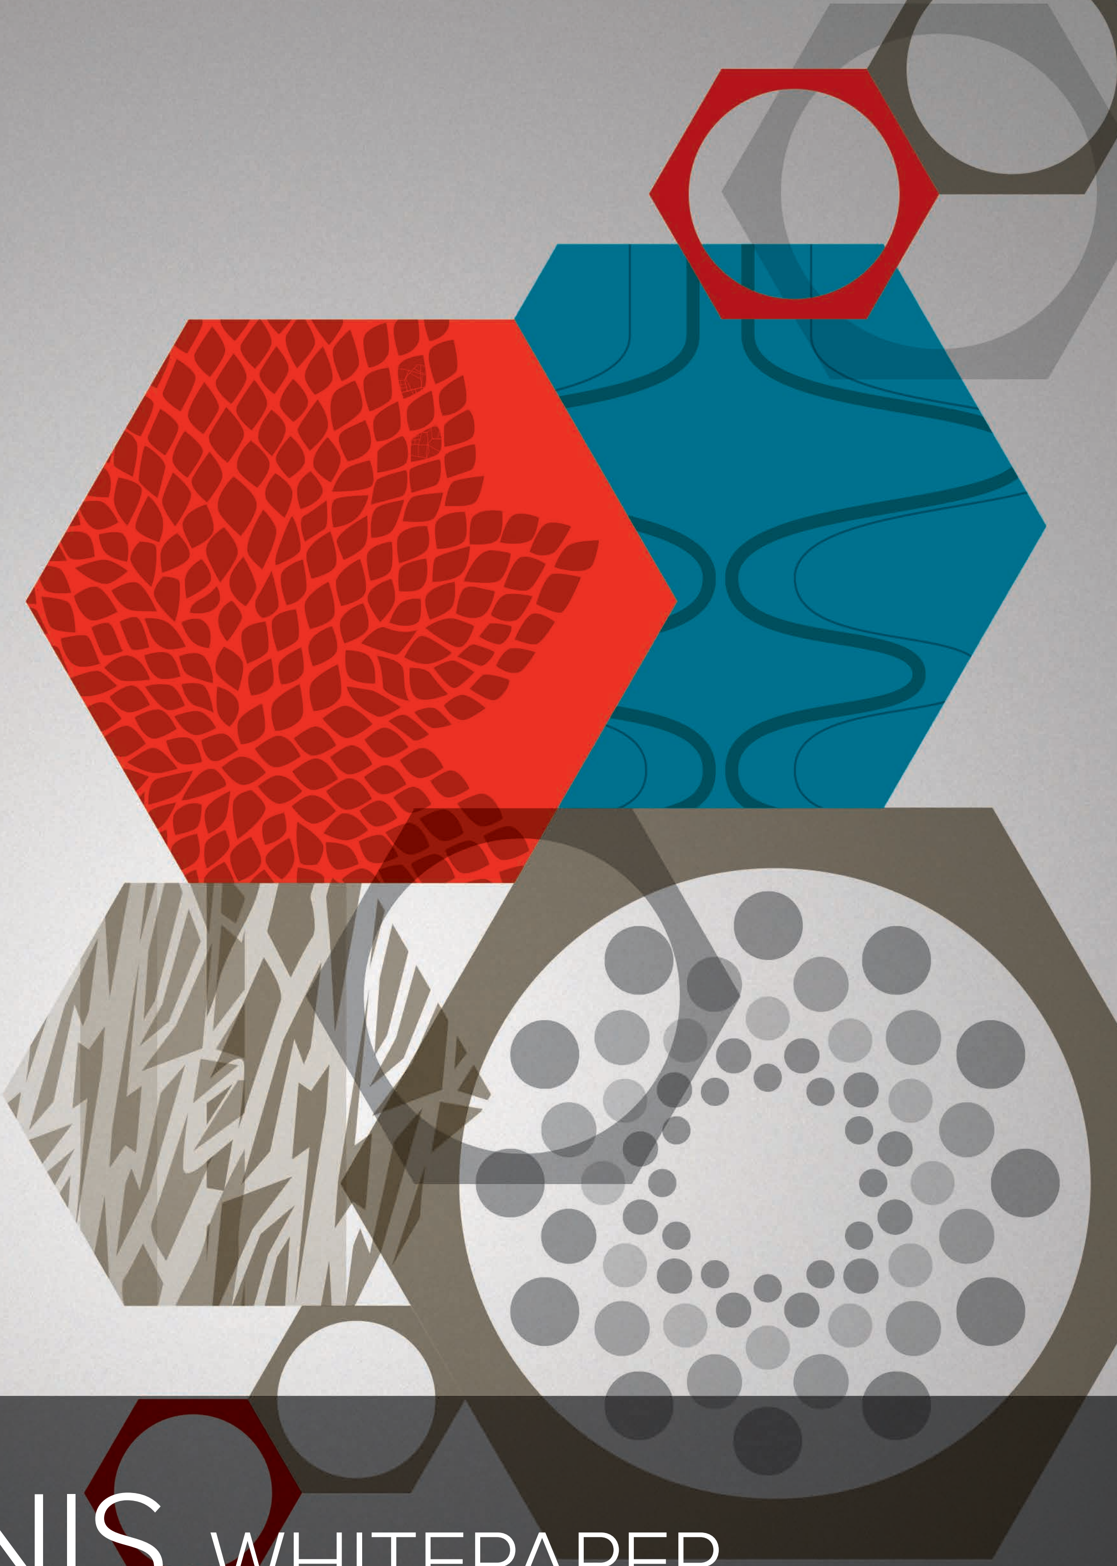

VARONIS WHITEPAPER

11 Things IT Should be Doing (But Isn't)

CONTENTS

OVERVIEW 3

TOP 11 THINGS IT SHOULD BE DOING 4

- Audit Data Access 4
- Inventory Permissions and Directory Services Group Objects 4
- Prioritize Which Data Should Be Addressed 4
- Remove Global Access Groups from ACLs (like “Everyone”) 4
- Identify Data Owners 4
- Perform Regular Data Entitlement Reviews 4
- Align Security Groups to Data 5
- Audit Permissions and Group Membership Changes 5
- Lock Down, Delete or Archive Stale, Unused Data 5
- Clean Up Legacy Groups and Access Control Artifacts 5
- One More Thing - Get Control of Public Cloud Services 5

11 IT MUST DO’S? NO PROBLEM WITH VARONIS 6

ABOUT THE VARONIS METADATA FRAMEWORK 7

VARONIS DATA GOVERNANCE SUITE 8

- Varonis DatAdvantage for Windows 8
- Varonis DatAdvantage for UNIX/Linux 8
- Varonis DatAdvantage for SharePoint 8
- Varonis DatAdvantage for Exchange 8
- Varonis DatAdvantage for Directory Services 8

VARONIS DATAPRIVILEGE 9

VARONIS IDU CLASSIFICATION FRAMEWORK 10

VARONIS DATA TRANSPORT ENGINE 11

VARONIS DATANYWHERE 12

11 THINGS IT SHOULD BE DOING (BUT ISN'T)

OVERVIEW

When it comes to protecting spreadsheets, documents, images and other data on file servers, SharePoint sites, and in Exchange mailboxes and public folders, most organizations readily admit that their current processes and risk profiles are less than ideal. Unfortunately, IT personnel – rather than the people that own the data – are the ones making many of the decisions about permissions, acceptable use, and access review. Since IT personnel do not have the business context behind the growing volumes of unstructured and semi-structured data, they're only able to make a best-effort guess as to how to manage and protect each data set. Until organizations start to shift the decision making responsibility to business data owners, IT, despite its best efforts, will continue to struggle to keep file permissions current and correct as data grows and user roles change.

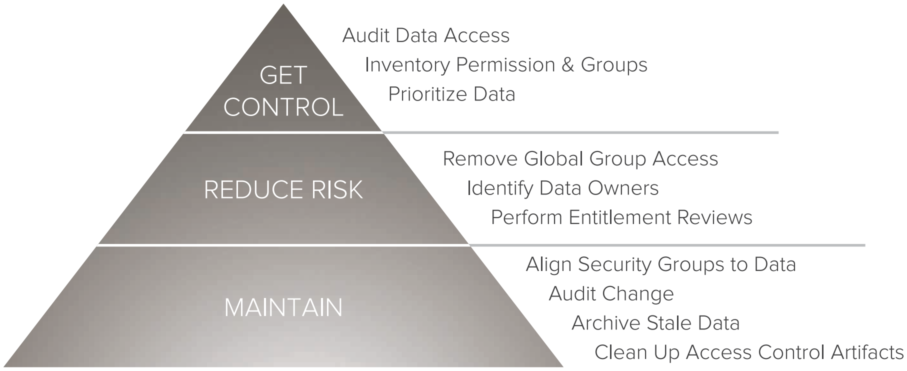
GET CONTROL
Audit Data Access
Inventory Permission & Groups
Prioritize Data
REDUCE RISK
Remove Global Group Access
Identify Data Owners
Perform Entitlement Reviews
MAINTAIN
Align Security Groups to Data
Audit Change
Archive Stale Data
Clean Up Access Control Artifacts

The principal of least privilege is a well-accepted guideline for managing access controls—i.e., only those that have an organizational need to access information should be allowed to do so. However, for most organizations, achieving a least-privilege model is almost impossible, because data is generated far too quickly and personnel changes are too numerous. Even in small organizations, the pace of organizational changes often exceeds the IT department's ability to keep up with access control lists and group memberships. Ideally, all organizations should automate the management tasks outlined below so that their access control processes can scale to the organization's needs, and can be conducted as part of a daily data management routine. Nevertheless, here are the 11 must-do's for maximizing unstructured and semi structured data protection.

TOP 11 IT MUST DO'S

1. AUDIT DATA ACCESS

Effective management of any data set is impossible without a record of access. Unless IT staff can reliably monitor data use, they can't spot non-use, misuse, or even abuse. Without a record of data usage, it's difficult to answer critical questions—from the most basic ones, like "who deleted my files, what data does this person or people use, and what data isn't used?" to more complex questions, "like who owns a data set, which data sets support this business unit, and how can I lock down data without disrupting workflows?"

2. INVENTORY PERMISSIONS AND DIRECTORY SERVICES GROUP OBJECTS

Effective management of any data set is also impossible without understanding who has access to it. Access controls lists and groups (in Active Directory, LDAP, etc.) are the fundamental protective control mechanism for all unstructured and semi-structured data platforms, yet too often IT can't easily answer fundamental data protection questions, like, "Who has access to a data set?" and "What data sets does a user or group have access to?" Answers to these questions must be accurate and accessible for data protection and management projects to succeed.

3. PRIORITIZE WHICH DATA SHOULD BE ADDRESSED

All data, of course, should be protected. But for a quick win, IT should focus initially on what might be considered "sensitive data." Some data sets have well-known owners and well-defined processes and controls for their protection, but many others are less understood. With an audit trail, data classification technology, and access control information, organizations can identify active and stale data, data that is considered sensitive, confidential, or internal, and data that is accessible to many people. These data sets should be reviewed and addressed first to reduce risk. Automation that moves, archives and deletes data based on content, activity, permissions and other metadata should be considered.

4. REMOVE GLOBAL ACCESS GROUPS FROM ACLS (LIKE "EVERYONE,") ESPECIALLY WHERE SENSITIVE DATA IS LOCATED

It is not uncommon for folders on file shares to have access control permissions allowing "everyone," or all "domain users" (nearly everyone) to access the data contained therein. SharePoint has the same problem (with authenticated users). Exchange has these, as well as "Anonymous User" access. This creates a significant security risk: lax directory access settings means that any data placed in a folder will also inherit those "exposed" permissions by default. When sensitive data, like PII, credit card information, intellectual property, or HR information are in these folders, the liabilities to companies can become very significant. Global access to folders, SharePoint sites, and mailboxes should be removed and replaced with rules that give access to the explicit groups that need it.

5. IDENTIFY DATA OWNERS

IT should keep a current list of data business owners and the folders and SharePoint sites under their responsibility. By having this list "at the ready," IT can expedite a number of the previously identified tasks, including verifying permissions revocation and review, and identifying data for archival. The net effect is a marked increase in the accuracy of data entitlement permissions and, therefore, data protection.

6. PERFORM REGULAR DATA ENTITLEMENT (ACL) REVIEWS AND REVOKE UNUSED AND UNWARRANTED PERMISSIONS

Every file and folder on a Windows or Unix file system, every SharePoint site, and every mailbox and public folder has access controls assigned to it which determine which users can access the data and how (i.e., read, write, execute, list). These controls need to be reviewed on a regular basis and the settings documented so that they can be verified as accurate by data business owners and security policy auditors. Users with access to data that is not material to their jobs constitute a security risk for organizations. Most users only need access to a small fraction of the data that resides on file servers. It is important to review and then remove or revoke permissions that are unused.

7. ALIGN SECURITY GROUPS TO DATA

Whenever someone is placed in a group, they get file system access to all folders that list the group on its ACL. Unfortunately, organizations have completely lost track of what data folders contain which Active Directory, LDAP, SharePoint or NIS groups. This uncertainty undermines any access control review project, any Role Based Access Control (RBAC) initiative. In Role Based Access Control methodology, each role has a list of associated groups, into which the user is placed when they are assigned that role. It is impossible to align the role with the right data if the organization cannot verify what data a group provides access to.

8. AUDIT PERMISSIONS AND GROUP MEMBERSHIP CHANGES

Access Control Lists are the fundamental preventive control mechanism in place to protect data from loss, tampering, and exposure. IT requires the ability to capture and report on access control changes to data – especially for highly sensitive folders. If access is incorrectly assigned or changed to a more permissive state without a good business reason, IT and the data business owner must be quickly alerted in order to execute remediation.

Directory Groups are the primary entities on Access Control Lists (Active Directory, LDAP, NIS, etc.). Servers also have their own "local" groups that should be audited. Users are added to existing and newly created groups on a daily basis. Without an audit trail of who is being added and removed from these groups, enforcing access control processes is impossible. Ideally group membership should be authorized and reviewed by the owner of the data or resource to which the group provides access.

9. LOCK DOWN, DELETE, OR ARCHIVE STALE, UNUSED DATA

Much of the data contained on unstructured and semi-structured platforms is stale. By archiving stale or unused data to offline storage or deleting it, IT reduces risk that stale data will be accessed by inappropriate parties, and makes the job of managing the remainder simpler and easier, while freeing up expensive resources.

10. CLEAN UP LEGACY GROUPS AND ACCESS CONTROL ARTIFACTS

Unneeded complexity slows down performance and makes mistakes more likely. Organizations often create as many groups as there are users, leaving many groups that are empty, unused or redundant. Some groups contain other groups, which contain other groups, with so many levels of nesting (that they sometimes create circular a reference when they contain a group that contains itself). Access control lists often contain references to previously deleted users and groups (also known as "Orphaned SIDS"). These legacy groups and misconfigured access control objects should be identified and remediated.

11. ONE MORE THING – GET CONTROL OF PUBLIC CLOUD SERVICES

With millions of users now using Dropbox and other public cloud collaboration services for work, organizations cannot only have data stored in repositories without controls or oversight, they run the risk of losing their data entirely. Users demand file synchronization with their laptops and workstations, mobile device support, and an easy way to share files with third parties. Organizations either need to choose a sanctioned, private cloud service that meets organizational compliance and security requirements, or extend their existing infrastructure to provide the public cloud experience so that users are no longer tempted to bypass IT policies and infrastructure, and continue to collaborate using their organizations' controlled infrastructure.

11 IT MUST-DO'S? NO PROBLEM WITH VARONIS

Varonis DatAdvantage automates these 11 IT must-do's. DatAdvantage delivers the visibility and automated auditing you need to determine who can access your unstructured data, who is accessing it, who should have access, and what is likely to be sensitive.

Continuously updated information drawn directly from your environment (with no performance impact for your servers) showing you the individual users and the groups they are part of, every folder on your file and SharePoint servers, every mailbox and public folder on your Exchange servers, and each data access — open, delete, rename, mail sent received, etc. — for every user. All permissions and group changes are logged, and can be sent to the data-owner for initial approval and/or review.

Click on a folder, site, or mailbox to see exactly who has access to it, what type of access they have — read, write, execute, etc., and where their permissions came from. Varonis DatAdvantage shows you detailed data access behavior and makes recommendations about whose access can be safely revoked. Once the owner is identified using the access activity and analysis in DatAdvantage, the owner can be automatically involved in authorization decisions and reviews via DataPrivilege.

ABOUT THE VARONIS METADATA FRAMEWORK™

Ongoing, scalable data protection and management require technology designed to handle an ever-increasing volume and complexity—a Metadata Framework.

Four types of metadata are critical for data governance:

- User and Group Information – from Active Directory, LDAP, NIS, SharePoint, etc.
- Permissions Information – knowing who can access what data in which containers
- Access Activity – knowing which users do access what data, when and what they've done
- Sensitive Content Indicators – knowing which files contain items of sensitivity and importance, and where they reside

The Varonis Metadata Framework™ non-intrusively collects this critical metadata, generates metadata where existing metadata is lacking (e.g. its file system filters and content inspection technologies), pre-processes it, normalizes it, analyzes it, stores it, and presents it to IT administrators in an interactive, dynamic interface. Once data owners are identified, they are empowered to make informed authorization and permissions maintenance decisions through a configurable web-based interface—that are then executed—with no IT overhead or manual backend processes.

The Varonis Data Governance Suite will scale to present and future requirements using standard computing infrastructure, even as the number of functional relationships between metadata entities grows exponentially. As new platforms and metadata streams emerge, they will be seamlessly assimilated into the Varonis framework, and the productive methodologies it enables for data management and protection.

VARONIS DATA GOVERNANCE SUITE

Varonis provides a complete metadata framework and integrated product suite for governing unstructured data on file servers, NAS devices, Exchange mailboxes and (semi-structured) SharePoint servers. Varonis DatAdvantage, DataPrivilege, and the IDU Classification Framework provide organizations the ability to effectively manage business data through actionable intelligence, automation of complex IT tasks, and sophisticated workflow management.

VARONIS DATADVANTAGE FOR WINDOWS

VARONIS DATADVANTAGE FOR UNIX/LINUX

VARONIS DATADVANTAGE FOR SHAREPOINT

VARONIS DATADVANTAGE FOR EXCHANGE

VARONIS DATADVANTAGE FOR DIRECTORY SERVICES

DatAdvantage provides a single interface through which administrators can perform data governance activities.

VISIBILITY

- Complete, bi-directional view into the permissions structure of unstructured and semi-structured file systems:
- Displays data accessible to any user or group, and users and groups with permissions to any folder or SharePoint site
- User and group information from directory services is linked directly with file and folder access control data

COMPLETE AUDIT TRAIL

- Usable audit trail of every file touch on monitored servers
- Detailed information on every file event in a normalized database that is searchable and sortable
- Data collection performed with minimal impact to the file server and without requiring native Windows or Unix auditing

RECOMMENDATIONS AND MODELING

- Actionable intelligence on where excess file permissions and group memberships can be safely removed without affecting business process
- Model permissions changes without affecting production environments

DATA OWNERSHIP IDENTIFICATION

- Statistical analysis of user activity effectively identifies business owners of data
- Automated reports involve data owners in data governance processes
- Facilitates round-trip data owner involvement via DataPrivilege

VARONIS DATAPRIVILEGE

DataPrivilege automates data governance by providing a framework for users and data owners to be directly involved in the access review and authorization workflows. A configurable web interface for data owners, business users, and IT administrators automates data access requests, owner and IT authorization of changes, automated entitlement reviews, and business data policy automation (e.g. ethical walls). A complete audit trail ensures that data governance policies are in place and being adhered to.

AUTOMATED ENTITLEMENT REVIEWS

- Data owners are provided scheduled entitlement reviews with recommendations for access removal (generated by DatAdvantage)
- Reviews can be scheduled based on business policy

ACCESS CONTROL WORKFLOW

- Users can request access to data and group resources directly, providing explanation and duration
- Data owners and other stakeholders are automatically involved in authorization process
- Permissions changes are carried out automatically once approval requirements are met
- Permissions revocations are carried out automatically on their assigned expiration

BUSINESS POLICY IMPLEMENTATION

- Multiple levels of authorization provide automated implementation of business and IT data governance policy
- Ethical wall functionality enforces data access policies

COMPLETE SELF-SERVICE PORTAL

- Data owners can view and manage permissions on their data and groups without requiring elevated access privileges, if desired
- Data owners can view access activity and statistics about their data, if desired

COMPLETE AUDIT TRAIL AND REPORTING

- All workflow events are recorded for audit and reporting which can prove the enforcement of governance practices
- Authorizations, entitlement reviews, and other management reports provide evidence of process adherence

VARONIS IDU CLASSIFICATION FRAMEWORK

The Varonis IDU Classification Framework gives organizations visibility into the content of data, providing intelligence on where sensitive data resides across its file systems. By integrating file classification information—from either the included classification engine or from a third-party classification product—alongside the rest of the Varonis metadata in the DatAdvantage interface, the IDU Framework enables actionable intelligence for data governance, including a prioritized report of those folders with the most exposed permissions and containing the most sensitive data.

- Actionable Intelligence
- Classification information provides visibility into business-critical content from within the Varonis IDU
- Organizations can see where their most sensitive data is over-exposed along with actionable recommendations on where that access can be reduced

EXTENSIBLE ARCHITECTURE

- The provided data classification engine provides a powerful and flexible method for classifying sensitive data through regular expressions and dictionary searches
- The IDU Classification Framework can also integrate content classification data from third-party classification and DLP products, extending the ability of both
- Intelligent, fast
- True incremental scanning is attained with DatAdvantage real-time knowledge of all file creations and modifications–only new data is classified
- Produces rapid-time-to-value results that have a clear remediation path or “next step”
- Produces results dramatically faster than traditional approaches

LEVERAGES EXISTING INFRASTRUCTURE

- Can use either its built-in classification engine or those already deployed
- Uses the unique metadata layer created by the Varonis Intelligent Data Use (IDU) Framework
- Builds on the foundation of the Varonis IDU Framework, with no need for additional servers or storage
- Results flow into Varonis DatAdvantage and Varonis DataPrivilege (future)

EASY, POWERFUL CLASSIFICATION RULES

- Rules match a combination of content AND metadata conditions (e.g. creator, accessing user, permissions sets)
- Prioritization based on Varonis metadata (e.g. scan the most exposed folders first)
- Files are searched for keywords, phrases and/or regular expression patterns
- Dynamic/auto-updated dictionary matching capabilities

VARONIS DATA TRANSPORT ENGINE

With a choice between manual processes and primitive utilities, migrating and archiving data has long been a time-consuming nightmare for IT. We have always been able to describe quite clearly what we want to happen during a migration, but ensuring what we want to happen actually does happen has always required massive amounts of planning, testing, tweaking, verifying and finger-crossing.

By harnessing file system, permissions, access activity, and content metadata across UNIX and Windows file shares, SharePoint, and Exchange Mailboxes and public folders, the Varonis Metadata Framework provides critical intelligence make data migrations more efficient and more secure, such as which data is active or stale, which content may be sensitive or regulated, and which permissions may be excessive or broken.

Now, with an intelligent rules engine and scheduling mechanism, the Varonis Data Transport Engine (DTE) allows IT personnel to set dynamic criteria to identify the data that should be moved, where it should end up, when it should be moved, whether the permissions should remain effectively the same or be changed (for the better), and then executes the migration automatically—end to end. The Varonis Data Transport Engine automates all the heavy lifting: copying data and metadata while adhering to maintenance windows and other scheduling constraints, automatically synchronizing source and destination with incremental copies even if the source data is still “live,” translating permissions across platforms and domains, and reporting on progress every step of the way.

Because DTE is built on top of the Varonis metadata framework, you can make sure that all data is managed and protected, where only the right people have access, all use is monitored, and abuse is flagged—before and after a move. You’ll know which users are happily using your new server or platform to collaborate with their data, and which ones haven’t read the memo.

With the Varonis Data Transport Engine, IT finally has an intelligent system that can be told what an ideal migration looks like, and it will take care of all the scary details for you. Say goodbye to the weekend-shifts and all-nighters—just describe your ideal migration, simulate it before committing, and automation will make it happen quickly, and securely.

VARONIS DATANYWHERE

Varonis DatAnywhere extends the corporate infrastructure so that remote employees can access traditional LAN resources with the same robust experience provided by cloud-based file synchronization services. With DatAnywhere, employees can:

- Automatically and securely sync files from corporate file shares to and from their laptops (future: Smartphones and tablets)
- Authenticate with corporate directory services (e.g. Active Directory)
- Securely share files with other employees and external business partners

Varonis DatAnywhere proxies and streamlines access to existing corporate infrastructure, translating legacy LAN based protocols (e.g. CIFS) into a protocol optimized for today's remote employees, leveraging secure https communications and seamlessly integrating with corporate directory services, access controls, and data protection and management processes.

Organizations can provide the cloud experience without moving their data from existing file servers, without reconfiguring access control lists and groups, and without new costs and management headaches from separate, incompatible infrastructure.

When combined with the rest of the Varonis Data Governance Suite, organizations have a complete solution to automatically optimize access controls involving the correct data owners, audit all activity, and flag suspicious activity across all unstructured data, whether it is used by in house employees on the LAN or by remote ones using cloud-like services. Wherever it resides, regulated or sensitive content can be identified and safely locked down.

ABOUT VARONIS

Varonis is the leading provider of software solutions for unstructured, human-generated enterprise data. Varonis provides an innovative software platform that allows enterprises to map, analyze, manage and migrate their unstructured data. Varonis specializes in human-generated data, a type of unstructured data that includes an enterprise's spreadsheets, word processing documents, presentations, audio files, video files, emails, text messages and any other data created by employees. This data often contains an enterprise's financial information, product plans, strategic initiatives, intellectual property and numerous other forms of vital information. IT and business personnel deploy Varonis software for a variety of use cases, including data governance, data security, archiving, file synchronization, enhanced mobile data accessibility and information collaboration.

Free 30-day assessment:

WITHIN HOURS OF INSTALLATION

You can instantly conduct a permissions audit: File and folder access permissions and how those map to specific users and groups. You can even generate reports.

WITHIN A DAY OF INSTALLATION

Varonis DatAdvantage will begin to show you which users are accessing the data, and how.

WITHIN 3 WEEKS OF INSTALLATION

Varonis DatAdvantage will actually make highly reliable recommendations about how to limit access to files and folders to just those users who need it for their jobs.

START YOUR FREE TRIAL

WORLDWIDE HEADQUARTERS
1250 Broadway, 31st Floor, New York, NY 10001 **T** 877-292-8767 **E** sales@varonis.com **W** www.varonis.com

UNITED KINGDOM AND IRELAND
Varonis UK Ltd., Warnford Court, 29 Throgmorton Street, London, UK EC2N 2AT **T** +44 0207 947 4160 **E** sales-uk@varonis.com **W** www.varonis.com

WESTERN EUROPE
Varonis France SAS, 13-15 rue Jean Jaures (1er Etage) 92800 Puteaux **T** +33 184 88 56 00 **E** sales-france@varonis.com **W** sites.varonis.com/fr

GERMANY, AUSTRIA AND SWITZERLAND
Varonis Deutschland GmbH, Welserstrasse 88, 90489 Nürnberg **T** +49(0) 911 8937 1111 **E** sales-germany@varonis.com **W** sites.varonis.com/de