Enterprise Authorization

Scott Thorne

Why is Authorization Important?

Authorization management is ordinarily not at the top of a higher education institutions list to improve, but the use of clean and regulated data is both critical and foundational to student success and employee productivity. Before diving into the details of how to best determine and give authorization, let's first define what authorization is. The definition of authorization is determining what allowances someone has, once they are identified.

Data is exploding on campuses and making proper use of it is a must. Certain core data such as Financial, Student, and HR, is sensitive and needs complete and proper management. All institutions have the responsibility of controlling sensitive data. You hear in the news of hackers breaking in and stealing data, so it is often dismissed as a technical problem that a firewall can guard against, rather than a business issue. Yet, there is also a significant risk of an employee's actions leading to a data spill or misuse. One of the first steps to reduce this potential misuse is to restrict access to the data itself. Many business processes make use of this data, so many employees have a legitimate need to access it.

Using data effectively is a key to student success and your institution should make it easy for employees to access the records they need.

Goal and Benefits of Authorization

There are many possible processes used to grant access to important data, from very informal to strict and rigid. Improving how users gain access can have many unexpected benefits.

What are the characteristics of an optimal authorization process?

- It should be easy for an employee to request access to data needed to do their job effectively
- It should be easy to determine who has access to sensitive information
- Authorizations should be easy to understand
- Access to data should need explicit request, not the data itself
- Access should not be broader than needed
- There should be ways to restrict access to certain sets of information
- The person responsible for the information should be deciding who has access to it

- Grant authorization to approved users
- Record all authorization decisions
- Authorizations should be consistent across systems

Why are these characteristics important? It's easier to look at examples where this is not the case, to see what the effects are.

Employees Should Easily Request Access to Data

This seems obvious but is not always the case. What happens when it is difficult or time-consuming to get access to data?

Employees will often find a way around difficult processes. They might ask someone they know who has access for a copy. When they do get access, they are more likely to save it in case they need it again. They are more likely to use an alternate source of the data which might not be as accurate as the truth so these practices work against good data practices. So one of the quickest and most straight forward ways of improving data practices is making getting access to the right data easy.

It's hard to eliminate inappropriate data sharing and storing if the alternative is not simple and easy.

Easily Determine Who Has Access to Data

Record Authorization Decisions

- It should always be possible to determine who has access to what data, but often this involves inspecting a technical implementation to see what is in effect. It should be simple to find out the answer to a question such as, Who currently has access to the Social Security Number (SSN) of Students? document the authorization decision, so that it is clear who made the decision and when, and then there is a way to compare this to what is currently in effect. This also enables analysis of whether access employees need access for the job. It's better to have this thought out before the Auditors ask the question.

Give Authorization to Employees Who Need It to Do Their Jobs Successfully

If all authorizations are not explicitly granted, it will always be harder to determine who has legitimate access.

Employees often prefer to avoid the work of defining and then recording the authorization decisions. The most common practice is to use existing lists, or job titles to derive access. This method is complicated and often leads to frustration. For example, using an existing mailing list to grant access, can later allow someone to put a person on the list to get the mail without realizing that it also grants access.

Before you attempt to infer access based on something else, try to determine if in all cases it is true. Anytime you can imagine an exception it can lead to workarounds, which will undermine the process. There are cases, such as the person who supervises an account needs access to it, where it might be appropriate. It's
typical to also need the ability to grant this access to someone who is not the supervisor of the account. It's better to think of having an explicit authorization for access to an account, and then have a process to make sure that when a person becomes a supervisor they receive this authorization. Automation is ideal during this process. This makes it easier to later determine who has access from the authorization rules, without having to then apply various exceptions which are not written down.

The Employee Responsible for the Data Should Own Access Control

Who should decide who gets access to data? This is a tricky question and varied widely with each use case, so let's look at some examples.

Financial data is usually stored in a financial system under the control of the CFO. In order to maintain clean records and accurate data, it is the CFOs discretion to decide who should gain access to see those records.

In higher education, each School or Department has a series of accounts set up for their use only. Assuming that the financial information in this subset of accounts is theirs to manage, it may be more appropriate for someone in that local organization to be responsible for access. Take the case of a research project that is managed by Department A. If it is a collaborative project, employees from another department may need access to the account information. Who is in the best position to understand the request and grant it? The CFO or central financial person might not know the details of the project, so someone in Department A who is responsible for the Project might be the best to decide. Obviously when someone higher up at the Institution wants

to see everything including Department A's accounts, then Department A doesn't need to be included in this decision.

In many cases, creating authorizations should be done in a physical central location, close to the actual business process. They have to rely on information others give them. There is often no way for the employee who should know to review data what has been worked on and manage the work appropriately. This disparate process often leads to oversights and errors. The more that the process can have the decision maker record the access decision directly, the better. If that decision maker can easily review who has access, there is a better chance that data will remain accurate and uncompromised.

Make Authorizations Easy to Understand

Access rules must be easy to understand, implement, and expressed in business terms, not technical ones. For example, "Sally can view salary information for the School of Science", instead of granting access to a series of database tables or a system screens. When granting or reviewing authorizations it should be easy to understand what effect they have. It is common for the implementation of access rules to be implemented by IT, who have to interpret what a business person intends. This is error-prone. Employees often end up with more access then they need, but management may not notice.

Best practice is: recording the authorizations in a three-part structure that mimics an English language sentence. For example, include granting access to a business function. Person A can view salary information for the School of Science. Person A "can view salary information," which is the business function and is limited to the scope of the qualifier, in this case, the "School of Science". By pattering authorization rules this way most access rules can be clearly represented. Having them all take the same form makes them easier to understand. Anything more complicated gets confusing and is harder to execute. Rules at this level are easy for both business users and technical folks to understand.

Keep Access Tightly Maintained

Fine grain access controls enable the largest amount of proper data access since only intended data can be accessed. Without the ability to limit the scope of authorizations; either very few employees will get to see everything, or employees will get access to more than they need.

For example, imagine you have the authorization to see a class list, so that either you can see all or none of them. There are some employees whose job requires that they see them all, so this works for them. But what about limiting the employees who should only view certain class lists only. If there is no way to do that then either they won't have access or they'll have too much access. In the case of them not having access, there is a request for the data they need. This

creates a manual process, as someone has to respond and create a dataset, usually a spreadsheet, and then send them a copy.

Request Access to Data

Creating data sets and passing them around causes a number of problems, so reducing this behavior should be a goal. Data sets have a short shelf life and typically go stale upon creation. If there are changes or updates to the underlying data, they won't be included. Working with old data is fine for several tasks, but it's far better to work as close to the source as possible. When a data set gets created, it is usually filtered. This might be appropriate, but if the data is recycled it may cause confusion for a second user. If there is a discrepancy in data cleanliness, reverse engineering the error becomes difficult. Compounding this difficulty is a human error because using people as gatekeepers to create datasets on request is not a good strategy. The only way to avoid this situation is to have procedures in place to grant access at exactly the right level.

Authorizations Should be Consistent Across Systems

Often the same data exists in more than one system, such as the SOR and a Warehouse. Ideally, a person's authorization should be consistent between these systems. It makes no sense for someone to be able to access information in one system but not in the other. The best way to achieve this is to record the authorization rule in a central system and then drive the authorization checks from both systems from this. If you only maintain authorizations at the individual system level, then it will take extra work to ensure the consistency of the rules.

How do you know if you have a good Authorization process?

Start by answering these questions:

- Is there a documented process for getting access to data?
- Is it easy to find out who has what access to data?
- What information is considered sensitive at your institution?
- Are authorization and data access decisions recorded?

It takes a considerable team effort to achieve clean data that is properly maintained, but once the value is transparent and buy-in is achieved, the steps to student and employee success are well underway.