



## White Paper: The Business Case for Data Governance

### CONTENTS OF THIS WHITE PAPER

Data Governance .....	1
Why Today's Solutions Fall Short .....	2
Use Cases .....	3
Reviewing Data Permissions .....	3
Reviewing Data Permissions with Varonis .....	3
Reviewing User and Group Permissions .....	4
Reviewing User and Group Permissions with Varonis .....	4
Providing Audit Data .....	5
Providing Audit Data With Varonis .....	6
Changing Access Permissions .....	7
Changing Access Permissions with Varonis .....	7
Identifying Data Owners .....	8
Identifying Data Owners with Varonis .....	8
Summary .....	8

## The Business Case for Data Governance

Data governance is one of the biggest initiatives in IT today. It encompasses the people, processes, and procedures to create a consistent, enterprise-wide view of a company's data in order to:

- Increase consistency in entitlement decision making
- Decrease the risk of data misuse
- Improve data security

Data governance and the business case for it is clear— if you can ensure that the right people have access to only the data that they need to perform their jobs, you will automatically improve security, reduce risk, and take a significant step toward least privilege access compliance.

### DATA GOVERNANCE

The issues in bringing data governance into the enterprise, however, come in the steps one takes to do so because there is a fundamental mismatch between the access controls that business owners assume are possible and those that can actually be established and enforced in the IT world. Business owners, for example, may consider their data assets to be protected as if they were physical, much like assets stored in a bank vault. These assets are the heart of the bank's business, and certainly not every employee would have access to them. Only those individuals whose jobs required access to the vault would have the ability to get there, and such rights might become more granular as the vault was further divided into safe deposit boxes, customer deposits, monetary reserves, and so on. Additionally, if individuals moved to branch offices, their ability to access the vault or any section of it in the main bank would automatically be terminated.

When you translate this simple example into IT terms, though, the analogy falls apart. One of the basic reasons is that most IT processes are oriented around providing employee access to resources, not restricting it. After all, if employees don't have access to the resources that they need, they cannot do their jobs. As well, while solutions such as Identity and Access Management may help with provisioning, it is ultimately the

business owners' responsibility to notify IT of changes to entitlement, and there is no real process in place in most enterprises to enable that action. This situation is made still more complex as employees enter and leave the enterprise, change roles or projects, and most importantly, by the astounding rate at which data volumes and access needs are growing. Still more complex are the problems faced when one considers unstructured data. Unstructured data is that which is outside of a database. Files stored in file shares are unstructured data, and such data in most cases is protected much more loosely than data stored in structured format. The problem is compounded by the fact that some of the structured data, thought to be protected, finds its way to the file system and is shared without the level of security expected by the data owner.

To get an idea of the scope of amount of unstructured data one merely has to consider a 2003 DM Review Magazine in which Merrill Lynch estimates "...that more than 85 percent of all business information exists as unstructured data."

### WHY TODAY'S SOLUTIONS FALL SHORT

Virtually all attempts to tackle data governance problems have resulted in only partial success because the proposed solutions are fundamentally static, while the problem changes all the time. At best, these techniques result, in a series of "snapshots" that try to capture the inherently dynamic nature of enterprise business owners, users, groups, and data. A real solution requires a dynamic method of enabling the following actions:

- Examining the permissions of users and groups to data
- Determining accuracy by tracing how permissions were granted
- Visualizing user/group permissions to folders
- Reviewing user/group actions on data
- Recommending entitlement changes
- Determining business impact by testing permission changes prior to enactment

Varonis® DatAdvantage™ offers a revolutionary way to accomplish these key data governance tasks. With DatAdvantage, enterprises can instantly see existing permissions by user/group or by folder, expedite IT operations and helpdesk response to entitlement requests and changes, identify the actual business owners of data, and greatly improve audit and forensic capabilities. DatAdvantage is based on Varonis' patent-pending Intelligent Data Use (IDU) analytics engine which aggregates user, data and access event information in a comparative matrix and applies algorithms that determine the groupings of users and data that belong together based on business need. The engine computes the relationships to show exactly who has appropriate group membership and who does not. This simple and elegant premise, which is executed by very complex statistical modeling, is wholly unique to Varonis and forms the core of the company's comprehensive platform for data governance.

In this paper, we will compare the actions that IT must take to complete typical data governance related tasks without Varonis and contrast the effort needed when compared to using DatAdvantage. Many of these tasks are part of the daily work of IT, helpdesk, security, and Active Directory management staff. We will analyze:

- Reviewing data permissions
- Reviewing user/group permissions
- Providing audit data
- Changing access permissions
- Identifying data owners

## USE CASES

### *Reviewing Data Permissions*

The task of reviewing data permissions is common during daily IT operations, such as the creation or modification of subdirectories and regular security audits. In most cases, an even cursory examination will show that Active Directory (AD) contains users with permissions in the Access Control Lists (ACLs) that don't match their business requirements. The task of reviewing and cleaning up these permissions is very inefficient and time-consuming using standard tools and is made even more difficult when there is a group hierarchy within AD. Typical steps required to review set permissions on a dictionary include:

1. Browse to the directory.
2. Go to the properties screen.
3. Check whether the directory inherits the permissions from a parent directory, which can be up to several levels above making it very difficult to trace the source of the inheritance.
4. Review the list of users and groups currently in the ACL.
5. Open the Active Directory Users and Computers Microsoft Management Console (MMC), for each group to identify who the members of the group are.
6. Drill down to get to the actual user list if there is a group hierarchy. This may require reviewing additional groups.
7. Repeat this process for each subdirectory, since some may not inherit the permissions of the parent directory and may have a different Access Control List, with different groups and group memberships.

The effort needed to accomplish this task is dependent upon the size of the company and its directory structure, and there is no direct correlation between number of directories that must be analyzed and the number of users, size of the company, or other concrete metrics. To get a sense of general scope, a typical enterprise will have tens or even hundreds of thousands of directories on each file server, while large file servers can contain a million or more directories each. It can take hours or even days to come up with an accurate list of the users who have access to a single directory hierarchy and the process must be repeated for every directory. Given the ever changing nature of user roles, new user additions, and growing data, by the time this information is gathered, it is almost guaranteed to be outdated.

### *Reviewing Data Permissions with Varonis*

Using DatAdvantage, the current permissions set for a directory is visible by simply clicking on the Directories Pane in the Work Area. One click will show you all the groups who can access the data, including the hierarchy of the groups and the members within them, color-coded by their permissions. You can also see which permissions were inherited instead of specifically granted for some purpose, and you can easily track the source of the permissions within seconds. Or, if you prefer, you can simply choose to run a directory permission report, which will provide the same information. This report can be set up as a subscription, or can be exported and sent to data owners and auditors.

### ***Reviewing User and Group Permissions***

Reviewing permissions by users or groups is also a very common access control task. This task is required when any change in permissions is requested or needed, and is usually performed by helpdesk or other operations personnel. Because this situation depends upon human input not performed by a dedicated resource, it contains many opportunities for small errors that can propagate throughout the entire enterprise. By simply extending default permissions without being aware of the access rights these permissions enable, for example, a simple action can have resounding consequences. Yet, it would be virtually impossible for the helpdesk or operations staff to obtain the information required to make a knowledgeable decision about individual permissions, because that data is not readily available to them.

As described above in the data permissions use case, understanding the relationship between users and groups is not a trivial task, especially given the hierarchical nature of Active Directory. A user may be granted permissions to data based on group membership at any level of the hierarchy, making it necessary to spend a significant amount of time mapping all the groups to which a single user belongs. If this task is multiplied by the number of users in the entire enterprise the scope of the problem quickly becomes overwhelming.

The following steps are required to identify all the groups that a user belongs to:

1. Open the Active Directory and Groups Computers Microsoft Management Console (MMC).
2. Find the user.
3. Select properties and review the list of groups in the "Member of" tab.
4. Repeat the process above for each group to find out if it is part of a hierarchy.

Once that task is accomplished, the reviewer now needs to identify all of the directories on every share or every file server that the user can access, based on the group membership information that has just been gathered. As described above, there is no way to accomplish this manually without spending a great deal of time comparing the Access Control List on each directory with the list of groups to which the user belongs. One possible alternative is to write a script or application to scan the file systems and perform the comparison, which requires a sizable time investment in coding and maintenance. Even with such a script perfected and in place, however, the process will still need to be repeated for every user at regular intervals or, every time permissions change.

### ***Reviewing User and Group Permissions with Varonis***

Just as in reviewing data permissions, correlating Active Directory users and groups with the directories that they can access is a matter of one mouse click. The information can be presented very easily, whether you begin with users and groups and drill down to individual permissions or start with directories. This makes it very easy to get a comprehensive picture of exactly who has access to what, and to how those permissions were granted. To simplify the process even further, by clicking on the user, you can see at a glance all of the groups to which he/she belongs, without having to sort through AD hierarchies. As well, report and log tabs can gather this information for you and present it in a variety of graphical formats.

**Providing Audit Data**

Especially in the current environment of regulatory compliance, IT is required to provide a wealth of data to be used as part of audits. Some examples may include:

TYPE OF DATA REQUIRED	TYPE OF AUDIT
Access permissions on folders, particularly those containing sensitive information, such as Finance, HR, or Legal	Sarbanes-Oxley, HIPPA, Gramm-Leach-Bliley, Data Protections Act, and many other state- or industry-specific regulations
Access data for specific files or folders	Internal protections of Intellectual Property
Data integrity information	Operational and general security purposes
Historical records of data access	Forensics

Providing the information required for an industry or internal audit can involve a combination of all of the steps delineated in the review processes above. As we have seen, such actions can prove difficult or even impossible to accomplish. This is especially true in the case of forensic audits, because the individual has a vested interest in not being detected. In situations such as these, the timeliness of the information is paramount, but is not feasible with current methods.

The situation is made still more complex by an often misused group in Active Directory itself: the “Everyone” group. This group type allows unrestricted connections to the server. Microsoft itself warns that use of the “Everyone” group can pose a serious threat to the organization if not used correctly, and advises that rather than allowing this default group to be used in your access control lists, it should be removed and each specified user explicitly added. This advice is often overlooked in today’s busy IT environment, however, resulting in a “ripple effect” of incorrect and sometimes dangerous access permissions being deployed as part of user setup.

Use of the “Everyone” group can cause a myriad of problems, many of which show up as part of any audit. If access to a directory is provided as part of the “Everyone” group, IT is reduced to examining log data on network traffic by IP address, or to interviewing users to determine which users were even aware that they had such permissions. Even in cases where Windows auditing was enabled, coming to a any sort of concrete conclusions requires digging through a mountain of log data. (It is worth noting here, that Varonis can show detailed folder access, even for those folders open to “Everyone” and without turning on Windows Server auditing functions.)

Still another audit related issue is forensics. If data is deleted, it may require a full audit to determine who could even access the data, let alone who deleted it.

### **Providing Audit Data With Varonis**

Not only does Varonis DatAdvantage make it easy to generate audit data, such data is actually provided for you as a standard part of the product. Logs and reports include:

#### Logging

- Detailed Access Summary – Displays a detailed log of daily events.
- Sensitive Detailed Access Summary – Displays a detailed log of attempts to access files that are specially monitored by DatAdvantage.

#### Statistics

- Directory Access Statistics
- User Access Statistics
- Tactical Access Statistics.
- Sensitive Files Access Statistics
- Total Files Accessed

#### Active Directory Group Membership

- Group Members
- User or Group Parents

#### File System Permissions (ACLs)

- Resource Permissions for User or Group
- User or Group Permissions on Directory
- Sensitive Files Permissions
- Global Access Analysis

#### Analysis Engine and Editing

- Summary of Changes for a Directory
- Summary of Changes for User or Group
- Access Denied Errors

#### Alerts

- Alerts Detailed Report – Displays all alerts from the selected day.

#### Archive

- Inactive Users
- Inactive Directories

The ability to generate access activity reports, by either user or data, becomes vitally important in the case of forensics. User access to or removal of data can often be stopped if the enterprise is aware of the activity while it is going on. It is also easy to determine who deleted a file by just examining an activity report.

### ***Changing Access Permissions***

Due to the hierarchical nature of Active Directory, changing user and group permissions is one of the most potentially difficult tasks that IT can face. While AD's feature of inherited and nested permissions are a significant benefit in provisioning, they can easily make data governance very difficult. A seemingly minor addition of permissions can result in sensitive data being exposed to the entire organization.

Inadvertently exposing data to users and groups that do not require it certainly raises risk of misuse but blocking warranted access has huge immediate impact in terms of lost productivity. Few things will create as much havoc as locking a group of users out of mission-critical data; yet, the hierarchical nature of AD makes it a very easy thing to do inadvertently. Many IT departments simply choose to stay on the safe side, and seldom reduce access permissions unless absolutely required to do so. It's an easy decision to understand, when one examines the steps required to reduce user/group permissions safely.

Like auditing, reducing access permissions requires performing all of the tasks outlined above, then taking them further. In addition to determining the current permission set, which involves correlating users and groups to directories, IT also has to determine what the correct permissions are. To do that, IT must either know the contents of the directories themselves or must have a way to involve the data owners in the process.

Each of the individual actions required to safely and correctly reduce access permissions are onerous tasks using available tools, if they are possible at all. Combining the tasks over large file shares quickly becomes extremely time-consuming and inefficient to the point of being useless in today's dynamic enterprise. After all, what good is it to establish who should access data if the users have changed in the time required to complete the analysis? What is the use of examining data use if the data itself has already disappeared? Why take actions if they are obsolete before they are completed?

The "Everyone" group rounds out the already bleak picture of user/group permissions changes. Virtually the only thing that the enterprise can do to reduce the problems caused by misuse of the "Everyone" group is to remove it completely and field the resulting complaints. Not only is this process inefficient, it is extraordinarily costly when the price of lost productivity is combined with helpdesk overload and manual corrections. Because most users don't access all the data that they need every day, the process can go on for months, effectively slowing the enterprise to a halt. As well, there is no way to ensure that the new permissions being granted are any more appropriate than those that have just been taken away.

### ***Changing Access Permissions with Varonis***

Because Varonis provides you with a real-time view of permissions by user, by group and/or by directory, changing permissions is relatively simple. But DatAdvantage uses the information of users and data access to take the value a step further with actionable recommendations. Using its foundation IDU technology, DatAdvantage can show you not only who can access data, but also who is actually doing so. Then, by correlating typical user a group activities, DatAdvantage can point out users whose behavior is not inline with others in the group, and display that information in the Recommended Users and Groups pane of the Work Area. IT can act on those recommendations or simply note them. Access permission review and modification becomes a matter of minutes.

Should IT wish to act on a permissions change, however, it's not necessary to take the action then wait for user response. By simply clicking on a user, a group, or a directory, and then changing user permissions in the Recommended Users and Groups pane in the Control Window, you can see at a glance what will happen before it does. In other words, Varonis can show you the impact of the changes in a sandbox. Once verified, you can then use the DatAdvantage "commit" function to push the changes in the live environment.

### *Identifying Data Owners*

This is another common task for IT, when data must be moved or archived, and when access permissions need to be modified. Using current methods, it is very difficult to identify the business owners of unstructured data without a significant amount of effort. One way is to review the Access Control List (ACL) and then contact every user listed. Alternatively, IT could examine the contents of the files within a directory and try to guess who the owner might be. Not only are these procedures extremely time-consuming and an inefficient use of IT time, they can pose a security risk at the least and violate regulatory compliance at the most.

### *Identifying Data Owners with Varonis*

While DatAdvantage will not specifically call out the business owner of the data, the visibility of users and access that it provides can narrow the field to a few likely users in just a few clicks. By reviewing actual usage, accessing information to a directory, and correlating that pattern with business roles, the business owner generally becomes obvious. The most active user is usually either the data owner themselves, or can readily identify who is. The entire process of reaching a likely owner candidate takes a few minutes, a review of information, and possibly one or two phone calls.

## **SUMMARY**

In today's enterprise, the data is the business. As IT processes have grown up around the data, though, many fail to take into account the dynamic nature of that data, which is compounded by the ever-changing user base who need to access it. Data governance is inherently difficult to ensure because today's methods use time-consuming, inefficient methods to reach static conclusions that are outdated almost before they are reached.

Varonis DatAdvantage provides an innovative new way to tackle data governance, by correlating Active Directory users and groups with the file systems and folders that they access. Information can be provided by examining the user's side of the equation to determine what groups they are part of, what permissions they have, and how they got them, or what directories the users and groups are accessing. This information is then combined using the Varonis Intelligent Data Use analytics engine, to provide recommendations on who should be allowed to access data, while a wealth of statistical information makes reviews simple and enables IT to take — and keep — control of data governance.