

# *Access Control based on weighted claims*



General Information: [info@cionsystems.com](mailto:info@cionsystems.com)  
Online Support: [support@cionsystems.com](mailto:support@cionsystems.com)

CionSystems Inc.  
6640 185<sup>th</sup> Ave NE  
Redmond, WA-98052, USA  
<http://www.CionSystems.com>  
Phone: +1.425.605.5325

## **Trademarks**

CionSystems, CionSystems Inc., the CionSystems Inc. logo, CionSystems Active Directory Manager Pro are trademarks of CionSystems. Other trademarks and registered trademarks used in this guide are property of their respective owners.

Contributors – Zubair Ansari, Paul Cayley, Alejandro Buschel

## WHAT IS THE CURRENT SITUATION?

Traditional access control systems make binary decisions, based on predefined rules, by evaluating users' or processes' identity and group membership. This model is no longer viable. The move to cloud-based services is creating multiple authorities for identity and access, multiple providers of functionality and resources, and a highly decentralized nexus of control.

Binary authorization systems are typically static design which evaluates a single or small number of attributes to make an authoritative decision. This requires a great deal of upfront design and planning. Also, static, binary weighting requires making tradeoffs. The decision process and scenarios where such tradeoffs are appropriate is separated from daily operations. This creates risks and complicates establishing a secure, well managed environment.

“Yes or no” is not always the right answer with today's growth, evolution, and expectations. “Yes or no” decision forces arbitrary and inflexible decisions. This becomes problematic when transitioning from centralized control by authoritative sources to decentralized ownership by peers. Worse yet, “yes or no” business practices tend to get circumvented or modified without proper analysis.

Modern environments must authenticate both internal and external users, and then provide fine grained access control. This is needed so that data owners can meet statutory requirements, provide appropriate levels of access and prevent data leakage. IP protection is extremely important; particularly in the emerging business models that are dependent on IP sharing and monetization. Access must be modulated based on characteristics of the requested resource and the requestor. IT architects and implementers cannot be expected to understand subtle nuances of an internal and external user community, let alone dynamically changing byzantine business rules driving access requirements.

Federation is simple in concept: The identity life cycle is delegated to a trusted party. This includes adding, modifying, and deleting users. It also includes routine housekeeping, such as password resets, and attributes maintenance. The implementation that is the dominant method in use today is to allow claims asserted by the identity provider to be used to authenticate the user. Claims are also used to identify roles and privileges held by the user.

## AUTHORIZATION AT THE SPEED OF BUSINESS

Claims-based federated access control provides a dynamic mechanism for granting access to resources. However, the real power of this approach remains untapped. Authentication and access decisions are calculated, dynamically, from a rich set of assertions. This allows a gradated “degree of certainty and trust” to be calculated. The calculation is not Boolean – it is a weighted value, representing the risk presented by the entity requesting access.

We will refer to this computed value as “the Trust Level”(TL)

Thus, a variety of factors can be evaluated – in real time – when granting or denying access. Resource owners can weigh the identity provider, the source of the request, what claims are or are not present, and, of course, the content of the claims themselves. In addition, previously computed trust level, current level of threat, metadata for the target resource, permissions being requested, volume of requests, time of day and possibly agent type can be factored in as well.

Current access control methodologies don't adequately differentiate failed access events as either (passive) lack of trust or (active) distrust. This is a crucial bit of information. In the simplest case, consider someone attempting to enter a building. They swipe a card, and it is not authorized for that building, so they simply do not get access. Now consider someone who previously had access, and this access was revoked. This second person swipes their card and alarms go off, either locally, remotely, or both.

This simple analogy directly mirrors IT use cases. Unfortunately in the IT world, audit and alerting are layered onto the infrastructure in the hope of identifying these scenarios. So the deep integration and proactive use of these technologies is not built into applications. Further, business processes and operational realities often result in both audit and notification being used for forensics purposes, rather than to protect, prevent, and respond quickly.

However, with weighted authorization, we detect suspicious activity in real time. When an access attempt results in a user being actively distrusted, remedial action is automatic. For instance, the user can be redirected to an alternate site. This "honeypot" does not host sensitive information and is closely monitored. This arrangement protects sensitive information, improves visibility regarding active attacks, supports evidence gathering, and does not alert interlopers that their current attack vector has been detected. Thus, attacks are deflected.

Not only is the datacenter moving to the cloud, the workforce \*has\* left the building. Home workers, either full or part time, are now spread across the globe. Trust, granting access, identity, and audit now must extend to the interconnection fabric, the system, the consuming application, security devices, protocols, and any human that might or might not be present. It simply is not possible to express every possible scenario as a Boolean expression involving groups and user identities. This has already caused bloating in the number of groups, group memberships, and partitions holding resources. We must adopt frameworks for access control that support more advanced flexible and reactive scenarios. Designers must weigh the impact of false negatives against business goals and profitability. This means that the resultant implementations allow a few "bad guys" to get through because the alternative would be blocking access to legitimate users. The rationale is "audit and monitoring" will detect and prevent damage.

Experience has shown that is not the case. In making these decisions, the implementers are weighing risks that may be unknown, or dynamically changing overtime. In either case, the risk assessment is not based on an accurate, known level of risk and impact.

For example an increasingly common scenario in today's world is where one Identity Provider is compromised. The recent breaches at social media sites like LinkedIn are a case in point. Many

applications now accept authentication and authorization claims from multiple social media sites. If a site is compromised, traditional, binary (yes and no) mechanisms can no longer safely base access on the compromised site's attestation. However, this is not necessarily the best approach. If identity was based on weighing multiple assertions, then valid identity and access claims from another site could be combined to grant access. In fact, going back to the trusted vs. explicitly untrusted scenario, discussed above, the user who is authenticated by the compromised Identity Provider (IDP) but fails on the intact IDP becomes highly suspect!

Identity lifecycle management is a critical and significant undertaking for most enterprises. A great deal of time and effort is placed on the upfront design of provisioning processes, which are critical to onboarding and operations. De-provision access is critical for security and compliance.

The ITAR, or International Traffic In Arms Regulations, is one of many examples where the export or viewing of data is statutorily controlled. It is important to note, that these regulations typically do not factor in intent, or accidental disclosure when assigning culpability. That is, if information is inadvertently transmitted to a proscribed location or presented to an contravened individual, the ramifications can be severe.

Changes in variables like citizenship, credit score and criminal convictions can affect the level of access a person can have to specific applications or documents. For security and compliance, as well as operational efficiencies, having

Using a flexible, extensible authorization framework can bring multiple systems into the equation. This make the decision to grant or deny access a more granular and dynamic operation.

Similar exposure exists for Medical (HIPPA), Financial (PCI), Trade Secret, and Personal information. Having the most effective mechanism available to prevent accidental disclosure or intrusion is essential

## Personnel and Background Checks

Office 365 ITAR-support plans ensure that Microsoft personnel who have access to customer-owned data are U.S. citizens in accordance with International Trafficking in Arms Regulations (ITAR), which restricts access to ITAR- controlled data to “U.S. Persons” or “U.S. Citizens.” In addition, all customer-owned data is processed and stored in Microsoft data centers that are located only in the United States.

All Microsoft personnel who have access to customer content that is hosted in Office 365 ITAR-support plan environments undergo the background checks and screenings that are described in Table1.

**Table1. Microsoft Personnel Screening and Background Checks**

Check	Description
<b>Employment History Check</b>	7-year employment history verification.
<b>Education Verification</b>	Verification of highest degree attained.
<b>Social Security Number (SSN) Search</b>	Verification that the provided SSN is valid.
<b>Criminal History Check</b>	A 7-year criminal records check for felony and misdemeanor offenses at the state, county, and local level, as appropriate, as well as at the federal level.
<b>Office of Foreign Assets Control List (OFAC)</b>	Validation against Department of Treasury list of groups with whom U.S. persons are not allowed to engage in trade or financial transactions.
<b>Bureau of Industry and Security List (BIS)</b>	Validation against Department of Commerce list of individuals and entities barred from engaging in export activities.
<b>Office of Defense Trade Controls Debarred Persons List (DDTC)</b>	Validation against Department of State list of individuals and entities barred from engaging in export activities related to the defense industry.
<b>Fingerprinting Check</b>	Fingerprint background check against FBI databases.

Source: <http://www.microsoft.com/en-us/download/details.aspx?id=23910>

In this case, there may be additional physical controls to limit the physical location from which the person can effect changes on the environment. Now, imagine that after establishing a trust level for a person to access the control environment, it would be possible using weighted claims, to verify one or multiple of the checks defined above, like the DDTC, in real time.

In this scenario, the concept of continuous monitoring that is being extended to network traffic will be applied to human actors, and their requests for access.

People are being increasingly targeted by phishing, water holing, and clever targeted attacks. Moving to two-factor authentication could minimize these risks, but it is an effort that must reach all applications.

Instead, the authorization framework can be augmented by capturing the multiple variables available today when an authorization request is made and grant access based on 'trust level' there by limiting security exposure.

## Conclusion

Having an authorization framework that captures and weighs multiple factors provides a far more elegant and powerful authorization, security and governance framework. Not only can access be granted based on the assurance level, failed access attempts can be evaluated in a meaningful way. The normal operational issues associated with failed logins, forgotten passwords, and system unavailability will not mask the more serious incidents, indicative of attack or compromise.

More importantly, the continued evolution and adopting of cloud computing make rigid, coarse grained, authorization and security mechanisms problematic. Adopting a more intelligent, resilient, self-healing, and defensible Identification, Authentication, and Authorization framework is essential. Deploying weighted authorization will drive financial rewards, market differentiation, and an ability to better address modern needs.