

Whitepaper

Auditing Unstructured Data

Identity-Aware Storage, File Activity Monitoring, and Compliance Reporting Across Platforms

Table of Contents

Introduction	2
Why Audit Unstructured Data?	3
Architecture	7
File Activity Monitoring	10
Content and Context	11
Reporting	11
Conclusion	12
Do You Have an Auditing Problem?	13
A Ten-Step Program for IT Auditing Success	14

By **Steve Hoenisch**,
Likewise Software

Executive Summary

IT auditors face file servers that contain a rapidly growing amount of unstructured data, 40 percent of which tends to be sensitive information: intellectual property, confidential data, and company secrets.

Auditors' problems are made worse by a heterogeneous network in which the unstructured data of Unix and Linux users is typically stored separately from the data of Windows users, making it difficult to apply a uniform set of security policies for access control and to conduct audits across the systems.

Incompatible identity management systems further compound the problem, creating another obstacle to standardized audits that link users to identities. Meanwhile, compliance regulations, disclosure laws, and risk mitigation require security that identifies threats based on user identities and privileges.

This white paper argues that a multiprotocol file server or NAS system with an integrated cross-platform identity and access management service is the architectural basis for solving many problems in protecting and auditing unstructured data.

First, it frees you from the silos of platform-specific storage, enabling you to audit all the stored data without regard for storage protocol.

Second, it secures the unstructured data by applying a common security model to it, enabling your audits to associate access with user identities and organizational roles.

Third, it establishes the foundation for a high-performance file activity monitoring system that can audit unstructured content in a security-aware context of user identities, patterns of access, and file change events.

The result is an identity-aware, cross-platform storage system that makes it easy to secure unstructured data from internal threats, monitor user access, track changes to sensitive files, and generate reports that demonstrate regulatory compliance with evidence.

Introduction

At the center of an IT auditor's complex matrix of compliance regulations, legal requirements, and internal security policies stands a mountain of unstructured data – files such as spreadsheets and documents that frequently contain sensitive information.

Unstructured data, industry analysts say, is growing faster than all other types of data and will increase by as much as 800 percent during the next five years. Respondents to an industry survey by the Aberdeen Group estimated that 40 percent of their sensitive data is in unstructured formats such as PDF and Microsoft Office files like Word and Excel.

To protect unstructured data, the Aberdeen Group recommends that you prioritize security control objectives for unstructured data “as a function of risk, audit, and compliance requirements” and that you standardize audit, analysis, and reporting.

Compliance regulations and disclosure laws for unstructured data are increasingly pointing toward the need for systems that can secure unstructured data by identifying and remediating threats based on user identities and privileges.

Frequently, however, the file servers and network attached storage systems holding sensitive unstructured data lack integration with an identity management system. When an identity management system is tied to the file server, it can control access based on user privileges while establishing a powerful identity-aware auditing and reporting framework.

The complexity of heterogeneous networks and storage systems that mix Windows computers and Unix machines further complicates the issue. There are NFS-based file servers for Unix users and CIFS-based file servers for Windows users. The inability to interoperate between the two protocols turns cross-platform storage into a complicated network of mixed systems containing duplicated data – systems that are nearly impossible to audit with a common auditing and reporting system. Ad hoc systems that add an auditing layer for file events frequently result in performance issues and fail to link file events to user identities.

This white paper describes a number of problems that make it difficult to audit unstructured data, lays out a set of key system requirements for IT auditors, and prescribes a solution to help securely manage sensitive unstructured data so you can audit it with ease.

Why Audit Unstructured Data?

The main reasons that IT auditors and records managers undertake initiatives to audit unstructured data are typically as follows:

- Demonstrate compliance with regulations, legal requirements, industry standards, and internal security policies.
- Store and protect company secrets with a data management policy.
- Mitigate the risk of security breaches, data loss or exposure, fraud, noncompliance, and legal problems.

Demonstrating compliance tends to be a key objective, leading to the deployment of security information and event monitoring tools (SIEM). SIEM tools, however, can fall short in closing the compliance gap because they likely lack integration with an identity management system. As a result, the tools cannot monitor user access and activity and detect exceptions based on one of the most important security factors, an authenticated identity.

Stored secrets, meantime, are a hidden problem waiting to surface. If a competitor obtains documents containing intellectual property, for instance, it could hurt your business prospects or undermine your competitive advantage. The goal is to control the secrets so they don't get into the wrong hands.

The practice of records management provides an additional motivation to control unstructured data, ensuring that it can be audited later. You might need to store unstructured data for years for such reasons as complying with government regulations and obtaining patents.

The need to audit unstructured data to comply with regulations and to keep secrets gives way to a final, overarching objective: mitigating risk. Auditing can detect potential sources of data loss, fraud, inappropriate entitlements, access attempts that should not occur, and a range of other anomalies that are indicators of risk – especially when the audit can associate data access with user identities.

Problems in Managing and Auditing Unstructured Data

The following problems, use cases, and requirements should be considered before implementing a system to audit unstructured data. Most of the problems and example use cases stem from the need to conduct risk assessments and control

evaluations based on a model of information security management or auditing, such as the infosec triangle, that takes into account confidentiality, availability, and integrity. Some infosec models are extended to include authentication and access, privacy, ownership and distribution, data retention, and auditability. Additional problems arise from the four basic tasks of information life cycle management (ILM): identify, collect, classify, and control information.

- **Locating unstructured data that contains sensitive or confidential information.** The data is typically distributed throughout the enterprise and segmented in disparate file servers. Here's an example use case: There are files that you believe contain highly sensitive proprietary information – engineering plans – on the engineering department's file share, but since the file share is a Unix system using the NFS protocol, you cannot access it from your Windows system, which is using the CIFS protocol. Even though you may have rights to access the data, your access is effectively blocked and you cannot determine whether the file server contains sensitive data, let alone audit it to ensure compliance with internal policies that mandate the protection of intellectual property.
- **Collecting unstructured data after it is identified.** Example use case: If your data is distributed among multiple file servers and you want to consolidate it into centralized storage, you might have to do it manually – but you must still allow it to be accessed by both Unix and Windows users, who are controlled by multiple identity management systems.
- **Classifying directories and files as containing sensitive information.** Example: As an IT auditor at a large police department, you need to be able to mark as sensitive the directories on a file server that contain the incident reports of police officers so you can secure access to them and audit them later. The reports frequently contain personal information about citizens who have not been arrested or charged with a crime.
- **Determining who can access sensitive data, determining whether the set of users and groups that can access it is the right set, and linking user identifiers with people and their organizational roles.** Example: As an auditor, you want to be able to generate a report that lists the identities and entitlements of those who can access sensitive data on a file server, and you want to be able to map their identities to their organizational roles so you can vet it for inappropriate access rights.
- **Controlling who can access the data and make changes to it by linking identity to the entitlements that allow users or groups**

to access and potentially change directories and files – and doing so regardless of whether the users and groups are from the Unix or the Windows side of the shop. Example: As a records manager, you must be able to select who can access what sensitive data and you must be able to modify their access rights when their roles change or the content transitions to a different state.

- **Showing what directories and files were changed when and by whom, especially in the face of identities that are inconsistent across access control systems.** Example: As an auditor, you cannot correlate the identities of some Windows users with the identities of some Unix users, even though you believe they might be the same people.
- **Tracking and correlating changes to files, especially ones marked sensitive, in a way that can show exceptions that might indicate inappropriate access.** Example: As the security manager at your company, you want to make sure that sensitive files are protected from access by those without permission to view them.
- **Generating a security alert when sensitive files or folders are accessed or modified.** Example: As the records manager charged with storing and protecting sensitive files, you want an email alert when certain files are accessed or modified.
- **Generating an audit trail to show who accessed which records.** Example: As a compliance officer managing the physical therapy department at a large hospital that recently began storing the records of physical therapy patients in electronic form, you must be able to prove under HIPAA that only staffers with a need to know can access the records. You must be able to show that access is limited to individuals involved in a patient's care and if challenged you must be able to present an audit trail showing access.
- **Generating reports to fulfill specific compliance requirements, such as those of HIPAA, HITECH, ITAR, PCI DSS, Sarbanes-Oxley, GLBA, or FISMA.** In addition, you must be able to create custom reports, including exception-based reports, to meet specific internal policy requirements. Example: As an auditor at a U.S. defense company, you

know that the International Traffic in Arms Regulations, or ITAR, dictates that material pertaining to defense and military technologies may be shared only with U.S. persons unless exempted or specially authorized by the Department of State. You can face legal fees and heavy fines if a foreign employee views ITAR-protected information, and you must be able to report on access rights to certain file servers in a way which shows that some of your employees, who are foreign nationals, cannot view ITAR-protected data.

- **Showing chain of custody over tracked files, tracked directories, data on file access events, and data on file change events.** Example: To comply with Section 404 of Sarbanes-Oxley, you must be able to demonstrate that important documents – minutes of board meetings, financial reports, bank records, and so forth – are genuine and have not been altered. In short, as the auditor, you must be able to design and implement a control, prove its operational effectiveness, and monitor it to make sure that it continues to work as designed.
- **Produce reports over which you have chain of custody.** Example: As an auditor, you must be able to show that the data used in your reports as well as the reports themselves have not been tampered with.
- **Retaining and archiving records by, for instance, setting expiration times and similar conditions to trigger secure archiving and, eventually, deletion.** Example: As the records manager at your company, you want to trigger certain directories on your file server to be automatically archived at a set date and then deleted in seven years.
- **Inspecting application data on file servers.** Example: As an IT auditor, you find out that your IT department is increasingly migrating application workloads to filers because it is easier to provision file-based storage than block-based storage. As a result, you have a heightened need for exception-based methods of auditing the application data.
- **Providing context-aware security where the context is the intersection of content, event, access, and identity.** Example: As an IT security officer in charge of compliance at a defense company, you seek to design a collaboration hub on a file server for sharing classified ITAR-controlled documents and project information. You must be able to control who can access and view which content, see who changes what, monitor for exceptions, and produce reports to demonstrate compliance with ITAR. You must be able to prove that employees who are foreign nationals cannot view the documents.

There are some additional requirements that either fall out from the use cases above or are more general in nature. You will, no doubt, have other requirements to deal with your unique problems in complying with regulations, showing adherence to internal policies, and managing risk. Some of the following general requirements ensure that the infrastructure has the flexibility to conduct audits and generate reports in a way that fulfills a diverse, dynamic set of needs.

- Audit the data for patterns that would indicate the presence of confidential information like Social Security numbers and credit card numbers.
- Create custom reports and exception-based reports. You should, for example, be able to formulate your own queries, including searches using Boolean operators to formulate logical statements of inclusion, exclusion, and so forth.
- Audit the data for inappropriate access by analyzing events in which access or modification attempts are denied.
- Audit changes to sensitive or tracked files, such as modifications of content or security descriptors as well as attempts to delete files or content. This requirement is in effect file activity monitoring, or FAM, and can help comply with the file integrity monitoring stipulated in PCI DSS requirement 11.5 – raising an alert for unauthorized changes to content files.
- Use an analytics engine to discover hidden patterns that could detect fraud or reveal internal threats.
- Have an architecture that is flexible enough to accommodate complementary technologies for data loss prevention, such as interoperable storage encryption and document archiving.

Solving the problems discussed above and fulfilling these requirements by positioning content in its rightful security context, however, raises the following question: What kind of architecture for a file server would make identity-aware auditing a reality with a minimum of complexity?

Architecture

The following components provide the architecture for a file server that supports a universal approach to auditing unstructured data.

- A multiprotocol, cross-platform file server or NAS system that supports CIFS and NFS to allow connections from both Windows and Unix computers.

- An integrated authentication engine that can authorize users with Active Directory, NIS, or LDAP.
- An integrated application for marking and tracking sensitive folders and files.
- A secure event monitoring subsystem with collectors and forwarders that record, manage, and transmit file activity events.
- A NoSQL database for event processing and advanced analytics.
- A SQL data store for reports.
- An auditing and reporting console.
- An events dashboard.

Multiprotocol File Server Accessible by Windows and Unix

At the foundation is a file server that is multiprotocol and cross-platform: It supports both the SMB/CIFS and the NFS protocols, making it usable simultaneously by Windows and Unix or Linux clients. A cross-platform, multiprotocol file server solves the interoperability problem that often separates the data of Unix users from the data of Windows users, providing a consolidated approach to storage for users of all types of computers. As an auditor, it frees you from having to conduct separate audits on the unstructured data that resides in silos differentiated by operating system.

Cross-platform incompatibilities have also been a hindrance to applying a uniform set of security policies. In the past, just as there have been different, incompatible storage systems for Unix and Windows users, there have also been different, incompatible identity management systems for Unix and Windows users. Unix clients have tended to use NIS or LDAP, while the de facto standard for Windows clients is Microsoft Active Directory.

In this architectural schematic, therefore, the file server includes an integrated identity management service to authenticate users with Active Directory, NIS, or LDAP – a component that, when combined with the multiprotocol file server, lays the architectural foundation for solving many of the problems in protecting and auditing unstructured data.

The overall result is twofold. First, it frees your users from the bounds of platform-specific storage, enabling you to audit all the stored data from a single system. Second, it secures the unstructured data by applying a common security model to it, enabling your audits to associate data access with user identities and roles.

Secure Cross-Platform Access Control for Unstructured Data

The integrated identity service delivers direct, authoritative, robust security to control access by user or group, including blocking all external users and allowing only those internal users and groups that you specify. Simply put, you can control access to sensitive unstructured data and, as described below, use the built-in auditing framework to demonstrate those controls for compliance.

The file server's tight integration with the identity service also gives you visibility into the entitlements and permissions that are used to access and modify files. Because the identity service is integrated with the file server, you can both secure access by entitlement and generate reports to prove the entitlements secure the data.

Classify and Track Sensitive Files Tied to Identities and Owners

The integrated identity service lets you mark sensitive files, associate them with the identities of their owners, and track changes by user or group. Records managers who are charged with managing confidential information in unstructured files can limit access to specific users and groups and map changes to the files to those users and groups.

Collect Access Data and File Events for Analytics and Reports

The event collectors and forwarders form the event monitoring subsystem. On the file server, the event collectors record information about moving, copying, reading, modifying, or deleting directories or files. The collectors also capture changes to security descriptors.

Over a secure connection, the event forwarders send the file events on to the NoSQL database – which is the basis for a powerful, flexible analytics engine that can correlate content types, sensitivity levels, modification attempts, security descriptors, user entitlements, and access patterns.

Furthermore, an analytics system can use data about past access patterns and file activities to hypothesize about future patterns of data storage. The inferences of an analytics system can help identify files that might contain sensitive data and need to be flagged for inspection or tracking.

The NoSQL system, meanwhile, interfaces with a SQL Server database that segments frequently used data into columns and rows for reports, including custom queries.

The auditing and reporting console can be used to create custom reports or reports based on templates to fulfill compliance regulations such as SOX, HIPAA, PCI, and FISMA. You can choose your own data points. A simple interface makes it easy to

create new reports, including custom reports. The reports let you audit file access and events by directory or server.

Meanwhile, for threat monitoring, the dashboard displays file events correlated with permissions in near real-time so you can proactively monitor user access and changes to sensitive files and respond to policy violations, potential breaches, or other security incidents.

Performance

Millions of file events can easily overwhelm the network and the monitoring system. Because of the number of events that are generated in an enterprise as a multitude of users access and modify files, performance is a requirement that must be considered up front – but all too frequently is not, and it is only after implementation that performance issues emerge: networks slow down, databases overwhelm disk space, dashboards freeze.

The performance of the event monitoring system plays a key role in how fast and efficiently many of the end-user components that rely on the events will function. To be expedient and relevant, exception monitoring depends on how fast events are collected and correlated. The auditing tools also rely on the performance of the system to quickly produce up-to-date reports.

To ensure that events do not consume too much network traffic or bog down systems, monitoring ultimately should take place as part of the file server. When the monitoring is handled by the file server and is built with performance in mind, it can ensure that the system scales efficiently to deliver high performance in high-traffic environments.

File Activity Monitoring

In the architecture outlined above, the event monitoring subsystem makes possible a file server with integrated high-performance file activity monitoring, or FAM. Similar to database activity monitoring, FAM refers to an emerging tool set that can help identify and report on file access patterns that could be noncompliant, fraudulent, or illegal.

File activity monitoring is at its most powerful when it is tied to identity management (IAM). The integration of the identity monitoring system with the activity monitoring system is a precondition for exception monitoring – a highly efficient and effective auditing method that takes place at the nexus of user activity and access to resources.

Increasingly, industry analysts report that auditors are looking at database activity monitoring tools to comply with regulations and to manage security risks associated

with structured data stored in databases. Doing so, however, puts the security and auditing focus on structured data without placing a corresponding emphasis on rapidly growing file repositories. Thus, FAM technologies should likewise be evaluated and implemented to perform the same auditing functions – only in relation to unstructured data.

Content and Context

The importance of file activity monitoring highlights the shift in IT toward contextualized security – in this case, viewing content in the context of identity, entitlements, access patterns, sensitivity levels, file events, and other factors related to security.

When identity, access, content, and events are tracked at the file server, file activity monitoring is enriched by contextualized security data – the correlations that take place at the intersection of users with known roles and entitlements accessing tracked content to perform logged events.

As an auditor, the result is that you can audit the data in context to produce detailed reports and then use the information in the reports to lobby for identity and entitlement changes. As a records manager, the result is that you can receive an email alert when sensitive files or folders are accessed or modified.

Using Big-Data Analytics to Mitigate Risks

In the events that are generated when you track content in the context of identity and access, there lies a huge amount of data that describes patterns of access, activity, and change – data that becomes an input to an important use that progressive auditors can exploit to mitigate risk in the future: Analytics.

In an enterprise environment with 50 million objects stored across a 25-node array, for example, more than 2 million objects can be modified a day, with the number of events for access attempts and file views being much higher.

An analytics system can use the data about past access patterns and file activities to hypothesize about future patterns of data storage. Such inferences can help identify files that might contain sensitive data and need to be flagged for inspection or tracking. The data can be correlated in unexpected ways to produce innovative results.

Reporting

Finally, reporting can help mitigate security risks, identify security vulnerabilities before they are exploited, inspect access rights, show patterns of access and change, and double-check levels of protection – all of which can help prove compliance with regulations such as PCI, SOX, and ITAR.

Yet many organizations lack reports tied to security information and event monitoring (SIEM) tools. Even fewer organizations have integrated their reporting and auditing tools with their identity management and access control systems. Fewer still move beyond reports to use dashboards to monitor correlated file server events in real-time for exceptions.

For many regulations, the reporting system as well as the reports that are generated must be secured with access control, typically as part of a policy that addresses information security for all personnel. A reporting system that is integrated with the identity management system allows you to do so, effectively enabling you to show change logs and chain of custody not only for sensitive data but also for the reports themselves.

But reporting and auditing is not just about proving compliance, it's also about cutting costs. According to the Aberdeen Group, "The greatest financial gains for best-in-class organization will come from automating the enforcement of policies whenever reasonable, standardizing audit, analysis, and reporting, and driving continuous improvements by finding and eliminating root causes for exceptions, security events, and audit deficiencies."

Standardizing and automating reports at the confluence of storage, identity, and access radically improves visibility to possible data breaches, security threats, and compliance failures that expose your organization to risk.

Conclusion

A multiprotocol file server or NAS system that includes an integrated cross-platform identity management service to authenticate users and control access provides the architectural basis for solving many of the problems in protecting and auditing unstructured data.

First, it frees you from the bounds of platform-specific storage, enabling you to audit all the stored data without regard for storage platform or storage protocol.

Second, it secures the unstructured data by applying a common security model to it, enabling your audits to associate data access with user identities and roles.

Third, it establishes the foundation for a powerful high-performance file activity monitoring system to audit unstructured data in a security-aware context of user identities, patterns of access, and file change events.

The result is an identity-aware storage system that makes it easy to secure unstructured data from external and internal threats, monitor user access, track changes to sensitive files, and generate reports that demonstrate regulatory compliance.

Do You Have an Auditing Problem?

Many people jump into technology discussions without fully understanding the problems they are attempting to solve. Here's a list of questions to ask yourself to help determine whether your organization might have an IT auditing problem.

Before looking at the list, consider your compliance obligations for unstructured data – obligations that will vary by a number of factors, including the size of your company or organization, the industry you are in, the applicable regulations and laws, and so forth.

In general, ask yourself a couple of overarching questions: First, is sensitive unstructured data consolidated in a file server or storage system? If your unstructured data is not on a file server, what conditions would have to be met to place it there? Second, what assistance can an internal auditor provide to help meet your compliance obligations? The auditor, for example, may be able to develop compliance requirements that summarize your organization's regulatory, legal, and internal obligations.

- Have you had security breaches or data losses, whether internal or external, or other security-related incidents with regard to stored unstructured data?
- Have you had incidents of non-compliance? Do you have trouble producing information to comply with external requirements such as PCI DSS, ITAR, HIPAA, or SOX?
- Have you had incidents around unstructured data that violated internal security policies or procedures? Would you know if you did?
- Can you locate your the sensitive directories and files amid the mass of unstructured data on your file servers and tag the data for tracking so you can audit it later? Can you categorize your files by department and set their sensitivity level, both by directory and by file?
- Can you regularly audit patterns of end-user access to sensitive unstructured data and report on them?
- When you do identify an audit deficiency, how much time and money do you spend trying to address it?
- Do you have a records management system in place for managing sensitive unstructured data? If so, can you audit the data in it?
- Do you have a scalable mechanism by which you can audit the unstructured data on all your file servers, whether Windows or Linux?
- Can you identify who changed or deleted which sensitive files when?
- Can you link user and group accounts to IP addresses, departments, individuals, roles, and entitlements?
- Can you produce historical records that show who changed what when? Can you submit custom queries to generate reports for external auditors?
- Can you show chain of custody not only for your sensitive files but also for their event data and the reports that rely on the data?
- Can you set retention periods for sensitive files, archive older files, or set a sensitivity level to revert to non-sensitive after a set period of time?
- Do you have reports that correlate the roles and access privileges of users with their file access and attempted modifications?
- Can you conduct contextualized audits that take into account multiple planes of data: users, groups, entitlements, and business roles; patterns describing where, when, and how data is accessed; content owners, file types, sensitivity levels, permissions, and security descriptors; and types of operations carried out against the content, such as copying, modifying, and deleting? Can you control the correlation of these planes of data to get the auditing results you want?
- Have you standardized your auditing and reporting? In its survey on Securing Unstructured Data, the Aberdeen Group writes: "The standardization of audit, analysis, and reporting is an emerging capability with respect to unstructured data. Just 29% of Best-in-Class organizations indicate this as a current capability." The Aberdeen Group adds: "Standardization in this area improves visibility, provides a common point of reference, and reduces the ongoing cost of operations compared to non-standard, ad hoc methods."

The downstream consequences of an IT auditing problem can be legal fees, costly compliance violations, and reputation-damaging data breaches.

Consolidating storage of sensitive unstructured data to file servers governed by a common security model and access control system can help identify who has access to what data, providing you with a framework within which you can better regulate access at a granular level and audit the data to demonstrate compliance.

A Ten-Step Program for IT Auditing Success

Mismanagement of unstructured data can put your reputation at risk, lead to privacy violations, and result in incidents of noncompliance. In large organizations, administrators and computer users are frequently unaware of regulatory requirements for sensitive data. Unless automated systems are put in place to force adherence and to monitor for lapses, users will inadvertently subvert those requirements.

Here's a ten-step program to organize, protect, and audit your unstructured data.

1. Identify your IT auditing and reporting requirements in relation to compliance regulations, disclosure laws, privacy laws, industry standards, and internal security policies.
2. Find your secret, toxic, confidential, and otherwise sensitive unstructured data.
3. Consolidate your sensitive unstructured data to a cross-platform, cross-protocol, high-performance file server or NAS system in the data center that can be accessed by Windows as well as Unix, Linux, and Mac OS X clients.
4. Integrate the file server with an identity management system that can provide cross-platform access control to enforce the same security model for Windows, Mac, Linux, and Unix users.
5. Implement an identity-aware security incident and event monitoring tool, or SIEM, to monitor access to sensitive data and produce exception reports that link access to identities, roles, and privileges.
6. Monitor file activity so you can audit data for file events such as modifications and associate those file events with user identities.
7. Tightly integrate the monitoring system with the file server and the identity management system to ensure scalability. The monitoring system must perform well even in enterprises with heavy network traffic and a deluge of user activity.
8. Make sure the monitoring system includes a dashboard that can display near real-time security events and exceptions.
9. Include a module to generate compliance reports that provide evidence during audits. Make sure the reporting system can create custom reports as well as reports from predefined templates for regulations like Sarbanes-Oxley, PCI DSS, HIPAA, FISMA, and ITAR.
10. Use an analytics engine to aggregate all the monitoring data so you can look for new patterns to improve auditability. The analytics engine empowers you to audit the events in new ways to detect aberrant patterns and find innovative ways to mitigate risk.

Next Steps

For more information on Likewise, visit the web site at www.likewise.com. To contact the sales team, call (800) 378-1330 or email info@likewise.com.

About Likewise Storage Services

Likewise Storage Services provides a platform for cross-platform network access to files used in OEM storage products built on Linux- and Unix-based devices. Whether you're building a cloud storage offering, cloud gateway, a traditional NAS device, or another application where you need to provide secure access to files across a network, Likewise Storage Services can help. To learn more visit www.likewise.com/products/likewise_storage_services.

About Likewise

Likewise makes an integrated software platform, Likewise Storage Services, for identity, security and storage used by market-leading OEM storage vendors including Riverbed, EMC and HP. In addition, Likewise Data Analytics and Governance is an application which helps enterprise IT organizations mitigate risk and drive greater value from their unstructured data. Likewise Data Analytics and Governance ties identity and other contextual data to unstructured data as metadata for better analytics, governance and compliance, entitlement management, and performance management. Likewise enables organizations to provide both access to and control of their data across mixed network environments. More information is available at the company's website, www.likewise.com.

The information contained in this document represents the current view of Likewise Software on the issues discussed as of the date of publication. Because Likewise Software must respond to changing market conditions, it should not be interpreted to be a commitment on the part of Likewise, and Likewise Software cannot guarantee the accuracy of any information presented after the date of publication. The contents herein are subject to change without notice.

These documents are for informational purposes only. LIKewise SOFTWARE MAKES NO WARRANTIES, EXPRESS OR IMPLIED.

Complying with all applicable copyright laws is the responsibility of the user. Without limiting the rights under copyright, no part of this document may be reproduced, stored in, or introduced into a retrieval system, or transmitted in any form, by any means (electronic, mechanical, photocopying, recording, or otherwise), or for any purpose, without the express written permission of Likewise Software.

Likewise may have patents, patent applications, trademarks, copyrights, or other intellectual property rights covering subject matter in this document. Except as expressly provided in any written license agreement from Likewise, the furnishing of this document does not give you any license to these patents, trademarks, copyrights, or other intellectual property.

© 2011 Likewise Software. All rights reserved.

Likewise and the Likewise logo are either registered trademarks or trademarks of Likewise Software in the United States and/or other countries. All other trademarks are property of their respective owners.

Likewise Software
15395 SE 30th Place, Suite #140
Bellevue, WA 98007
USA