



Contents of This Paper

The First Frontier of Data Protection Is The File Server	2
Pragmatic Data Protection: Data Governance	2
The Ten Imperatives.....	3
Visibility	3
Control.....	3
Auditing	4
Security	4
Performance.....	5
Scale	5
Ease of Installation.....	5
Ease of Use	5
Ease of Integration	5
Low Total Cost Of Ownership	5
Conclusion	6
About Varonis	6

Preventing Data Loss: Ten Imperatives

Ever since corporate data losses and insider theft started dominating headlines, your charter has been to investigate the technologies and methods that will keep your organization from being the latest casualty. In your research, you've found a number of likely solutions, like putting rules at the perimeter to prevent sensitive data from leaving the network or locking down individual documents. Better yet, you can even lock down physical devices and memory cards so that if they are stolen, at least the data residing on them will be inaccessible.

So, which of these approaches is best to minimize the risk of data loss? The answer is probably all of them, in layers, because the problem is multi-faceted (i.e. sensitive data can leave the company any number of ways including as email, a print out or on a USB drive). The difficulty is that a broad technology rollout that comprises every sensitive document and every physical device will take a while to implement.

Is there anything to be done now?

The majority of unstructured data is aggregated on file servers and in data centers. It is from there that it makes its way to the devices of users. Protecting data at its source is a fundamental first step that can significantly reduce the risk of loss and misuse. This white paper describes the Varonis approach to data protection. It centers on the notion of data governance – a comprehensive means to see, control and audit all aspects of unstructured data access. The ten “must haves” for data protection and comprehensive data governance are detailed in the discussion that follows.

The First Frontier of Data Protection Is the File Server

Consider that the majority of your data, between 80 to 90 percent, resides on file servers. Now think about how you are controlling access to those shares. Most organizations find themselves with overly permissive access controls. Employees join and leave the organization frequently, and roles, responsibilities and project teams change quickly as well. All this leads to more access permission granted than revoked, since it is nearly impossible to manually keep up with the changes. The result is that most folders on file shares are oversubscribed in terms of access by well over 70%. By fixing broken access control to your file servers, you can significantly reduce the probability of data misuse in your environment.

Any program to reduce the probability of data loss and misuse has to start with rightful and warranted access controls. Ensuring that only the right people can get to the right data at all times not only reduces the odds of misuse, it also makes any subsequent safeguards and loss prevention techniques more cost effective and pragmatic to deploy. Consider a folder containing confidential data. If it is open to “everyone” or to a large number of individuals then (1) anyone can access and misuse the data, and (2) access by everyone must be monitored and audited – which is not a realistic undertaking. Alternatively, limiting access to those who actually need the data, and reporting on their access patterns, is realistic and a practical way to ensure that data access permissions are not abused.

Pragmatic Data Protection: Data Governance

Data file access controls are fundamentally flawed. File server permission mechanisms do not provide administrators with the ability to easily determine which users are accessing which files, how frequently and for what business purpose. This means that the documents, spreadsheets, presentations and other unstructured data that resides on file servers are at risk from overly permissive or ineffective access.

The first and fundamental step in protecting this data is a plan and a system for data governance. Data Governance (DG) is an industry term that defines the framework of people, processes and permissions or access rights that are employed to ensure proper data use. All enterprises need a strategy and process that governs appropriate and authorized use of business data. This approach significantly reduces the risk of data misuse. A comprehensive Data Governance infrastructure delivers:

- Effective and scalable data access control
- Increased and consistent data protection
- Reduction in the cost and complexity of data control
- Comprehensive and granular audit of data use

The Ten Imperatives

When considering which technology to implement to realize your data governance objectives, it is important to gauge the effectiveness of the solution against the ten imperatives of data protection. These imperatives are: visibility, control, auditing, security, performance, scalability, ease of installation, ease of use, ease of integration and low total cost of ownership. Additionally, any system for controlling access to unstructured data has to provide sufficient automation to make the process continuous and accurate. A detailed discussion of the ten imperatives follows.

Visibility

Any solution for unstructured data management and control must provide a clear visual representation of the access settings to the data as they are currently defined in the existing network. This visual must show, in an aggregated and searchable fashion:

- All users including their group memberships, Active Directory attributes and data permissions
- All folders and sub folders within a file server as well as the Microsoft NTFS permissions to this folder for any user or user group who is part of the domain
- Filtered views that allow queries based on username, group name or folder/data name
- Automated updating of views to reflect changes or new data within Active Directory (i.e. user to group membership) as well as within the file server (i.e. new data, deleted data, renamed data)

Control

Any solution for unstructured data management must include all mechanisms to define, test, enact and reverse file and folder permissions. Specifically the system needs to provide:

- The means to “push” or commit changes to access permissions directly onto the file server. The mechanism should include an option to push changes explicitly with system administrator intervention or in an automated fashion via a scheduler.
- “What If” capabilities, otherwise known as a sandbox where changes to folder permissions can be carried out in a simulated fashion in order to determine what, if any, the impact to access will be. For instance, the system shall allow the revocation of an entire group’s permissions in a sandbox. The system should indicate clearly which legitimate users will be affected negatively and allow for mitigation of that condition prior to live push.

Auditing

A detailed audit must be provided for all aspects of data use. The presentation of the information should be easily comprehensible and searchable. Specifically, the audit record should include:

- All file touches for a given Active Directory user (i.e. open, delete, rename)
- All access by access type (i.e. open, delete, rename)
- All access activity by folder
- All access detail to sensitive folders
- All inactive users
- All inactive data sets
- All administrative changes including security configuration changes by administrators
- On-hand searchable audit record for a period of no less than 12 months
- The information listed above should be available as reports in different formats, and should be exportable.
- The delivery of reports to subscribers should be automated and able to be scheduled
- The audit information should be searchable with support for complex Boolean (e.g., “and”, “or”, etc.) search conditions

Security

A system for unstructured data governance needs to provide an automated means for the revocation of data permissions. Specifically the system should:

- Identify by name all users whose access to a given data set should be revoked
- Re-compute revocations as changes to Active Directory and file servers occur
- Provide the means to test the recommended revocations prior to enacting on the servers for enforcement
- Provide revocations with accuracy greater than 3 nines (99.9%)

Performance

Any proposed solution for unstructured data management should not impede the performance of file servers, the user access experience or business traffic flow. Specifically, the system should not require Windows auditing in order to deliver its core functionality for data control

Scale

Because most organizations add additional file servers over time, and unstructured data can grow rapidly, the system has to provide room for growth. A data governance solution should be able to scale to accommodate unstructured data doubling in volume every 12 months.

Ease of Installation

A practical data protection solution cannot disrupt business operations or traffic flow. A solution should install quickly (e.g., within 5 business days), without the need for specialized professional services, and without assigning dedicated IT staff.

Ease of Use

A solution should not require specialized off-site training in order to operate. Any necessary training should be simple, and something the vendor can deliver themselves, on-site. Of course, the user interface should be intuitive.

Ease of Integration

Data protection solutions need to support a range of file servers and storage devices including Windows Server 2003 and network attached storage (NAS) from leading NAS vendors.

Low Total Cost of Ownership

A solution for data protection has to demonstrate quantifiable benefits in time and resource savings. Be sure to look for automation in the following areas, which are often the most manually intensive:

- Data permission revocations
- Data audit report generation
- Data entitlement review
- Stale data identification
- Data business owner identification
- Data migration

Conclusion

Varonis Systems is a software company unilaterally focused on data governance. Our solutions deliver on the ten imperatives for protecting unstructured data by showing exactly who has access to its, how individuals are using their permissions and who should have their access revoked. And, Varonis dynamically adjusts as changes to either directories or file servers occur, so that access controls to shared data are always warranted and based on business needs. With Varonis in place, the fundamental step to data loss prevention is addressed: limiting what data makes its way to laptops, printers and USB drives. That way, efforts to further protect data via filtering, encryption, etc., can be focused on only those items that are valuable, sensitive and actively being accessed.

About Varonis

Today Varonis is the foremost innovator and solution provider of comprehensive, actionable data governance solutions. The company's installations span leading firms in financial services, health care, energy, manufacturing and technology worldwide. Based on patent-pending technology and a highly accurate analytics platform, Varonis' solutions give organizations total visibility and control over their data, ensuring that only the right users have access to the right data at all times.

Varonis Worldwide Headquarters

499 7th Ave., 23rd Floor

New York, NY 10018

Phone: 877-292-8767

www.varonis.com