices, like deciding at what point your users have their e-mail accounts  
22 created, is easy to adjust as you go along. Other things like defining a useful and extensible set of  
23 roles to determine what users can access is much harder to adjust once in place. Regardless of  
24 what part of your identity infrastructure you're laying out, careful thought and planning up front  
25 can save a lot of headaches later.

26 Another aspect that goes hand-in-hand with provisioning is auditing. Now more than any time in  
27 the past, auditing of systems for compliance with legal regulations is a huge priority. Sometimes,  
28 however, it's just as important to audit systems to make sure that local institution policies are  
29 being properly enforced. Doing this manually, though, is next to impossible. Things work best  
30 when your auditing and reporting mechanisms are built into and work alongside your  
31 provisioning systems.

32 This cookbook seeks to jumpstart your planning with easy-to-follow advice and trusted practices  
33 that others have found work well. It's written from the perspective of provisioning in a higher-ed  
34 environment, but the advice applies in other settings, too.

35 We've created this cookbook with a simple, straightforward format. The first section defines  
36 concepts and terminology that's helpful to understand before diving into the specific recipes. The

37 sections that follow each cover a specific area of provisioning and deprovisioning. Each section  
38 starts out with an overview of that concept, followed by the specific guidance.

39 We offer our guidance in “DO”/”DON'T” form: each item starts with “DO” or “DON'T”. Items  
40 that aren't necessarily a definite DO or DON'T but that are worth keeping in mind start with  
41 CONSIDER.

42 There's no such thing as one size fits all in provisioning. Much of the architecture depends on  
43 your organization's needs, culture, and size. Institutional policy will be required (and, perhaps,  
44 evolved) to shape IAM business processes and system architecture. While this cookbook can't  
45 tell you exactly how to architect everything, it explains best practices and the concepts needed to  
46 get the job done. Our goal is to help you understand what you need to know so that you can  
47 make your own informed decisions appropriate for your organization.

## 48 **2. Problem Statement**

49 When it comes to identity management, higher-ed institutions are complicated places. A single  
50 person can have multiple roles all at once: student, alumnus, staff member, and parent of a  
51 student, for example. That same person can disappear from any active affiliations with the  
52 institution, only to reappear several years later with a whole new set of roles. Managing the  
53 digital identity, credentials, services, and roles of that individual is no small task.

54 Consider further that higher-ed institutions on-board hundreds or thousands of new users each  
55 year as new students enroll. This means that, not only do the policies and practices that a school  
56 has in place need to be well-defined, but the systems that enforce them need to be scalable.  
57 Institutional roles need to be able to express affiliations concisely and accurately so that systems  
58 can easily perform the work to create all of the resources that those students will need to learn  
59 and participate.

60 Then, every spring, hundreds or thousands of students will graduate. Those same systems will  
61 need to deprovision resources. Sometimes, some of the resources will be maintained if the school  
62 offers continued services to alumni. Other pieces of information will be maintained to create an  
63 auditing trail or to allow easier on-boarding if one of those students comes back later for another  
64 degree or a job.

65 In addition to the usual identity lifecycle of students, faculty, and staff, most institutions see a  
66 large number of guest or sponsored accounts. Visiting scholars, contractors, K-12 students  
67 attending summer camps – they all need a digital identity while they're interacting with the  
68 institution. Unlike longer-term engagements like that of a faculty member or student, though,  
69 knowing when to provision and deprovision these users can be very challenging. Regardless of  
70 the user population, how do you know when specific users need to have access and what services  
71 they need to access?

72 With so much data flowing, being able to audit the full system and make sure that nobody  
73 is being incorrectly given access is crucial. A good provisioning system also has good auditing  
74 capabilities and good reporting functionality.

75 Once all that's done, you shouldn't overlook deprovisioning. Especially in these days of cloud-  
76 based services and per-user licensing, failure to deprovision a set of users who no longer need  
77 service can become costly over time. When you deprovision and what you do with the data in  
78 identities and services before they're provisioned depends on your exact situation. For some  
79 cloud-based services, you might even choose to let your users convert their accounts to personal  
80 accounts that they can continue to own but that are no longer affiliated with your institution. The  
81 important thing is that these decisions are made up front.

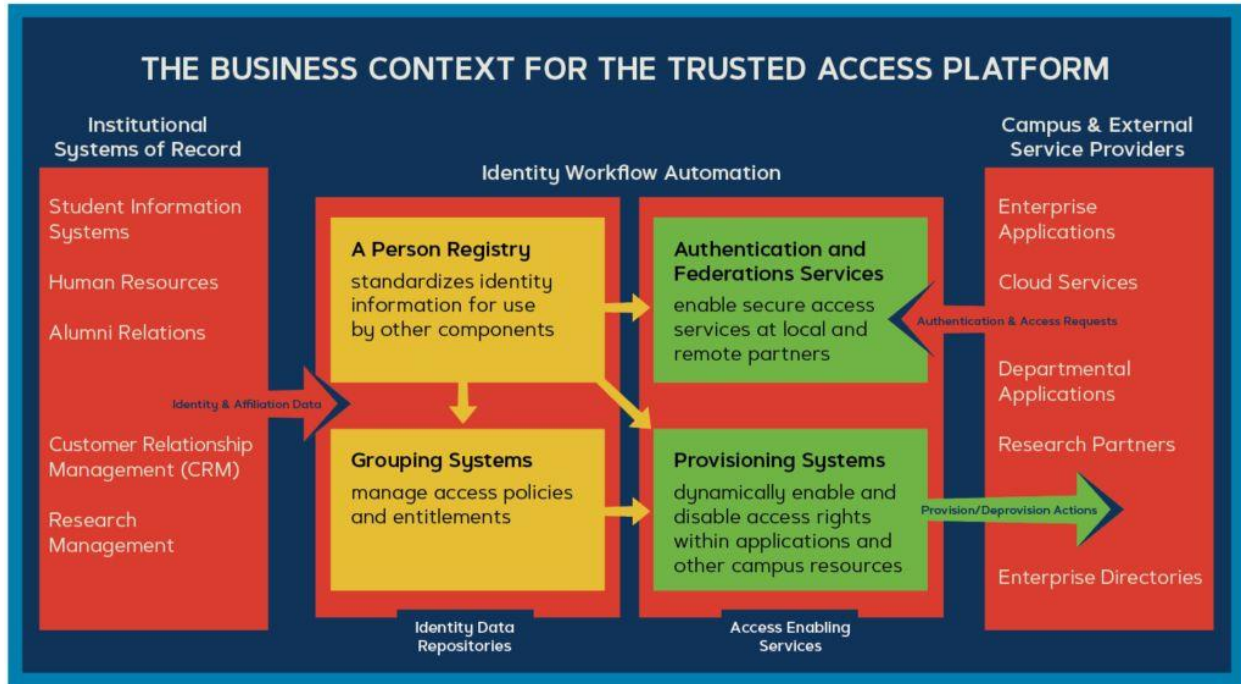
82 Obviously, the big picture can quickly become overwhelming. There are lots of things to  
83 consider and lots of decisions to make, each one that will significantly impact your users  
84 directly. We'll tackle this challenge one piece at a time.

## 85 **2.1. The IAM Business Function and the IAM System**

86 Addressing these issues is generally addressed by an organizational unit within the institution,  
87 usually central IT. In partnership with other units, such as Payroll, Registrar, Library, research  
88 projects, student organizations, *etc.*, the IAM unit is responsible for the business functions that  
89 maintain the institution's unified repository of its community: students, staff, faculty, library  
90 patrons, researchers, members of the chess club, *etc.* The IAM unit is also responsible for the  
91 controlled dissemination of selected information from that repository to online resources to  
92 facilitate access control and personalization.

93 In order to accomplish this for the constantly changing hundreds of thousand members of the  
94 community, IAM Systems are used to automate the lion's share of these processes. The following  
95 diagram shows the interrelationships among affected organizational units and components within  
96 the IAM:

97

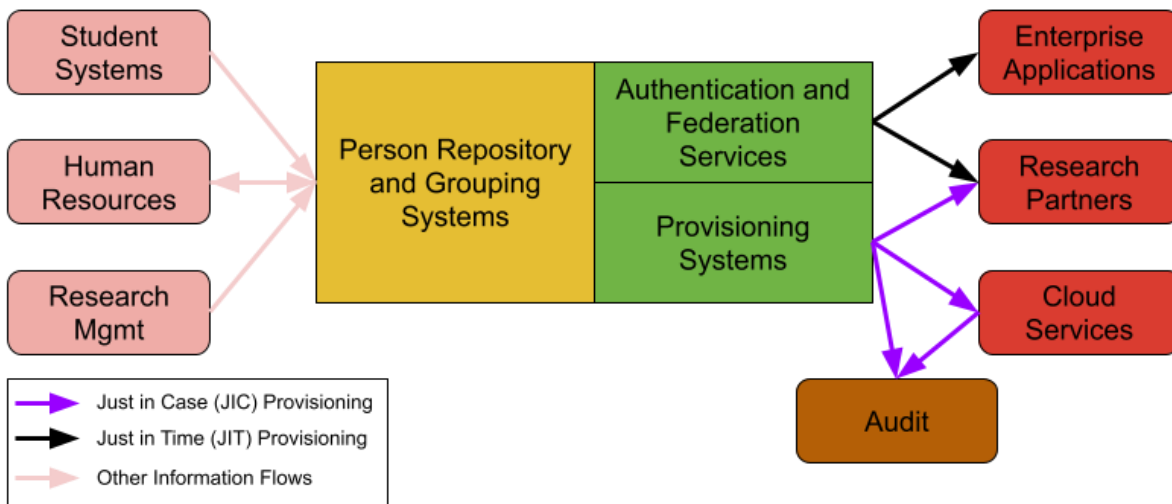


98

99

100 This document concerns itself primarily with the provisioning of data out of the IAM system, but  
 101 also with flows from Institutional Systems of Record, as highlighted in the following diagram:

102



103

104

105 Provisioning flows out of the IAM system to service providers facilitate access control decisions  
106 and personalization by online resources, providing those resources have the information they  
107 require to make those decisions.

108 The provisioning flows out of the IAM System can be classified as:

- 109 • "Just in Case" (JIC), flows that occur in case one of the recipient services needs the  
110 information at some time in the future, or
- 111 • "Just in Time" (JIT), flows that occur at the moment the information is needed by the  
112 recipient service, generally as part of an authentication event.

113 For example, when identity information is transferred at the start of a session with a Service  
114 Provider, it is JIT provisioning, generally accomplished by the authentication system (*e.g.*, as a  
115 SAML assertion or a collection of OIDC claims). When identity information is transferred as the  
116 result of some other event (*e.g.*, a change generated by a System of Record), it is JIC  
117 provisioning.

118 Audit systems also receive provisioned information from the IAM System to assure that access  
119 control policies are consistent with reality, Note that audit systems may also retrieve information  
120 from the service providers to enable comparisons between the service provider and the IAM  
121 System to find discrepancies.

122 While only secondarily within the scope of this document, information flows between Systems  
123 of Record and the IAM System may be JIT or JIC flows, driven by changes in state within the  
124 systems of record, to keep the IAM System's Person Repository and Grouping Systems current  
125 with data about community members from the Systems of Record operated by partner  
126 organizational units, such as Registrar and Human Resources. Note that these flows may be bi-  
127 directional, when the IAM System holds information of use to a System of Record. In this case,  
128 such flows are driven by changes in state within the IAM System. Examples of such information  
129 include user names and directory information.

## 130 **3. Basic Concepts**

### 131 **3.1. Identity and Subject**

132 In the context of the field of Identity and Access Management (IAM), *identity* refers to the set of  
133 information that pertains to a *subject*. A subject is typically a person but may be an institution, a  
134 well-known service, *etc.*; it is the thing with which the identity is associated. The information  
135 that comprises an identity may include identifiers, group memberships, entitlements, roles,  
136 names, and other characteristics.

### 137 **3.2. Identifiers**

138 *Identifiers* uniquely distinguish an identity within a given *domain*. Typical domains might be a  
139 university or other organization, or a federation of universities and other organizations. Examples  
140 of university identifiers might be called netID, login name, principle, or username and have  
141 values like `susan.smith`, `jrrtolkien`, or `a983k50299`. Electronic mail addresses may also be  
142 used as identifiers, although the practice is discouraged, as electronic mail addresses may  
143 change, may be reassigned to other people, may be used by more than one person, *etc.* In  
144 summary, it is usually undesirable to manage electronic mail addresses in a manner  
145 commensurate with their use as user identifiers.

146 *Federated identifiers*, such as the [SAML General Purpose Subject Identifier](#) (subject-id) or the  
147 eduPerson Principle Name (ePPN), are unique within their federation and are typically structured  
148 by appending a unique institution identifier to a person identifier that is unique within the  
149 institution. For example, if John Doe's identifier within the University of Illinois is `jdoue`, then his  
150 federated identifier might be `jdoue@illinois.edu`.

151 (See [Understanding Federated User Identifiers](#) for more information.)

### 152 **3.2.1. Uses for Identifiers**

153 Identifiers are used in multiple ways:

- 154 • Internally to link identities among multiple services and directories
- 155 • Externally with federated vendors
- 156 • By the user as part of an authentication event
- 157 • To retrieve attributes or other information associated with an identity

### 158 **3.2.2. Types of Identifiers**

159 Identifiers have multiple characteristics.

- 160 • They may be usable by any service provider, or they may be *pair-wise*, generated  
161 uniquely by the identity provider to be usable by only a single service provider.
- 162 • They may be *human-readable* to be easily recognized when read, or they may be *opaque*.
- 163 • They may be *short-term*, for example for the lifetime of a session, or they may remain  
164 valid for longer than a single session, always unique to the same identity. How long an  
165 identifier remains valid is also an important characteristic. Identifiers are considered to be  
166 *stable* or *immutable* if they remain valid over a very long period of time (*e.g.*, longer than  
167 the existence of the identity within its domain).
- 168 • An identifier may be re-assignable to a different identity, or it may not.

169 It should be noted that identifiers that are used to control secure access to resources or services  
170 should be stable and not re-assignable to mitigate unintended loss of access by an otherwise  
171 authorized subject, as well as unauthorized access by a subject who receives a re-assigned  
172 identifier. Also, unless there is a need for multiple services to coordinate (*i.e.*, track) their  
173 offerings for the same subject, pair-wise, opaque identifiers are generally preferred to enhance  
174 privacy.

175 More information regarding *federated* identifiers can be found in [Understanding Federated User](#)  
176 [Identifiers](#) and [Choosing the Right Federated User Identifier](#).

### 177 **3.3. Affiliations, Roles, Groups, and Other Attributes**

178 There is a rich set of information that can be used to determine a person's permissions to use  
179 services and resources.

- 180 • *Affiliation* identifies how or why a person is a member of the institution's community.  
181 Examples might be student, staff, or faculty. A person can have multiple affiliations  
182 simultaneously.
- 183 • *Role* identifies what a person does for or with the institution. Examples might be  
184 instructor, principle investigator, or IAM administrator.
- 185 • *Group* is a collection of people that share a common characteristic. Examples might be  
186 people with the instructor role, members of the chess club, or Physics 101 students.
- 187 • *Attribute* is a broad term for just about any information associated with a person.  
188 Examples are telephone number, electronic mail address, or job title. Not all attributes are  
189 useful for assignment of permissions, but some are.

### 190 **3.4. Provisioning Models**

191 There are two types of provisioning, *Just in Time* and *Just in Case*.

- 192 • *Just in Time* (JIT) provisioning occurs when identity information is provisioned at the  
193 time it is required. Just in Time provisioning is usually implemented as part of (or  
194 immediately after) authentication for a session.
- 195 • *Just in Case* (JIC) must occur before the identity information is required. Just in Case  
196 provisioning occurs as the result of some external event, such as enrollment of a student  
197 or assignment of a new role to an employee, or as regularly scheduled processing, such as  
198 overnight reconciliation between the IAM and a service.

## 199 **4. Identity Provisioning**

200 At universities, it is typical for Admissions and Registrar to determine whether an individual's  
201 affiliation with the institution is as a student, Human Resources determines whether that  
202 individual's affiliation is an employee, and other organizational units to determine other  
203 affiliations. It is also possible for an individual to have multiple affiliations. For this reason the  
204 IAM function needs to form partnerships with those other organizational units to obtain the data  
205 of who, currently, are students, employees, *etc.* Operationally, this is implemented by identity  
206 provisioning interfaces between each of your partners' Systems of Record and the identity  
207 registry. An identity matching algorithm is then used to determine when an individual having  
208 multiple affiliations receives records from multiple sources.

### 209 **4.1. Identity Matching**

210 Many organizations, particularly universities, have multiple source systems representing multiple  
211 segments of the population. These include staff, faculty, students, parents, physicians, friends of  
212 the library, *etc.*, to name but a few. These segments, however, are not mutually exclusive. A  
213 student can be a staff member, for example. For this reason, records retrieved from identity  
214 source systems must pass through an Identity Matching algorithm to ensure that a) one person  
215 does not appear to be two people in your IAM system, and b) two people do not appear to be one  
216 person.

#### 217 **4.1.1. Do: Use a scoring system that separates new identities into positive** 218 **matches, unmatched, and potential matches.**

219 When creating a new electronic identity, it's important to see if the individual for whom  
220 you're creating already has an electronic identity at your organization. Not checking  
221 carefully can lead to someone getting two identities or, worse yet, incorrectly assuming  
222 that two people are the same person. A scoring system works well to solve this challenge.  
223 Compare common elements of the new identity with those already in the system such as  
224 date of birth, gender, name components, and whatever else you have available. An exact  
225 match might get a high score unless it's a common name: there could be two individuals  
226 named Elizabeth Jones both born on March 1, 1980. So, common names might score  
227 lower. Use all of the information you have, score each comparison intelligently, and decide  
228 how high of a score counts as a match. Where practical, engage directly with the person  
229 being processed through self-service processes to verify knowledge of data you already  
230 have, such as student ID, course information and grades, employee ID, *etc.*

231  
232 You may want to consider categorizing certain data elements used in the matching as High  
233 Assurance, Medium Assurance and Low Assurance. For example :

- 234 •
  - 235 ○ High Assurance: IDs from systems of record, the registry's unique ID, SSN,
  - 236 driver's license, state ID, passport, NetID
  - 237 ○ Medium Assurance: name, date of birth
  - 238 ○ Low Assurance: email, address, phone

239  
240 These are, of course, only examples; you will need to determine each of your data  
241 elements' contribution to the matching algorithm, based on the characteristics of your  
242 population and the quality of the data you receive from the various sources.

#### 243 **4.1.2. Do: Establish a process for putting potential matches in a suspense** 244 **group for manual review and reconciliation.**

245 It can be easy, using the above system, to know what to do if you definitely have a match  
246 or if you definitely don't have a match: high and low scores can be handled automatically.  
247 But what about scores in the middle? This is where a good workflow process comes into  
248 play. Create a queue for the "maybe" cases, and assign staff to review it manually to  
249 research if a new identity is the same as an existing one or really a different individual.



250 **4.1.3. Do: Establish a process to correct mistakes made either in the**  
251 **automated or manual identity matching processes.**

252 Mistakes are unavoidable. You may not detect them until long after identity assignments  
253 have been made.

254 **4.1.4. Consider: Products and processes that facilitate matching such as**  
255 **phonic and fuzzy name matching, and address standardization.**

256 Non-exact names might still be a match: a new identity for Beth Jones might be for the  
257 same individual as the existing identity belonging to Elizabeth Jones. Also consider that  
258 data entry can contain typos and a name could have an error in it because someone entered  
259 it wrong. Using software that can do non-exact comparisons of names or look at phonic  
260 comparisons will help immensely in identity comparisons. Similarly, changing addresses  
261 to a standardized format where geographically available can give another clue if someone  
262 is a match.

263 If you can engage the actual person whose new data is being processed to participate in a  
264 self-service to match with the existing data you already have on the existing record such as  
265 Student Id, Course Information and grade, Employee Id. Something similar to credit card  
266 report validation.

267 **4.2. Institutional Username Assignment**

268 Your users' institutional usernames (also known as *netid*, *login id*, or *userid*) are the unique  
269 identifiers they use for your single sign-on system, facilitating access for many resources. It may  
270 also be used for institutional electronic mail addresses, although many institutions do not do this,  
271 as it reveals part of what is needed to authenticate. It is also the case that service providers may  
272 or may not use the institutional username 1) because the institutional username is not in a format  
273 that the service provider can use, or 2) as with electronic mail addresses, it may reveal  
274 information that is better kept secret.

275 **4.2.1. Consider: Self selection vs. assigned usernames.**

276 There are good reasons to automatically generate usernames for your users. There are just  
277 as many good reasons for letting them choose their own usernames. It's up to each  
278 institution to do what's best for them, but consider the pros and cons carefully. Letting  
279 users pick their own username generally leads to nicer looking and easier to remember  
280 usernames than an algorithm can select, but there's always the risk that a freshman might  
281 pick a questionable username. When that freshman later discovers that "it seemed like a  
282 good idea at the time", you're opening yourself up to more username changes later. Some  
283 institutions have had no problems with students choosing their own usernames. Others  
284 choose to generate usernames for students while letting faculty and staff choose their  
285 usernames. Regardless, if you're going to assign usernames automatically, think carefully  
286 about your algorithm for username generation to try and maximize the use of your

287 username namespace with reasonable looking usernames. Nobody wants a username that's  
288 an unpronounceable string of consonants followed by a random-looking string of numbers,  
289 especially if it's also part of their institutional email address. It's also possible that a  
290 generated username could have cultural sensitivities or be otherwise offensive to an  
291 individual. So, if you do choose to use assigned usernames, consider having a documented  
292 exception request process and using a vanity/alias for the email address.

#### 293 **4.2.2. Do: Allow changes for good reason.**

294 Though generally painful finding all of the provisioning and authorization implications in  
295 downstream systems, username changes are inevitable. Marriages, adoptions, and other  
296 life events that result in name changes are often good reasons to change a person's  
297 username. Ensure you have a process that preserves the previous username so it cannot be  
298 reused by another user, and keep as much information as is practical on downstream  
299 systems that may have a copy of that username provisioned to them.

#### 300 **4.2.3. Don't: Reassign a username to a different person.**

301 While it may seem tempting, especially in a space where all the good usernames are  
302 already taken, to give the username of a long-graduated or long-retired user to a new  
303 student, there are very good reasons not to do that. There's the obvious possibility that the  
304 original owner of that username might return. There are also resources and permissions  
305 that might still belong to that username which you wouldn't want to transfer inadvertently  
306 to a new person. Non-reassigned usernames are of particular importance when dealing  
307 with cloud services that might still have data belonging to the existing username. In  
308 general, assign a new username instead of reusing a dormant one.

#### 309 **4.2.4. Do: Make sure the namespace is large enough to not run out for many** 310 **years.**

311 This is common sense, especially given the practice of not reassigning usernames. The  
312 question is how large is large enough? That depends on the size of your institution and the  
313 number of users you're on-boarding. There's no magic formula for how many characters  
314 long you should allow for the shortest and longest usernames and what characters (letters,  
315 numbers, certain symbols) you should include. You just need to find the right balance for  
316 your population. Naturally, usernames that are too long become unwieldy. Usernames that  
317 are too short lead to indecipherable strings. The right length and, if appropriate, the right  
318 username generation algorithm can go a long way, and your future self will thank you.  
319 Potential strategies that can help include:

- 320 •
- 321 ○ including middle initials
- 322 ○ appending sequence numbers when there are conflicts (Beware, though, that some  
323 numbers have potentially negative cultural significance.)
- 324 ○ enabling self selection of usernames

325 **4.2.5. Do: Check for a user's existing identity at the institution before**  
326 **assigning a new username.**

327 There's no better way to burn through your username namespace too fast than to use it up  
328 assigning multiple usernames to the same individual. With institutions assigning  
329 usernames for guests, summer program participants, faculty, staff, students and many other  
330 populations, it's extremely common for a given user to pass through multiple affiliations,  
331 perhaps leaving and returning between each one. It's not only kinder to your username  
332 namespace to check for existing usernames before assigning a new one, but it's also kind  
333 to the user to give them back the username that they had last time they were part of the  
334 institution. See the information on identity matching for advice on how to best detect when  
335 a new user is the same person as an already existing identity.

336 **4.2.6. Do: Have a list of usernames that should not be used.**

337 Reasons for inclusion on the list include being misleading (*e.g.*, a username of  
338 "President"), being culturally inappropriate, or violating a trademark. The list would apply  
339 to both usernames that are assigned or selected by users.

340 **4.3. Identifiers for Services and Target Directories**

341 Some of your service providers will require you to provision a username to meet their specific  
342 requirements. Here is some advice.

343 **4.3.1. Do: Maintain an opaque identifier that won't change over an entity's**  
344 **life cycle.**

345 Users will want or need to change their usernames, and for many good, and some not-so-  
346 good, reasons. It doesn't take long to discover that a username as a stable identifier is  
347 anything but stable. If a user's identity is going to stay tied together across multiple  
348 directories and services, you'll need something better than a username to associate  
349 between those directories and services. The days of social security numbers for this  
350 purpose are long gone. Some institutions are comfortable using a university ID number for  
351 this purpose. Others choose to generate a random, reasonably long, unique string for this  
352 purpose when the user's electronic identity is created. Such an identifier doesn't need to be  
353 human readable or ever displayed to the user. The user doesn't even need to know that the  
354 identifier exists. It's purely for linking between systems.

355 **4.3.2. Do: Consider an external identifier different than the internal**  
356 **identifier used by in-house applications.**

357 There are good reasons to have multiple user identifiers for a given user, each one with its  
358 own purpose. While it's important to not create too many as that becomes challenging to  
359 manage, there are good reasons to create new ones. One such case is for external services  
360 and systems. Especially since these types of systems are more out of the institution's

361 control, there's always the possibility that a bad actor could get your users' data. With the  
362 same stable identifier used everywhere, activities and data could easily be correlated  
363 between multiple systems. The best way to avoid this is to have a different identifier for a  
364 user for each service that they access. Rather than creating a new identifier for each service  
365 when the user is provisioned into it, consider an algorithm that could encrypt a string  
366 containing the user's internal stable identifier and the name of the service they're  
367 accessing. The external identifier could then be built on-the-fly each time it's needed. This  
368 type of identifier is called a pair-wise identifier.

## 369 **4.4. Social IDs**

370 Most people already have IDs from one or more social media platforms. These can often be  
371 linked with institutional usernames in your IAM System to provide an alternative login method.

### 372 **4.4.1. Do: Consider where account linking of social IDs to institutional** 373 **accounts might be appropriate and where it might not.**

374 Social credentials such as those issued by Google or Facebook can be a great choice for  
375 users who interact with non-regulated data. For example alumni returning after a long  
376 period to request a transcript, are much more likely to remember their Facebook login that  
377 they linked to their institutional identity, than to remember a username and password  
378 issued by your institution. Users are also quite likely to notice if their Facebook or Google  
379 credential has been compromised, which they won't for an infrequently used institutional  
380 credential. Common social identity vendors such as Google are very good at notifying  
381 users about changes to their accounts and credentials. Be careful though when using social  
382 credentials for regulated data as you lose control of credential quality, and social providers  
383 generally won't tell you if multi-factor authentication was performed, so you may not be  
384 able to guarantee some of the controls that are required by regulations.

### 385 **4.4.2. Do: Consider whether social ID can be a step in** 386 **onboarding/offboarding.**

387 Social identities can be a great way to begin interacting in an authenticated session with  
388 users while you build up enough information to do identity proofing and other identity  
389 onboarding steps. Consider, for example, an applicant for a position in your HR system, it  
390 could make sense to use a social identity to allow the applicant to upload information  
391 during the application process, and even during background checks and other steps  
392 required before they are considered an employee in your HR system.

### 393 **4.4.3. Do: Consider Level of Assurance LOA when using social IDs.**

394 Understand the identification, registration, authentication, and related processes employed  
395 by the social providers whose IDs you use.

396 **4.4.4. Don't: Assume all services do proper authorization and make them**  
397 **aware of the LOA concept.**

398 Work with your service managers to assure they understand the risks and benefits of using  
399 social IDs for their service.

400 **4.4.5. Do: Identity matching, even with social IDs.**

401 It is fine for a user to have multiple social ID's, such as one with Facebook and one with  
402 Google. When registering those social ID's for use as a credential in your systems, keep  
403 track of the tie to a single institutional identifier. You don't want users trying to remember  
404 which social ID they used for which of your services.

## 405 **5. Identity Lifecycle**

406 Lifecycles of users in higher education environments are complicated. A given individual can  
407 appear as a high school summer camp attendee, an applicant, a student, an employee, a retiree, or  
408 any of a number of other roles. Further, that individual can have more than one affiliation at the  
409 same time. Keeping track of those affiliations and when they start and end is essential for  
410 providing and revoking access to services. The recipes in this section will help to define the  
411 parameters and processes needed to manage identity lifecycle.

### 412 **5.1. State and Affiliation Changes**

413 Many events, both large and small, affect identities and the access permissions associated with  
414 them. For example, a graduating student who is hired as an employee will gain permissions as  
415 appropriate to their new position and job responsibilities. That person will also lose permissions  
416 granted only to students. All of this is driven by information obtained from the source systems.

417 **5.1.1. Do: Capture changes in affiliations/roles that matter for service**  
418 **entitlements.**

419 The most important thing to do with user affiliations is to track when they change and to  
420 communicate those changes to downstream services in a timely fashion. Whether a user is  
421 being on-boarded, off-boarded, or just gaining or losing one of multiple affiliations, track  
422 this information.

423 **5.1.2. Do: Work with service providers to ensure service entitlements are**  
424 **being handled correctly.**

425 Once you're tracking affiliations, you need to help service owners to map those affiliations  
426 to specific service access. A service owner, for instance, might say that they want their  
427 service to be available to all students. But what does student mean? Full-time as well as  
428 part-time? On-campus as well as remote? Should they gain access to the service at the start

429 of their first semester, or should it happen when they've registered for classes? Help  
430 service owners to understand all of the different transitions and options for understanding a  
431 given affiliation so that they aren't giving too much or too little access to their service.

### 432 **5.1.3. Don't: Overdo state changes.**

433 If a state change doesn't change an entitlement, it may not be necessary to distinguish it. A  
434 student, for example, who has registered but not yet started classes might need to be  
435 distinguished as some services might want to grant access before classes start. There are  
436 other cases, though, where state changes might not change any access such as a staff  
437 member changing positions within a given team. As mentioned above, carefully consider  
438 what you consider an affiliation so that you don't end up with thousands. There's a balance  
439 between future proofing your affiliations and information overload.

### 440 **5.1.4. Do: Account for users with multiple overlapping affiliations.**

441 This scenario is very common in higher ed, but it also happens in other organizations.  
442 Someone can be a staff member taking classes or a student with part-time employment. A  
443 retiree can come back as a student. There are lots of possibilities. Ensure that your system  
444 can assign multiple affiliations to an individual that can be separately assigned or removed.

### 445 **5.1.5. Do: Designate a "primary" affiliation for each user.**

446 Most services will care only if a person does or does not have a specific affiliation. Some  
447 services, however, need to know a person's primary affiliation with the institution. For  
448 example, a staff member taking classes might have a primary affiliation of "staff," but a  
449 student employee's primary affiliation might be "student." Also,  
450 [eduPersonPrimaryAffiliation](#) is useful in some federation scenarios. Institutional policy  
451 will be required to determine which of a person's affiliations is primary.

## 452 **5.2. Grace Periods**

453 Revoking access permissions adversely affect your users, even when it is warranted. Grace  
454 periods can give people time to prepare.

### 455 **5.2.1. Do: Make grace periods.**

456 In most cases, you won't want an affiliation change to remove a user's access immediately  
457 so that access isn't prematurely terminated. For instance, staff might need access to some  
458 services for a couple of weeks after leaving the university, and students might need to  
459 retain access for the summer after graduation. Building a grace period into the transition  
460 between affiliations can make for a better off-boarding experience.

### 461 **5.2.2. Do: Work with stakeholders to determine how long a grace period 462 should last.**

463 How long a grace period for a certain state transition should be varies depending on your  
464 organization, the state change, and the services that it impacts. In some cases, immediate  
465 loss of access won't be a problem. In others, it could lead to lots of confusion, extra work  
466 for service owners, and help desk calls. Consider each state transition and work with  
467 service owners to choose timing appropriately.

### 468 **5.2.3. Don't: Overextend a grace period if it compromises security.**

469 It's easy to go too far on how long you leave things active. The longer an account stays  
470 active, especially if it's unused, the more likely that a compromise to that account will go  
471 largely undetected. Weigh this against user convenience when selecting how long access  
472 should be retained.

### 473 **5.2.4. DO: Allow for immediate deactivation when necessary.**

474 There will, of course, be cases where a grace period is not only unnecessary, but it's  
475 dangerous. The immediate termination of an employee or the departure of a high profile  
476 individual from the organization are often reasons to immediately remove access. Include  
477 a mechanism for bypassing the grace period when needed to immediately remove access.

## 478 **5.3. Deactivation**

479 Here is some guidance for when the grace period is past, and it's time to deactivate.

### 480 **5.3.1. Do: Retain minimal data when deactivating an identity.**

481 Just because someone's leaving your organization doesn't mean they're gone for good.  
482 Consider a situation like a student graduating, going elsewhere for grad school, then  
483 returning to take a job. For the purposes of identity matching to assign the same username  
484 or university ID number if available, or at least for building a complete history of a given  
485 user, it can be helpful to at least keep a stub entry with information about the previously  
486 active identity. What data is contained in that stub entry depends on your organization's  
487 environment and what you might want to track or reinstate.

### 488 **5.3.2. Do: Establish policies and processes to reinstate disabled identities.**

489 Just because you keep a stub entry for a disabled or departed user doesn't mean that re-  
490 activating the entry upon the user's return will be straightforward. Carefully consider how  
491 to most easily and efficiently turn the stub entry back into an active entry, confirm the  
492 user's identity, and possibly make sure that there are no lingering permissions granting  
493 them access to resources from their previous affiliation.

### 494 **5.3.3. Do: Communicate with service providers to inform them of timelines 495 for deactivating identities.**

496 When a user is scheduled for deactivation, the results of that can be far reaching. Not only  
497 might they lose access, but accounts will be deprovisioned from services, data deleted, and  
498 group memberships terminated. Whether you're doing this for a one-off staff member who  
499 has left or the entire graduating class of last semester, service owners should be informed  
500 ahead of time. In general, it's enough to make sure that service owners know your  
501 schedule: staff are deactivated after this many days, students after this many, *etc.*  
502 However, in cases where you're going to deactivate a large batch such as the example of  
503 the graduating class, an extra communication to service providers to remind them of the  
504 specific upcoming event can help prevent a lot of potential problems.

#### 505 **5.3.4. Don't: Deactivate or delete identities without communicating.**

506 It may seem obvious, but it gets overlooked far too often. Before you terminate a user's  
507 access because of an affiliation state change, let them know. Provide information on things  
508 they might want to do such as graduated students setting up email forwarding or staff  
509 members downloading any relevant content that they're allowed to take with them upon  
510 their departure from the organization. Communicating ahead of time can also help to  
511 prevent mistaken deactivation such as a student who you believe has left but was just late  
512 registering for classes.

## 513 **6. Passwords, Multi-Factor Authentication,** 514 **and Provisioning**

515 Login credentials, such as passwords, "passwordless" hardware tokens, cell phone "second  
516 factor" apps, and combinations of the preceding (*i.e.*, multi-factor authorization), when coupled  
517 with a user's identifier, are the security mechanisms protecting that user from the risk of being  
518 impersonated by someone else. Here is some advice to help you deliver secure authentication,  
519 appropriately balanced against end-user pain and frustration.

### 520 **6.1. Password Rules and Policies**

#### 521 **6.1.1. Do: Limit the number of different passwords that users need to** 522 **remember.**

523 It's hard enough these days to keep track of the many passwords we have in our personal  
524 lives. Adding multiple passwords for work or school will inevitably lead to forgotten  
525 passwords, user frustration, and help desk tickets. Use a central password store that all  
526 services can authenticate. Have a single password that grants access to everything that a  
527 user accesses. Where that's not possible, consider synchronizing passwords between  
528 password stores.

#### 529 **6.1.2. Do: Encourage single signon.**



530 The only thing better than a single password is single signon. Rather than having to log  
531 into each service using the same password, architect your services such that a user only  
532 needs to sign in once in a given session. In a browser, things like SAML-based  
533 authentication against an identity provider that supports a persistent session can make this  
534 easy to set up and add security to your environment. On the desktop, setting up things to  
535 leverage and trust the user's desktop authentication can accomplish the same result. In  
536 addition to user convenience, this adds security as you won't have to type your password  
537 into lots of different sites and services.

538 **6.1.3. Do: Consider the password policy advice from NIST (currently**  
539 **Special Publication 800-63B).**

540 NIST has done extensive research and engaged numerous experts in developing these  
541 materials. There is good advice here.

542 **6.1.4. Don't: Require frequent password changes.**

543 It used to be a good, security-conscious idea to require users to change their passwords on  
544 a regular basis. With the rise of multi-factor authentication, this has become much less  
545 important. With multiple factors required to access a resource, a hacker can't do much of  
546 anything if they crack or steal the password. Furthermore, not requiring users to change  
547 their passwords on a regular basis avoids forgotten passwords which generate help desk  
548 calls and doesn't encourage a user to write down their password on that sticky note stuck  
549 to their monitor.

550 **6.1.5. Consider: Passwordless authentication.**

551 "Passwordless" authentication tokens, usually based on FIDO2 protocols, may be a good  
552 fit in environments where having possession of a physical token may be more appropriate  
553 for authentication than knowledge of a secret password. (See also "Assignment of  
554 additional authentication factors" below.)

555 **6.2. Initial Password Setting**

556 **6.2.1. Do: Transmit account claiming information securely using activation**  
557 **codes or short-lived links.**

558 Care should be taken the first time a user sets a password. The best approach is to provide  
559 some temporary, one-time-use item to the user to grant them access when initially setting  
560 the password. This can be a link that expires after a short amount of time or a random  
561 string of characters making up an activation code that's sent to them through some secure  
562 out-of-band means.

563 **6.2.2. Do: Perform additional identity proofing during the account claiming**  
564 **process.**

565 While a secure link or access code is a good start toward secure setting of initial  
566 credentials, it's only the first step. Just like multi-factor authentication uses items of  
567 different types such as something you have and something you know, first-time setup of  
568 credentials should do the same. In addition to secure link, ask the user to answer questions  
569 that a hacker wouldn't know the answers to.

## 570 **6.3. Assignment of Additional Authentication Factors**

### 571 **6.3.1. Do: Use multiple factors for authentication when possible.**

572 Multi-factor authentication is becoming commonplace. It can certainly make things more  
573 secure when you require a user to present more than a secret string in the form of a  
574 password to gain access. It also alleviates the need for highly-complex passwords, a  
575 benefit for end-users. There are many technologies built to open standards for multi-factor  
576 authentication, including FIDO2, as well as older standards, such as HOTP and TOTP.  
577 There are also commercial authenticator apps, such as Google Authenticator and Duo.  
578 Shop around, and choose the ones right for your institution in consideration of  
579 accessibility, the availability of software support in browsers and online services,  
580 avoidance of vendor lock-in, and support for service providers that require specific MFA  
581 technologies.

### 582 **6.3.2. Do: Use additional validation to password for adding or modifying** 583 **MFA.**

584 If a user only needs a password to gain access to the MFA configuration, then you might  
585 as well not have multi-factor authentication at all. A hacker who got a user's password  
586 could log into the user's MFA settings and change those settings to something the hacker  
587 could leverage instead. Make sure that, to change MFA settings, a user must perform an  
588 MFA authentication or some other backup method to verify that it's really them.

## 589 **6.4. Provisioning and Deprovisioning of Credentials**

590 Provisioning credentials into a service provider may be the only way to provide at least an  
591 approximation of a single sign-on experience for your users when the service provider cannot be  
592 federated (*i.e.*, must do its own authentication). This is, however, a security risk. It increases your  
593 risk of unauthorized exposure of the credentials, perhaps outside the scope of your direct control,  
594 if the service provider is not operated by you.

### 595 **6.4.1. Don't (if possible): Provision credentials when a federation option is** 596 **available.**

597 Syncing passwords with service providers increases the complexity of password changes  
598 and resets, and increases risk associated with password exfiltration. Protocols such as  
599 SAML and OIDC are widely supported and give service operators no access to  
600 institutional credentials. Integrations with LDAP, AD, or Kerberos allow you to

601 avoid syncing passwords, but may give service operators access to password cleartext, so  
602 SAML and OIDC should always be the preferred path.

603 **6.4.2. Do: Avoid service-specific passwords or any password on the service**  
604 **side whenever possible.**

605 Service specific passwords are confusing to users, and increase the complexity of  
606 password changes and resets. When using the SAML or OIDC protocols, ideally no  
607 password needs to be provisioned to the service. If the service requires a password, even  
608 though it won't be used because of a SAML or OIDC (Google and others) then setting a  
609 random complex password that is unknown to anyone can safely meet the requirement to  
610 assert an abandoned password value.

611 **6.4.3. Do: Use federated authentication with a unique (pair-wise) identifier**  
612 **for each service provider.**

613 Whenever possible use federated authentication to avoid the issues of password sharing.  
614 Unless multiple SPs need to share a common identifier (and policy allows such sharing),  
615 assign pair-wise identifiers to enhance privacy and to simplify transitions when identifiers  
616 must be changed.

617 **6.4.4. Do: Establish criteria to assure that the service provider's security**  
618 **measures for protecting credentials are comparable with yours.**

619 Document these criteria and hold the service provider accountable. Remember that a  
620 breach of their system will be a breach of your IAM System and, potentially, all other  
621 service providers that rely on it. Your security people will be a good resource for  
622 developing the criteria.

623 **6.4.5. Consider: Periodic audits of the service provider's compliance with**  
624 **your criteria.**

625 The need for this will depend on a number of factors. Consult with your security people  
626 for what is appropriate.

627 **6.4.6. Consider: Deactivating provisioned accounts, rather than deleting.**

628 You may need to reactivate quickly, preserving the user's resources and state. On the other  
629 hand, there may be licensing fees that must be paid, even when an account has been  
630 deactivated. (See, also, the [Deprovisioning](#) section below.)

631 **6.4.7. Do: Keep enough information around to, at the minimum, prevent**  
632 **reissuing a username.**

633 In addition to preventing reuse of a username, there may be files and other resources that  
634 may be needed in the future.

635 **6.4.8. Consider: In addition to the first factor, deactivate second-factor**  
636 **tokens if the account no longer has access to anything.**

637 Multi-factor authentication is great until a user returns to your institution several years  
638 later with a new phone number, email address, or whatever external entity you were using  
639 for the second factor. Suddenly, the user can't get in to set up their account because a one-  
640 time password or similar code is being sent to a phone number that they haven't had in  
641 years. There's little to no point to leaving multi-factor methods active once a user no  
642 longer has access to anything of use. It only leads to challenges when they return, and has  
643 the additional possible side effect of increasing licensing costs if the user is occupying a  
644 license slot with an MFA vendor. As soon as the user has been deactivated and no longer  
645 has access to anything of use, deprovision second factors.

646 Also consider the issues in the [Deprovisioning](#) section, below.

## 647 **7. Service Provisioning**

648 Many services require "out of band" provisioning, as opposed to "just in time" identity  
649 assertions. There may be solid business reasons why a service needs to know about a user before  
650 that user logs in (*e.g.*, to tell the user *how* to log in), identity assertions received at the start of a  
651 session being unable to enable this. It's also possible that a service is simply not technically  
652 capable of receiving new identity information at the start of a session. This section discusses the  
653 issues associated with service provisioning.

### 654 **7.1. Reconciliation**

#### 655 **7.1.1. Do: Ensure that source and destination are in sync.**

656 Deploy processes that ensure that synchronization errors, such as missed transactions, are  
657 corrected promptly.

#### 658 **7.1.2. Do: Have both targeted and full reconciliation (fully match accounts).**

659 Periodic full reconciliation for all information associated with all accounts is important to  
660 repair errors due to lost transactions. The need for targeted reconciliation (*e.g.*, for a  
661 particular user) can be used to repair errors or to enable (or disable) authorizations for a  
662 user, before the next full reconciliation is run.

### 663 **7.2. State Changes and Fine-Grained Authorization** 664 **(Continuous Access Evaluation Protocol)**

665 **7.2.1. Do: Look at more than institutional data for fine-grained state**  
666 **changes.**

667 Session data such as a mobile user's location also needs to be communicated (geofencing)

668 **7.2.2. Do: Keep up-to-date on emerging technologies in this area.**

669 Fine-grained authorization is a developing field.

## 670 **7.3. Communicating Updates to Service Providers**

671 **7.3.1. Do: Have a reconciliation process to correct missed messages or**  
672 **reconcile out-of-sync changes in the case that the service provider did not**  
673 **respond to the change.**

674 Incremental updates can, at times fail. A periodic reconciliation process will heal any  
675 damage.

676 **7.3.2. Do: Have a process to handle failed updates, connections errors, etc.**

677 Have a monitor to detect failures in connecting and processing updates both from your  
678 registry and source systems and provisioned service providers. Test transactions with  
679 expected results testing are a great way to catch non-retryable errors that effectively mean  
680 your data provisioning is offline. Also include monitors for retryable errors, many times  
681 sources are targets are temporarily unavailable, and detecting a transaction fail and then  
682 replying the transaction when the system becomes available can prevent silent drift of  
683 targets from the source of truth. Detecting a service has become available can be done by  
684 polling to retry the translation, or by using a test translation to make the target as available.

685 **7.3.3. Do: Have both incremental and full reconciliation.**

686 Having near real time integration typically means detecting incremental changes in your  
687 identity registry or source system and pushing just the changed entries to your directory.  
688 Occasionally an incremental change may fail or be incomplete, or despite trying to prevent  
689 it, data might get updated in a service provider. Maintain a process that is capable of  
690 reconciling all of the entries in a service provider and verifying that the data matches your  
691 "source of truth" identity registry and source systems.

692 **7.3.4. Do: Make sure to delete service provider entries before deleting in the**  
693 **person registry.**

694 If entries are ever removed from your identity registry, ensure you have a process that  
695 removes accounts in any service providers first, and the process assures the transaction  
696 succeeded. Orphan service provider entries are not only difficult to clean up, but can leave  
697 open application access that should have been removed.

698 **7.3.5. Don't: Permit updates from other sources to overwrite updates from**  
699 **the person registry.**

700 It is important that account and authorization data carried by registries reflects the source  
701 of truth that is your identity registry and other authoritative data sources. Attributes should  
702 not be changed by other systems, or be updated directly in the service provider. In cases  
703 where the target directory must allow direct updates for business reasons, such as is  
704 frequently the case with Active Directory, isolate via OU structures the accounts that are  
705 managed by your identity system and prevent any other updates in those OUs.

706 **7.4. Deprovisioning**

707 Deprovisioning is the removal of access to provisioned services and plays an important role in  
708 maintaining the security of electronic systems. A review of an individual's access should occur  
709 anytime that individual's affiliation with the institution changes. Common review times are  
710 when an employee terminates and when a student graduates but can also include times when an  
711 individual's role within the institution changes, such as an employee moving to a new  
712 department or a student switching majors. A best practice with respect to deprovisioning is to  
713 remove the entitlements and authorizations to services rather than just disabling the  
714 authentication credential. This decreases the potential for someone to inherit inappropriate access  
715 as would be the case if an institution recycles usernames and the new "johndoe" user  
716 automatically has access to all the services that the old "johndoe" had because the entitlements  
717 were never removed. It's important to have clearly documented and published rules and time  
718 frames regarding deprovisioning. Under what circumstances will I lose eligibility to Service X?  
719 How long will it be after I graduate or leave employment before Service X goes away? A best  
720 practice would be to send "pending service expiration" notifications or reminders to an  
721 individual during that grace period time frame. Not only will your users thank you but it could  
722 prevent the extra work of reinstating access if a mistake was made regarding eligibility or the  
723 user's eligibility changes during the grace period time frame.

724 Sometimes deprovisioning can also involve removal/disabling the accounts on the  
725 downstream/target service . For example, terminating an employee can result in removing their  
726 Box account (provided you have a policy for this).

727 You should consider the flexibility of supporting both deprovisioning via removal of access and  
728 also removal of accounts for the target services. Make sure the business rules and account  
729 mapping support both configurations.

730 **7.4.1. Do: Gather all data you need for reporting, auditing, reactivating, etc.**  
731 **before deprovisioning.**

732 There are many reasons why you may need to access information about deprovisioned  
733 access to services. The information may also be useful when it's necessary to verify a  
734 user's past ownership of the identity, as well as for forensic investigation. Make sure you  
735 have everything you may need before deleting.

736 **7.4.2. Don't: Forget deprovisioning, even when using just-in-time**  
737 **provisioning.**

738 Login events are typically used for just-in-time provisioning, but they cannot be used for  
739 deprovisioning, since you cannot expect a login event. Some other method, probably  
740 driven by periodic reconciliation processes will be needed when it's necessary to purge  
741 information stored by the service or, for example, to reclaim the service provider's end-  
742 user licenses or other resources.

743 **7.4.3. Do: Deprovision authorizations.**

744 It is sometimes the case that some, but not all, permissions must be removed for an  
745 individual's use of a service, based on changes received from identity sources or manual  
746 changes to group membership, *etc.* Full deletion is not always the correct action,  
747 particularly if other information is stored within the service, such as conversation  
748 transcripts, videos, files, *etc.*

749 **7.4.4. Do: Give user an opportunity to migrate data, tools, instructions, *etc.***  
750 **before they lose access (depatiation).**

751 When appropriate, be sure end-users are given sufficient warning to allow them to obtain a  
752 copy of their data from a service for which deprovisioning is pending.

753 **7.4.5. Consider: Whether you need a process for transferring ownership of**  
754 **the user's data to their unit/manager/*etc.***

755 The "user's" data is often really the institution's. When this is the case, you will need a  
756 process for transferring that data to a successor.

757 **7.4.6. Do: Plan for potential repatriation (in or out: claiming accounts**  
758 **created before provisioning, or extracting data at**  
759 **departure/deprovisioning).**

760 Deprovisioning is often followed in short order with a request for reprovisioning. Be  
761 prepared for this..

762 **7.4.7. Do: Set up processes ahead of time and provide users with an**  
763 **opportunity to preserve data (tools, instructions, *etc*) before deprovisioning**  
764 **occurs.**

765 **7.4.8. Do: Consider that accounts may exist before provisioning occurs and**  
766 **plan for dealing with it.**

767 This may or may not include permissions from previous eligibility that should have been  
768 cleaned up at deprovisioning.

769 **7.4.9. Don't: Accidentally restore old permissions that should have been**  
770 **cleaned .**

771 Repatriation may not include all of the old permissions. Restore only what is needed at the  
772 current time.

## 773 **7.5. Considerations for Cloud Services**

774 Cloud services, particularly commercial cloud services, present unique issues for provisioning  
775 and deprovisioning, due to likely use of proprietary technical solutions and difficulties with  
776 navigating corporate structures to find the "right" people. For more information see the  
777 [Provisioning and De-provisioning](#) section of the [Cloud Services Cookbook](#).

# 778 **8. Target Directory Provisioning**

779 Directories are an interesting form of service provider, in that they typically are used as identity  
780 sources for other service providers. This has implications for how you provision those  
781 directories.

## 782 **8.1. Linking Identities between Directories**

783 **8.1.1. Do: Have one or more attributes that are unique and immutable to**  
784 **link identities between source and target.**

785 It is important to be able to resolve accounts in various systems that belong to the same  
786 identity. Having one or more attributes that are in your identity registry, and in each of the  
787 target directories will allow you to quickly identify accounts as changes in role, activation,  
788 attributes, and entitlements change.

789 **8.1.2. Do: Create an opaque institutional identifier used solely for linking.**

790 Maintain an identifier to facilitate linking your directories and identity registry that doesn't  
791 contain strings that are meaningful to other aspects of the users identity. Strings like name  
792 based usernames, role based account naming conventions will change over time. The more  
793 opaque and limited the scope of linking identifier is, the less likely it is to change and  
794 break the association of a person across directories and registries.

795 **8.1.3. Don't: Use things like NetID that may seem immutable now.**

796 NetIDs are tempting identifiers since they are often recognizable, published, and known by  
797 the user, however they are generally not a good choice for linking. Often NetIDs are name



798 based or user chosen, and as with any non-opaque identifier are subject to change. Also,  
799 certain targets will have requirements that add scoping strings to the NetID for use as an  
800 authenticator, for example User Principal Name in Azure or Primary Email in Google. Not  
801 having a simple identifying attribute that is exactly the same in your target directories and  
802 you registry may hamper identity resolution.

803 **8.1.4. Do: Have a unique identifier for each of the target directories and**  
804 **make this mapping available to the provisioning process.**

805 This will be useful to differentiate between target directories that may share, for example,  
806 the same Relative Distinguished Name (RDN), such as LDAP and AD.

807 **8.2. Communicating Updates to Target Directories**

808 **8.2.1. Do: Use a reliable process or frequent deltas to push changes in as**  
809 **close to real-time as possible in the intended manner.**

810 Target directories are frequently the primary source for applications to learn about new  
811 identities and retrieve needed attributes and role information about the identity holder.  
812 Keeping the data in target directories as close to real time as possible means users gain,  
813 and just as importantly lose, access to applications that use directories for authentication  
814 and/or authorization.

815 **9. Authorization**

816 Authorization is an extension of institutions' processes for delegation of authority into digital  
817 services and resources. Institutional policies determine which people are allowed to access a  
818 service or resource, and what those people are allowed to do. Those policies may grant access  
819 entirely on the basis of institutional person data already known to your IAM system, such as  
820 affiliation, major, department, *etc.* (often referred to as "business" or "functional"  
821 roles). Alternatively, the policies might require an explicit decision by someone responsible for  
822 the service or resource.

823 A fundamental principle in IAM practice is "AuthN  $\neq$  AuthZ". In other words, a user's ability to  
824 authenticate (AuthN) should not imply authorization (AuthZ). Authentication only identifies who  
825 the current user is, not what that user is allowed to do (authorization). Authorization is  
826 determined based on information about the current user that is contained in the IAM system.  
827 That said, it should be noted that protocols such as SAML and OIDC appear to combine both  
828 authentication and authorization into a single step by identifying the current user and  
829 transmitting information about that user in a single transaction. Logically, though, the SAML IdP  
830 must perform authorization before gathering information to support an authorization decision,  
831 then transmit the results of both operations in the same transaction.

832 **9.1. Types of Roles**

833 **Business Roles** are generally used to describe high level affiliations or duties/functions within an  
834 organization. Examples are : Employee, Staff, Faculty, Staff, Student, Alumni, Admit Coming,  
835 Guest, Contractor, *etc.* These can be considered high-level or “course grained” roles, as they may  
836 only tell part of the story of an entity, and may be used for basic entitlements like software  
837 licensing or building access. Business Roles are usually determined by institutional person data  
838 obtained from the authoritative source systems.

839 **IT Roles** are generally used to assign access rights and entitlements for specific services and,  
840 therefore, are more fine grained. IT Roles can be further categorized as follows:

- 841 • **Application Roles** describe end users' functions and entitlements, and can be further  
842 divided into:
  - 843 ○ Admin Roles, such as “IdM admin tool” or ”Box Admin User”
  - 844 ○ End-User Roles, such as ,”Office 365 User” or "Box User”
- 845 • **Asset Roles:** This describes assets and devices that can be provided to end-users, such as  
846 hardware tokens, PC, Laptop , ID badge, mobile, *etc.*

847 Oracle's [Using Role Types to Design Flexible Roles](#) provides more advice for the structuring of  
848 roles.

849 In practice, roles are often implemented as **Groups** (collections of users), combined with any  
850 additional information needed to describe the associated access rights and entitlements.

## 851 **9.2. Methods of Authorization**

852 The most used methods for determining access rights are Role Based Access Control (RBAC)  
853 and Attribute Based Access Control (ABAC). There is a growing body of literature in this area.  
854 Wikipedia provides good starting points:

- 855 • [Role-based access control](#)
- 856 • [Attribute-based access control](#)

857 (When reading these references, note that the authors often assume a person has one role. In  
858 academia people often have multiple roles, such as multiple job appointments, or being both  
859 student and employee.)

860 You build a collection of roles (what a person does for or with the institution), possibly add more  
861 personal or environmental data (attributes), and authorize accordingly. For example, a student  
862 (role) in engineering (detail), currently on campus (environmental) is authorized to access the  
863 School of Engineering's student portal.

864 Whenever possible, authorization should be determined by Business Roles and identity  
865 Attributes that are determined through automated processes. This allows authorization and de-  
866 authorization to be automated, making sure that exactly the right people have access. If "every  
867 Registrar's office employee should have access to student records", or "every engineering student  
868 should have access to this file share", authorization can be granted and revoked as a person's job

869 or study field changes. When authorization is determined manually, de-authorization tends to lag  
870 behind, leaving people with access they should not have.

871 All authorization processes should have periodic re-examination, often called "attestation" or  
872 certification. If you have an automated process for "every Registrar's office employee should  
873 have access to student records", verify that the rule is still correct, as well as the logic for  
874 determining a Registrar's office employee. If authorization is manual, then each person who has  
875 access should be manually verified. (See [Assuring Provisioned Authorization Is Correct](#), below,  
876 for more information.)

877 Automated creation of roles from institutional data is good, but you will need to create roles that  
878 cannot be determined in this fashion. Not every access can be determined from institutional data.  
879 Also, the effort to automate may be greater than the benefits. Start automating where you can get  
880 the most benefit and reduce risk.

881 Authorization is usually implemented in two ways. Authorization decisions can be made in  
882 advance and stored within the application (periodic updating as changes needed) that will  
883 execute them, or the application can call an "AuthZ service" to request an authorization decision  
884 at the time of need. The former is the traditional and still most common method, but the latter has  
885 advantages and is growing in use. You will want to support both methods.

886 Commercial systems are available for implementing authorization. When investigating  
887 commercial systems, make sure that they can handle the complexity of your affiliations -- people  
888 who are both employee and student, multiple job appointments, *etc.* There is also a community-  
889 supported system, InCommon Grouper.

### 890 **9.3. Designing Roles, Attributes, and Groups for** 891 **Authorization**

892 Authorization decisions can be made on the basis of users' *roles*, *i.e.*, what they do for or with the  
893 institution; this is called Roll Based Access Control (RBAC). What a user does for/with the  
894 institution implies a set of responsibilities and authorities, and those authorities imply access to  
895 services and what those services will do for the user. *Groups*, lists of users, are used to represent  
896 the users that fill each role. As such, RBAC authorization decisions are relatively stable over  
897 time.

898 Other *attributes* can also be used to make authorization decisions; this is called Attribute Based  
899 Access Control (ABAC). Examples of such attributes include identity information, such academic  
900 department, but may also include attributes such as time of day, that are difficult to express with  
901 roles when criteria are relatively dynamic over time. Judicious use of RBAC and ABAC will  
902 result in a simple, yet robust design.

903 It should be noted that the distinction between the terms "role" and "group" can be fuzzy. Many  
904 computer systems provide access to a resource by making everyone who is authorized member  
905 of what they call a "group" or "security group". Some systems provide access similarly with

906 "security roles". This is so common that the terms "groups", "roles" and also "membership" may  
907 be used to describe authorization processes, regardless of the mechanisms actually used in a  
908 target system. In this document, we will try to be consistent with the definitions provided above.

909 Whenever possible, business roles and groups should be composed automatically from  
910 institutional data. Think of a person's data as determining group membership, and then granting  
911 access to "all members of this group".

912 At the highest (coarsest) level, roles describe a basic function in the institution-- faculty, student,  
913 guest, contractor, employee. Depending on your needs, you may want to split employee into  
914 regular and temporary, or add a campus for a large system (*e.g.*, FacultyMadison). At a  
915 university, it is common for people to have more than one affiliation. You can use these roles to  
916 drive eduPersonEntitlement, and to determine access to campus level services (*e.g.*, all students  
917 get Google mail).

918 The next finer level will be determined by the access needs of organization. If you need to grant  
919 access based on school or college, then define groups such as 'student in engineering' or 'faculty  
920 in public health'. Here too, people may be in multiple groups.

921 At successively finer levels, you may need to define groups such 'student in English 123' or  
922 'senior accountant in engineering'. These groups would enable you to automate access grants to  
923 course materials or the financial system.

924 The following are considerations for the design of your roles, groups, and attributes.

### 925 **9.3.1. Do: Start simple.**

926 The highest level roles and groups will allow you to determine access to campus-wide  
927 services. Their large size will help you work out mechanisms for computing, storing, and  
928 delivering authorizations.

### 929 **9.3.2. Do: Have a clear naming convention.**

930 This includes components to identify who owns the group, who or what the group  
931 represents, what resource the group has access to, and how that group has an effect on that  
932 resource. Naming varies based on the type of registry that is storing the data. For example,  
933 Grouper uses its folder (stem) path as part of the group name: "app:splunk:index:firewall2-  
934 read" shows the group "firewall2-read" as part of the stem "index", beneath the "splunk"  
935 stem, and thus under "app". This clearly indicates the ownership (Splunk service), the  
936 resource (some index named firewall2) and the effect or action (read).

### 937 **9.3.3. Do: Express RBAC for Business roles.**

938 These permissions usually map to specific operations on one or more resources. The idea  
939 of RBAC is that you map the role once to the specific set of permissions. The on-going  
940 maintenance is simply the membership of those entities in that RBAC role.

941 **9.3.4. Do: Express ABAC for your IT roles.**

942 Attribute Based Access Control can give you more flexibility to add additional information  
943 to the role. For example, A Business role “Staff” can be assigned multiple IT roles as  
944 attributes such as “Box Admin” “Office 365 User”. Similarly, a person with a location of  
945 "Communicable Diseases Lab" might be granted access to an application that admits  
946 guests to the facility. This will help you avoid role explosion. Reference NIST 800-162  
947 for more details. <https://csrc.nist.gov/publications/detail/sp/800-162/final>

948  
949 ABAC can facilitate complex rule sets from multiple identity attributes and allow for the  
950 creation of automatic groups based on those attribute values. For example, any student  
951 with two attributes “HomeCollege=Engineering” and “CurrentTermRegistered=true”  
952 could be the supporting details going into an ABAC group of all currently-enrolled  
953 engineering students.

954 **9.3.5. Do: Consider the tradeoff of adding, for example, location or**  
955 **department as roles, as opposed to attributes.**

956 Since a role maps to permissions, generally an organizational component (academic unit or  
957 building) does not map to a permission. For example, “The Math Building” doesn’t  
958 convey a permission, but rather a resource. You can create cohorts or groups based on a  
959 person’s location or department, which can be used in ABAC policies.

960 **9.3.6. Do: Follow common concepts provided in documentation such as**  
961 **InCommon’s Grouper Deployment Guide.**

962 The [Grouper Deployment Guide](#)'s advice is largely independent of your particular  
963 group/role management platform.

964 **9.3.7. Do: Have clear and complete documentation of the design of your**  
965 **authorization architecture.**

966 Your clients, co-workers, successors (and auditors) will love you for this.

967

968 **9.4. Implementing Roles, Attributes, and Groups for**  
969 **Authorization**

970 The following are considerations for the implementation of your roles, groups, and attributes.

971 **9.4.1. Do: Consider how you will handle authorizations that change *en***  
972 ***masse* with academic term/sessions.**

973 This is especially true if you authorize access based on enrollment in a class.

#### 974 **9.4.2. Do Consider allowing grace periods for de-authorization.**

975 Consider grace periods to ensure continuity of service (and to reduce the disruption of  
976 large volumes of changes smoothly) when de-authorization is likely to be followed by re-  
977 authorization in a short period of time. For example, authorizations granted to all students  
978 are probably best continued between the end of one term and the start of the next, unless it  
979 is know that a student is not returning. Also, consider how you will inform people that they  
980 have entered a grace period.

#### 981 **9.4.3. Consider: How to to grant access to someone who does not meet the** 982 **automatic criteria, or deny access to some who does meet them.**

983 Consider how can you do that in such a way that the exception will have time limits, or  
984 trigger periodically reevaluations.

#### 985 **9.4.4. Do: Consult with your institution's auditors on any requirements for** 986 **maintaining a history of access decisions.**

987 You will likely find your internal auditors to be strong partners as you build your IAM  
988 service.

#### 989 **9.4.5. Do: Keep authentication systems separate from enforcing** 990 **authorization decisions, where possible.**

991 Attributes and entitlements should be handed off to the service for it to handle its own  
992 authorization enforcement. For Services that lack proper authorization mechanisms, use a  
993 groups registry (*e.g.*, Grouper or groups in Active directory) to set an access policy where  
994 the IdP becomes the policy enforcement point by allowing/denying access at the IdP.

#### 995 **9.4.6. Do: Adopt product-specific best practices where applicable.**

996 A good example is the [Grouper Deployment Guide](#) for Grouper. We emphasize this  
997 model, as Grouper has grown to become a popular open-source solution to group and role  
998 management among higher-ed institutions. As always, avoid adopting practices that create  
999 heavy dependencies on your current platform.

## 1000 **9.5. Computing, Storing, and Delivering Authorization** 1001 **Decisions**

1002 Authorization decisions may be pre-computed and stored as roles, groups, and/or attributes, or  
1003 they can be computed "just in time," when a service or resource requests it. The choice depends  
1004 on a number of factors, most importantly whether the decision includes consideration of

1005 information that changes over time, like time of day or location, or if the decision includes only  
1006 information that is updated at known times by the IAM system. Other factors, such as storage vs.  
1007 compute cost and complexity may also be relevant.

1008 **9.5.1. Do: Reconcile and re-compute the decisions periodically when**  
1009 **authorization decisions are pre-computed.**

1010 Consider the acceptable lag time between when a person's information is updated and  
1011 when an authorization decision is computed and stored. This is often very short (*e.g.*, sub-  
1012 second), requiring event-driven computation, but even when it is not, regularly-scheduled  
1013 reconciliations should be established to meet institutional and end-user expectations, as  
1014 well as to recover from missed event-driven transactions.

1015 **9.5.2. Do: Implement "just in time" computation of authorization decisions**  
1016 **when policy dictates consideration of factors that change over time.**

1017 This is information, usually associated with the current session, such as time of day,  
1018 location, authentication method (*e.g.*, password, X.509 certificate, biometrics, MFA  
1019 token) or behavior patterns.

1020 **9.5.3. Do: Make sure the service provider is authorized to inquire about the**  
1021 **particular authorization.**

1022 Ensure that your attribute release policies support service providers' need to query for roles  
1023 and attributes that have been determined to be required for authorization decisions.

1024 **9.5.4. Do: Communicate fine-grained institutional changes to services when**  
1025 **appropriate.**

1026 Ensure that service providers have the information that it's been determined they need.

1027 **10. Assuring Provisioned Authorization Is**  
1028 **Correct**

1029 Your IAM system likely interfaces with multiple sources of thousands to hundreds of thousands  
1030 of identities to support tens to hundreds of services and resources. The advice in this cookbook  
1031 will help you deal with this large scale, but errors will, undoubtedly occur. This section discusses  
1032 attestation and audit, two IT governance processes designed to catch those errors and correct  
1033 them.

1034 **10.1. Attestation to Review Access Decisions**

1035 Attestation is a process whereby the people who are responsible for services are asked to verify  
1036 that the IAM system is correctly configured to support authorization decisions by their service.  
1037 This includes both how users are selected to be provisioned, as well as any entitlements that are  
1038 associated with those users. The attestation process should be designed to verify entitlements,  
1039 roles, and other attributes that are used to provision and deprovision service-specific permissions,  
1040 as described in the [Service Provisioning](#) section.

1041 **10.1.1. Do: Focus on what really needs a human to consider.**

1042 If access is granted by rule based on institutional data, only the rule needs periodic  
1043 attestation -- not the individual grants.

1044 **10.1.2. Do: Make attestations understandable, so it is clear what is being**  
1045 **attested to.**

1046 Make sure this is well documented. If possible create the documentation in partnership  
1047 with the person responsible for the service at the time the service is integrated with your  
1048 IAM system.

1049 **10.1.3. Consider: Attestations' effect on resulting deprovisioning processes.**

1050 Changes to a person's entitlements, roles, *etc.*, due to attestation, may initiate provisioning  
1051 and/or deprovisioning. This should be considered when establishing grace periods, *etc.*

1052 **10.1.4. Consider: Attestations that are not completed should result in access**  
1053 **suspension until completed.**

1054 If the party responsible for the attestation has not responded to the attestation review on  
1055 time, before the system can take action to remove access, it should first send escalation a  
1056 few times before access deprovisioning happens. The escalation should support multiple  
1057 levels, both horizontal (peers) and vertical (management).

1058 **10.2. Audit**

1059 Periodic audits of IAM processes is a more holistic review of your IAM system and its  
1060 relationships with the rest of your institution's IT ecosystem. It may include, for examples,  
1061 service managers' policies for the assignment of roles and entitlements, as well as the operation  
1062 of your IAM system, in light of institutional policy, regulatory requirements, and industry best  
1063 practice.

1064 **10.2.1. Do: Involve your institution's auditors early and often.**

1065 They can be a great help in designing your system and processes to facilitate regular  
1066 reviews to minimize the time incorrect information remains in your system.



1067 **10.2.2. Do: Schedule regular “full review” processes starting from**  
1068 **destination content and working backwards.**

1069 Audits require a reasonable amount of effort, easily postponed. Establishing a regular  
1070 schedule for this work will help ensure that it is not forgotten.

1071 **10.2.3. Do: Schedule regular review of authorization processes.**

1072 Because auditing data can add up so quickly, consider adding/exporting data to an  
1073 external database on a schedule compatible with the regular "full" and authorization  
1074 reviews. This will help performance significantly .

1075 **11. Product Lifecycle**

1076 Over time, an increasing number of services will rely on the provisioning functions of your IAM  
1077 System. When selecting provisioning products (software or service, open source or commercial,  
1078 on-premise or off-premise, SaaS or IaaS) for use in your IAM System, as well as the services  
1079 you are provisioning, it is important to consider the eventual end of life of that product, as well  
1080 as how you will integrate that product initially.

1081 **11.1. Product Onboarding**

1082 There are many issues you must address in the selection and deployment of a new product. The  
1083 following are specific to provisioning.

1084 **11.1.1. Do: Have a plan for loading your data into the product.**

1085 This is likely to be a combination of migrating existing provisioning-related data, as well  
1086 as new data that may be required by the new product. Services you provision may also  
1087 store user-owned data that must be loaded, particularly when migrating from a competing  
1088 product.

1089 **11.1.2. Do: Understand (and mitigate) the risks for sensitive data.**

1090 This is particularly true of cloud-based products.

1091 **11.1.3. Do: Understand the provisioning "hooks" in the product.**

1092 Understand what out-of-the-box connectors included in the products. Can the product be  
1093 customized to add or extend new connectors? This includes triggers to initiate provisioning  
1094 actions, service provider integrations, *etc.*

1095 **11.1.4. Do: Understand how you will map roles and attributes to**  
1096 **provisioned service provider permissions.**

1097 Maintain interface mapping documentation while onboarding service providers, keeping  
1098 tabs on the attributes, default values, transformations, schema validations (field type,  
1099 character lengths, nullable, *etc.*). Also consider how attributes affect provisioned  
1100 permissions. Can the mapping be configured (*e.g.*, via a GUI or by editing files) without  
1101 developing software code?

1102 **11.1.5. Consider: How the product's license terms may affect cost as your**  
1103 **institution's use of the product grows.**

1104 Is it a flat fee? Does it increase with population growth, number of "seats," storage  
1105 requirements, *etc.*? Does it consider alumni, or other internal or external users? Does it  
1106 use a standard size-ranking structure, such as IPEDS? Will the cost structure require you  
1107 to restrict usage or to offboard no-longer-eligible users with shortened grace periods?

1108 **11.1.6. Consider: Whether and how the product can be customized, and**  
1109 **who can do the customization.**

1110 IAM system requirements are often very unique to each institution, requiring unique  
1111 customization. This issue is particularly important for SaaS IAM services, where the  
1112 vendor has control over the deployment, as well as the software implementation.

1113 **11.1.7. Do: Consider all other technology acquisition issues that are not**  
1114 **specifically related to IAM.**

1115 This includes ongoing maintenance, vendor stability, ease of deployment, *etc.*

1116 **11.2. Product Offboarding**

1117 When selecting and onboarding a new product, consider your eventual exit strategy.

1118 **11.2.1. Do: Know how you will retrieve your data, workflows,**  
1119 **procedures, *etc.* from the old product.**

1120 Avoid products that do not have clear ways to do this.

1121 **11.2.2. Do: Know how sensitive data will be destroyed.**

1122 This is particularly important for cloud-based provisioning services.

1123

1124