WHITE PAPER

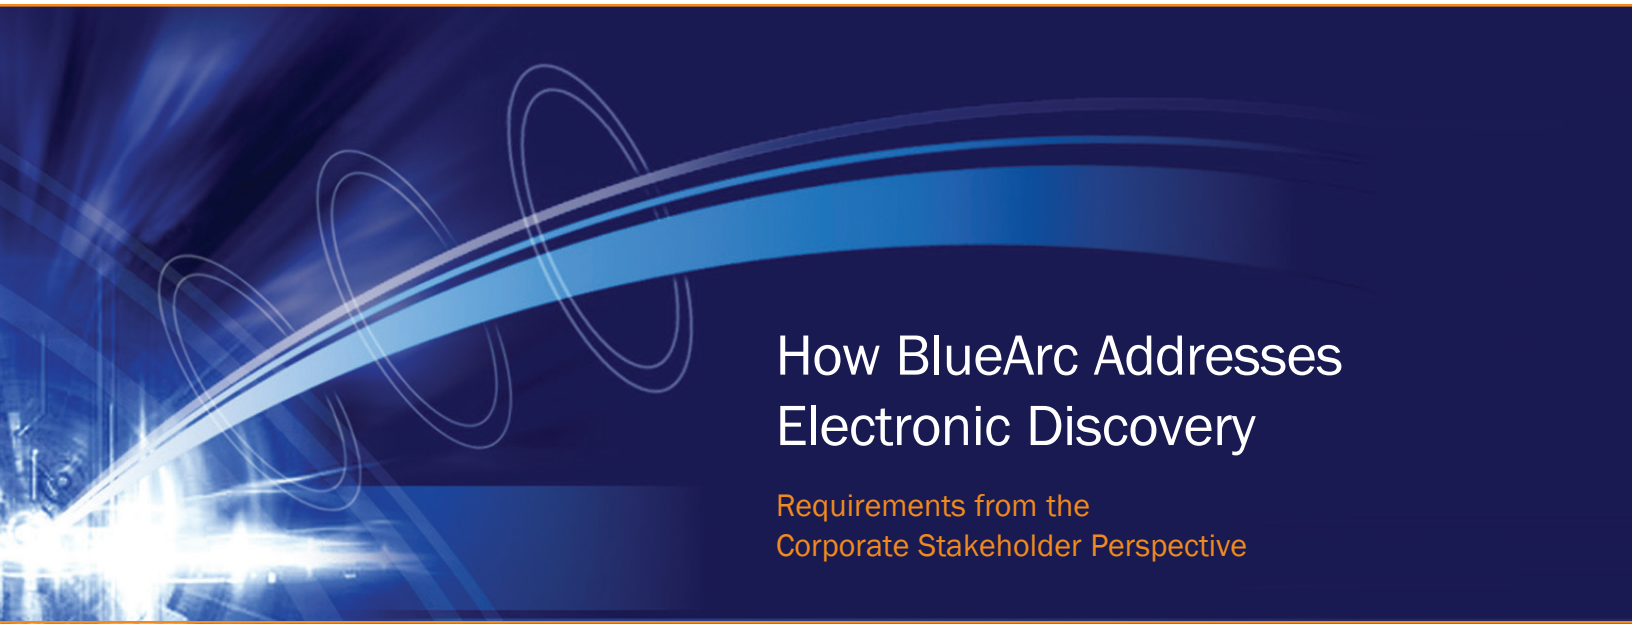

How BlueArc Addresses Electronic Discovery

Requirements from the Corporate Stakeholder Perspective

Abstract

Corporate governance-centric value propositions inherent in data management technology expenditures apply more to storage related expenditures than perhaps any other area of IT expenditure. This is so because of the heightened desire of organizations across the globe, to insulate themselves from the negative impact of out of control data management costs and the associated risks that can accrue. The mandate and responsibility for implementing risk management initiatives is shifting more into the CFO's realm whose role, as alluded to earlier, is being hybridized to incorporate more influence over IT and internal legal department strategies. Indeed, just as CFO's are becoming increasingly aware of the impact data management tools and infrastructure can have on mitigating organizational risk, CIOs, CTOs, and technologists are increasingly being pressured to more effectively quantify soft costs and discuss strategies that address data management risks and respective legal ramifications other stakeholders, like.

BLUEARC®

Introduction

When it comes to litigation discovery, compliance, records management and internal investigations, the operative concepts that best embody the driving force behind the technology decisions made by CIOs, GCs and CFOs are cost containment and risk management.

Today's business environment is characterized by unprecedented competitive and economic pressures. With each passing month, these pressures continue to mount for corporations as they find themselves increasingly mired with regulatory controls while attempting to achieve operational efficiencies. As record fines herald decreased thresholds of liability for data management malfeasance and corporate governance breakdowns, corporate CFOs have never been more focused on the corporate governance ROI of their IT expenditures. Under particular CFO scrutiny are those corporate expenditures that have the potential to positively impact shareholder confidence and facilitate complex corporate data governance initiatives. In economic and risk management terms, this means more organizations are seeking to derive increased ROI and risk management benefits by making sure that new infrastructure investments leverage highly scalable and cost effective technologies that that can help shore up the top line, increase the bottom line and reduce overall exposure.

Leading the charge in the effort to tie infrastructure capital (cap-ex) and operational expenditures (op-ex) to risk management ROI is the corporate CFO. While there are multiple stakeholders driving or involved in any corporate risk management initiative, CFOs are increasingly taking center stage when it comes to blessing decisions on technology directions. This is the case even if an organization's CIO may have already given a particular IT infrastructure cap-ex or op-ex a nod. As more governmental entities promulgate legislation requiring more exacting controls of information domestically and internationally, particularly consumer related non-public information (NPI), the CFO's role will continue morphing to encompass more tech and legal decision making responsibility. We will see board level mandates coming down that will require organizations to support technology acquisition decisions with statements from CFOs validating clearly articulated and demonstrable value proposition in IT storage and related infrastructure expenditures. This rationale is driven in part by:

1. Exploding data volumes.
2. Cost of technology.
3. The need for greater control and access to data for risk management and profit center activities.
4. Compliance requirements.
5. Pecuniary damages – significant fines and loss of market goodwill.

Corporate governance-centric value propositions inherent in data management technology expenditures apply more to storage related expenditures than perhaps any other area of IT expenditure. This is so because of the heightened desire of organizations across the globe, to insulate themselves from the negative impact of out of control data management costs and the associated risks that can accrue. The mandate and responsibility for implementing risk management initiatives is shifting more into the CFO's realm whose role, as alluded to earlier, is being hybridized to incorporate more influence over IT and internal legal department strategies. Indeed, just as CFOs are becoming increasingly aware of the impact data management tools and infrastructure can have on mitigating organizational risk, technologists are increasingly being pressured to more effectively quantify soft costs and discuss strategies that address data management risks and respective legal ramifications other stakeholders, like.

Fortunately for the corporate CFOs of the world, BlueArc, developers of one of the world's most powerful and affordable storage systems, has framed the problem and developed the answer.

The Data Benefit and Risk Quandary Every Enterprise Faces

From a legal discovery and CFO corporate governance perspective, the implications of "data state" relates to whether information is "readily accessible or "inaccessible" respectively for the purpose of compliance, litigation or internal investigation purposes. In the realm of storage management decreased disk costs have lead to trends where organizations are storing more data on-line to facilitate greater data accessibility. This creates the need for heightened controls as it relates to policies around data, particularly in the areas of legal discovery and compliance. As technology introduces additional layers of complexity to the policy mix, the forms that data takes is also becoming an issue. Data in most corporate organizations is generally a mix of:

1. Structured data – databases.
2. Semi-structured – email systems.
3. Unstructured – loose electronic files.

Some technologies today are making the distinctions between readily accessible and inaccessible data quite muddled. Take for example Microsoft Sharepoint technology. It is perhaps one of the most rapidly growing collaborative workflow, document management and records management technology commercially available. This technology makes data far more malleable in terms of the forms in which data can be stored, used and extracted. While these benefits are enormous, the model adds additional many variables to the data management equation.

It is clear that there are certain complex data management technology factors that are outside of the direct control of corporate stakeholders, however the one significant variable that is within a CFO's direct control is speed of access to information in the storage environment. Simply put, speed of access to data is a process agnostic multiplier. This means that it does not matter what the activity is, if you can remove the data I/O bottlenecks to the storage devices on which enterprise data resides, the result will be that profit center activities, i.e., access to marketing and sales data and risk management processes, i.e., internal investigations / NPI data remediation efforts take place in significantly less time than before. This assumes of course that storage the hardware infrastructure that accelerates information access is in place.

Other factors that affect data risk and governance in the enterprise are the distinctions between roles people play in the organization. Corporate data users involved in profit center activities want on-demand rapid access to the organizations information which they ostensibly leverage to its benefit. Data management custodians whose responsibility it is to protect data for the benefit of the organization want to audit, back it up and protect it quickly and transparently. Profit center users are compensated largely on how they leverage information for the benefit of the organization as opposed to data risk mitigation. Conversely cost center personnel are compensated on their ability to protect data and provide the profit center with timely access information as well as indentify, quantify and reduce the risk exposure in data stores. These divergent role based views of data exist in just about every corporate enterprise environment and often cause a fair amount of internal data management and policy contention. A classic example of this contention scenario is where data management custodians (IT Administrators) limit or cap the amount and age of email and or files users can have resident on network storage while users scream and throw tantrums about how they need to have additional storage space.

The IT DataGovernance Matrix

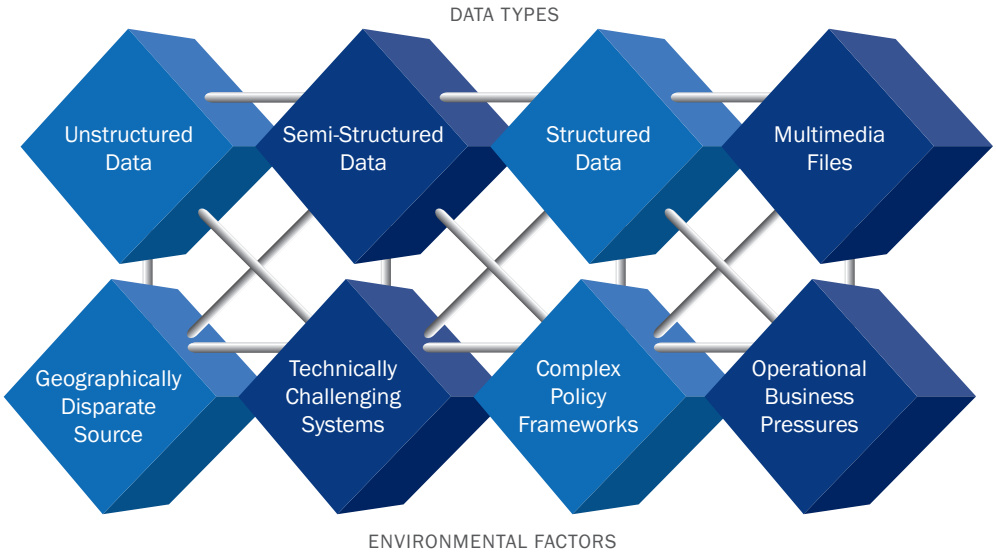

Figure 1

Figure 1. above, illustrates some of the high level environmental factors that create the data management challenges that exist in many corporate enterprises today. These challenges must be taken into consideration when it comes to deploying data storage platforms and associated data management technologies. Stakeholders, particularly CFOs, are being held increasingly accountable on the impact storage technology has on corporate governance and as a consequence they are demanding broad corporate governance applicability of technology deployed within their enterprise.

The ability to quickly pinpoint and access data both to leverage the value of information as well as ensure that it is subject to appropriate governance controls are steps that will keep the lights on in most corporate enterprises and keep C-level executives' legal fees to a minimum. Take for example concerns about corporate data that should be subject to various controls but remains unregulated within the environment. This "data below the radar" condition has an enormous potential to hurt organizations by damaging corporate goodwill as well as subjecting them to significant economic sanctions and fines. This situation has long been a concern for many C-level executives and the ability to assess and remediate offending conditions is a top management priority. Examples of information that is cause for concern include customer data that should be subject to defensible controls governed by regulations and laws such as Gramm-Leach-Bliley Act (GLB), HIPAA, Personal Information Protection Act (Law No. 57 of 2003 – Japan), European Union Directive 95/46/EC, but for one reason or another is not. Today, many organizations that are aware of possible offending data risk conditions choose to assume the risk as opposed to take steps to quantify their risks so they take a balanced and apprised approach to remediate the condition.

In many organizations, certain information is merely characterized as "inaccessible data," the hope and sentiment being that if information is designated as such, it is never going to pose a problem. The reality is that is an absolute corporate time bomb just ticking away. "Legacy data" often falls into this category. The truth is that if there is enough upside to the leveraging the "inaccessible" data, the data now takes on a different aspect and the corporate entity suddenly becomes willing to the data for a profit motive. BlueArc discovery experts know that the argument that data is inaccessible is rapidly weakening with each passing day as technology

makes it easier than ever to get to data on backup tapes or legacy systems. Simply put, "not knowing" what's where and invoking the battle tested "onerous and burdensome" argument is not a sufficient basis to bar discovery of this data. For example, if an organization receives a subpoena and has twenty terabytes of data on legacy storage that is potentially responsive, "not knowing" what's actually responsive or privileged there does not obviate the need to quickly find out what's there. In many instances an organization will wait until the last possible moment to address the issue. This can be a very costly approach to electronic discovery and is considered the basis of every CFO's worst nightmare.

To help CFOs understand the corporate governance implications (or lack thereof) of technology expenditures, it is important to qualitatively and empirically state and quantify the risks that the technology that they are being asked to purchase will ameliorate. The Litigation Data Lifecycle model is a good starting point when pointing out the value proposition of BlueArc.

The BlueArc Data Lifecycle Value Proposition

To see storage management efficiency and cost-containment opportunities in clear and uncertain terms, it is important to understand the Litigation Data Lifecycle. While not inclusive of every aspect of litigation data management, the model below provides a basic framework from which cost, workflow needs, matter requirements, technology and human capital requirements can be determined at various stages of an LDL process. The framework also lends itself to understanding the requirements of other data classification and management initiatives such as compliance, internal investigations and records management.

Unified Litigation Data Lifecycle

Functional Overview

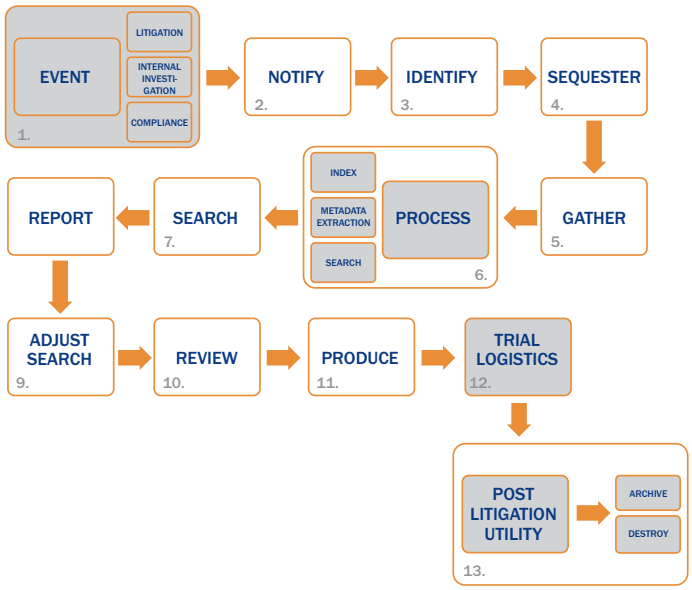

1. Event – in the context of litigation the "event" is the catalyst that initiates the discovery cycle. It is often referred to as a "triggering event." An "event" or "incident" in the context of data lifecycle management is the occurrence of one or more conditions that necessitate a response that may include but are not limited to the steps described in Figure 2. A common event could be a data breach where credit card numbers are lost or subpoena is issued for documents. As part of an incident response or event notification protocol, organizations should ensure that their corporate CFOs are one of the first ones notified when a triggering event occurs. Not only is this a best practice from a corporate governance perspective, but lowered thresholds of CFO liability should also be a consideration as to why CFOs should insist that this step be an integral part of the process.

A response to a triggering event is always easier when an organization has a high state of information awareness. Information awareness comes naturally to organizations that have the ability to quickly and continually classify their data in their storage environment in accordance with the organizations records management and other relevant taxonomies. As a precursor to data classification, an organization must have the ability to rapidly access and index multiple terabytes (and in some instances, petabytes) of data in storage, proactively, cost effectively and with minimal operational disruption.

Examples of the types of events that trigger discretionary and non-discretionary discovery responses are:

i. Litigation Discovery.
ii. Internal Investigations.
iii. Merger and Acquisition Due Diligence.
iv. Compliance productions – Department of Justic and Federal Trade Commission / related data collections.

While the main focus is litigation, there are many similarities in the data collection processes for initiatives i. through iv. above.

2. Notification – notification involves alerting the appropriate custodians who are owners of the data as well as "functional custodians," who are those individuals in the organization whose responsibility it is to provide, control and maintain the mechanisms that ensure the confidentiality, integrity and availability of data. These individuals and groups must be notified of the need to identify and preserve information in their custody and control as well as cease activity that may result in the destruction or alteration of information.

Increases in data management efficiency have created a paradigm shift for many organizations. The notification process is taking on a more procedural aspect in terms of its impact on the actual owners of data. Here's why; significant advances in storage I/O speed and throughput, long considered to be one of the most significant bottlenecks in getting to data, have enabled greater central control over data and the process described in figure 2. This is critical in terms of removing individual subjectivity and creating uniformity as to the information described in the preservation notice.

The impact of the speed variable on litigation data lifecycle management is significant. One of the foremost Federal Judicial experts on complex electronic discovery, New York Southern District Court Judge Shira A. Scheindlin, made it abundantly clear in her rulings that an organization has an obligation to periodically "refresh" its legal hold process (including but not limited to custodian hold notifications). Organizations that have the power to quickly control the data on their storage platforms are in a much better position to avoid sanctions that arise when they have the inability to meet legal hold requirements in a timely effective fashion.

3. Responsive Data Identification – early data identification is a "bounding process" wherein the parties seeking data, i.e. email, loose electronic files or database records, from an organizations storage environment. It requires developing "responsiveness criteria" that will bound the around the scope and form of the information sought.

Two common approaches to enterprise data identification are generally proactive and reactive. The proactive approach is considered by far to be the most efficient approach to enterprise discovery for litigation, records management, merger and acquisition due diligence or any other discovery intensive process. Unfortunately most organizations do not adopt a proactive discovery posture unless they are governed by non-discretionary compliance or regulatory requirements that force them into this mode. The prevailing institutional mindset among many organizations is "if it ain't broke, don't fix it." Given increased governmental and shareholder scrutiny (combined with lowered thresholds of liability for CFOs and other C-level

executives), adherents to the reactive discovery approach are finding their numbers dwindling. No longer is it acceptable practice to forego proactive discovery measures because of mistaken beliefs that it is costly or incredibly difficult to do. Stakeholders who prefer to wait for a compelling event before they engaging in early risk mitigation data discovery can have data collection costs 7-10 times higher than those organizations that have adopted proactive discovery best practices. In dollars and sense, this can translate to the difference between $300,000 and $3,000.000.

When an "Event" occurs, speed is of the essence. If systems on which discovery is being conducted are BlueArc enabled, the rate at which the identification of potentially responsive data is accomplished reduces operational impact and the risk of privileged document exposure.

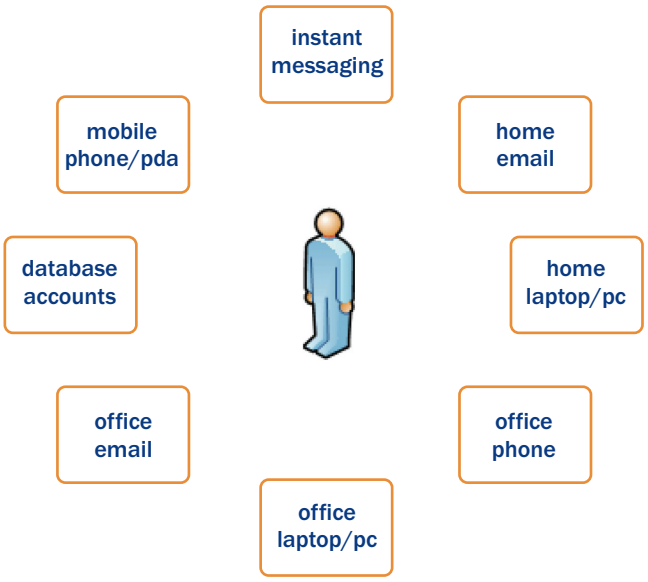

Figure 3

4. Gathering Responsive Information – effective early stage internal data collection presupposes the organization has a firm grasp as to the locations where information is to be found and the various forms the information may be in. As referred to earlier, data exists in unstructured, semi-structured and unstructured forms. Figure 3. above illustrates some of the human metadata considerations that may come into play when gathering data – these considerations are namely information dispersal and disparate data forms.

Locking data down in place is often the process that takes place behind the client firewall. This is when potentially responsive information is pooled to a staging area. Depending on the nature of the discovery initiative, the degree of forensic soundness may vary. It is important to preserve information about the data owner, source and original metadata to the extent possible. The figure below represents common human based data collection points from which data may be gathered for one or more of the initiatives mentioned in 1. above.

Given the increasing complexity of the data types that can be gathered for one or more of the aforementioned initiatives, the importance of the speed in the indexing and metadata extraction process cannot be underestimated. The eight user data points referenced in Figure 3, while not exhaustive, are representative of what we may find in many corporate environments today. Each user's data points consist of one or more unstructured, structured and semi-structured data types. Individual user data that is discoverable may exist in different geographic locations or on disparate systems, however to the extent the collection throughput is optimized, the cost and risk in the collection process can be significantly reduced.

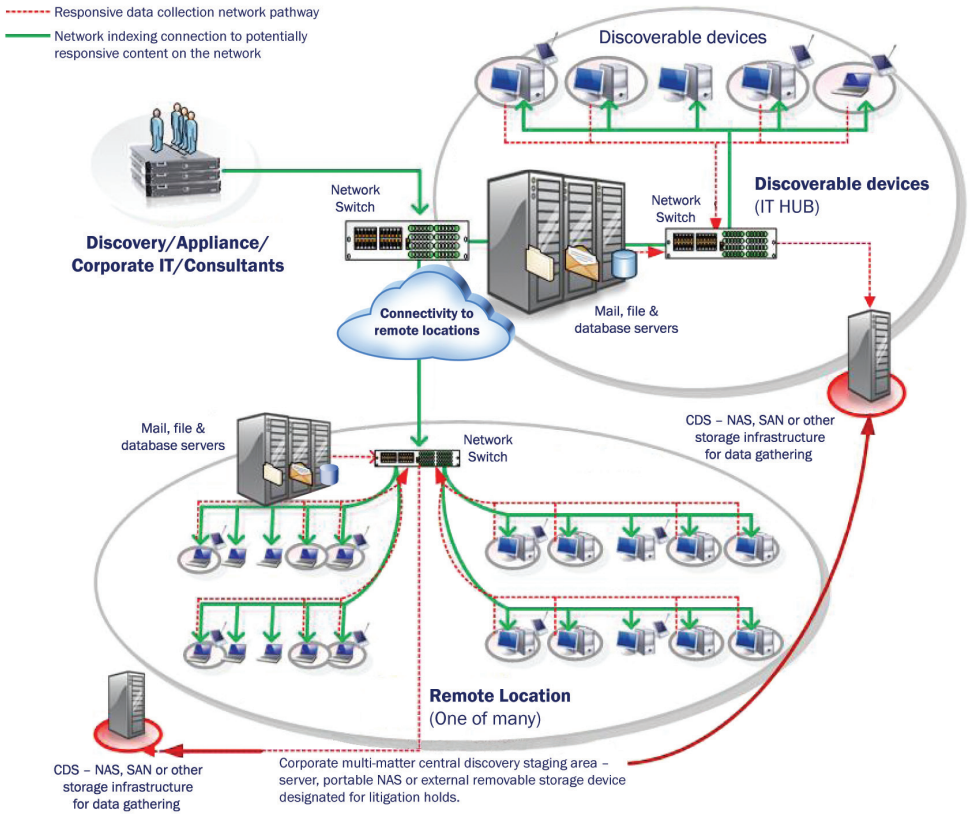

Figure 4.

Figure 4. describes the corporate environment. Gathering or aggregating data to a staging area is generally a precursor to shipping data off-site for vendor processing. An organization that has BlueArc enabled behind the firewall storage as a staging area can quickly move to stage 6 of the process in a much more cost effective fashion than they otherwise might.

5. Data Sequestration – a fancy term, but in the litigation context this is the litigation hold phase. Once the data in the target repositories have been searched and the data gathered, there are many ways to sequester information that meets the sought after criteria. In the context of litigation or investigations, we often refer to this process of locking data down as a "litigation hold." The analogous step in the compliance use case is the data archival process. All of the processes that have involved "touching" data (crawling, indexing, classifying, copying or moving) are very, very I/O intensive. The litigation hold process involves copying data with full metadata integrity to devices that physically or logical lock data down in a defensible and auditable fashion. Many storage technologies that play in this space tout themselves as "compliant storage." For the most part, they have built in encryption schemes that manage retention and prevent data from being altered. Depending on whether an organization has a proactive or reactive approach to the discovery approach to managing the content on their storage as well as the power and speed to index and classify content, the opportunity for data archival (as well as records retention/ management) which is shown in Figure 2 as a data end point, may actually occur much earlier in the lifecycle – it may be almost concurrent with data creation.

6. Processing – The term "processing" can vary greatly depending on the nature of the discovery initiative as well as the objectives of the persons or entities requiring the discovery. It is critical to define and understand what "processing" means from the stakeholder perspective. CFOs should understand the impact of processing because in general terms, it the most time intensive and costly non-legal fee part of discovery. The steps delineated below outline a few of the key points

that one should be aware of. There are many factors that determine the amount of time and cost, those factors include but are not necessarily limited to whether processing involves:

a. Metadata extraction – all unstructured files have metadata. There are three common types: system, application and user created metadata.
b. Hash value creation – Two common hashing algorithms used to generate file fingerprints are "MD-5" and "SHA-1." Hash values are calculated based on file content – larger files can take significantly longer to hash. This information is used for chain of custody and in de-duplication.
c. Data de-duplication: the full file hash is used for unstructured file de-duplication and for email (semi-structured data) a combination of time / date and to, from, cc (other metadata fields are hashed together).
d. Objective culling (system file removal or known non-responsive file removal) – removal of system files.
e. Deep indexing of content often referred to as full text indexing – this step is without question the most I/O intensive aspect of the lifecycle.
f. Specialized treatment – i.e.;
   i. Post tape restoration processing of emails to create archive load file formats for Legato, CA Message Manager, Enterprise Vault, etc.
   ii. Normalization of disparate file formats to a common format (pdf or tif) to facilitate review of data. Many document review tools work optimally with PDFs or TIFFs.
g. Searching - searching of large amounts of data is accomplished by searching on:
   i. Fielded system, application or user generated metadata.
   ii. "Free form" textual content.

The effectiveness of any search methodology is dependent on the indexing required by the tools that are applied post indexing. For example Boolean and concept searching require the underlying data on which the search is performed to be in a particular format in the index that the search engine uses. As is the case with data indexing, searching is an inherently I/O intensive process (particularly with federated searching) as requests are made to servers and the responsive content gathered and presented.

A "behind the corporate firewall" BlueArc enabled platform married with leading electronic discovery applications can reduce the cost of "processing" data by 80%. Leading remote backup and archive solutions can also facilitate the circumvention of many costly steps because of parallelization of steps in the LDL. The market price of "processing" discovery data without image conversion in step 5 is $500 - $750 / GB. Market price of processing with image conversion is $1500 / GB (endorsements as described in the Production phase, 11. is about $.02 - $.04 per page, keep in mind each GB generates about 65,000 pages). For corporate enterprises that have eDiscovery teams and the infrastructure necessary to process data, disbursements can be reduced by significant percentages.

7. Reporting – this is the process that provides "actionable intelligence" as to the information that flows through the process. It can and does in fact occur before and after many stages of the process, however in the context of the LDL, it is designed to provide detailed information about the responsiveness of the data germane to the initiative. This is a key phase that facilitates the "initial disclosure requirements" of Federal Rule of Civil Procedure (FRCP) 26(a) and the production scope negotiations that take place pursuant to FRCP 26(f) meet and confer. It allows a party to make the following statements to their counsel and adversaries: "I know the source of data that is responsive to this initiative, the data types and respective volumes as well as the age, ownership and extrapolated costs to production."

8. Adjust search – while functionally no different from the original search, the purpose of this searching is to provide iterative adjustment of criteria, fine tuning the results with the goal of obtaining the highest quality (least number of objects with the most responsiveness) results. This process allows what we in the industry refer to as "multiple bites at the apple." It is an

important phase in that once data is indexed and classified it can be re-used within the scope and purview of the primary data management initiative or entirely re-purposed for another.

9. Review – without question this is the single greatest cost component of the LDL. Attorney review fees are based on the corpus of information that needs to be coded:

a. For privilege status (attorney client privilege – communications between the client and their lawyers do not have to be turned over to the other side in discovery).
b. For responsiveness – responsiveness simply means "search hits that meet the criteria of documents sought in a litigation matter."
c. Subjective issues – attorneys will group documents in a way that helps them with their strategy. This is called "subjective coding."
d. To the extent that the amount information that needs to be reviewed can be reduced at data lifecycle points 3. through 5. In the model above, there will be a cascade effect in cost savings in that the amount of billable hours that attorneys need to spend reviewing data will be reduced. Another key factor that affects review time is the performance of the systems that the review platform is built on. Document review, like indexing and searching is a very I/O intensive process. To the extent that the systems that render and present the information is to be reviewed are faster, this will contribute to a reduction in the amount of time that it takes to view and classify data.

10. Data Production – while document productions on small scales are still done in paper, the overwhelming production of data both in the context of litigation and regulatory compliance are being done electronically. The I/O improvement achieved by utilizing BlueArc Titan Storage System can reduce the production time involved in getting electronic files normalized, usually by converting them from native electronic files to static images (PDF / TIFF) and endorsing the resulting images with Bates numbers and various protective designations, i.e., confidential, attorneys eyes only, etc. by half. By viewing the lifecycle as in interrelated continuum, production efficiency and cost control can be significantly impacted by leveraging each preceding step's efficiencies, namely those that take place in the processing and review phases.

11. Trial Logistics – this phase generally involves a trial presentation team working with a tiny fraction of the original data. Most litigation matters will not go to trial. It is estimated that of all the actions filed, only 5% actually make it to trial.

12. Post Litigation Data Utility – Without question, the ligation coding process that takes place in the review phase presents many organizations with the best opportunity they will ever have to classify large volumes of business related data. Corporate litigation document collections often end up in large databases that are external to the corporate data owner's infrastructure. While this scenario provides the vendors that host these data sets with a handsome annuity revenue stream, the utility of the data and related coding is often lost to the original owner. In most instances these documents are "coded" or classified in the context of litigation (or the initiative in question, i.e., records management taxonomy). These document collections represent one of the least recognized sources of institutional value for many corporations. The ROI in these document collections becomes obvious when subsequent litigation matters target them or when records managers realize that attorneys, experts, advisors and auditors have generated millions of dollars of value by classifying the documents irrespective of whether they were produced in the context of litigation and regulatory compliance. Practically speaking, once data leaves a corporate firewall, the organization often no longer has the ability control the security or proliferation of the information. It often becomes prohibitively difficult to find ways to leverage the information for profit center activities and business intelligence…imagine that.

Summary

In this paper we have reviewed some of the policy and environmental factors that impact legal data management efficiencies from a number of perspectives. Policy driven shifts and expansions in organizational roles, namely that of the CFO, are forcing all parties involved in the information lifecycle management model to re-think how and where value is generated in the organization. There are discretionary and non-discretionary variables that impact corporate governance and one of the key discretionary variables under the control of the CFO is IT infrastructure. CFO's have the ability and in many instances, the mandate to ensure that IT cap-ex and op-ex has broader corporate governance relevance than ever before. If one has the ability to spend funds in a way that accelerates information acquisition, use and disposition processes agnostically throughout the enterprise, then without question, value is generated throughout the corporate environment. While not intended to be an exhaustive treatment of factors that can impact litigation discovery or risk management, the topics covered herein are intended to provide corporate stakeholders, namely CFO's, with a different and highly relevant perspective on the potential impact of today's infrastructure investments.

Some of the issues that we have touched on that are addressed by BlueArc solutions relate to:

1. Reactive litigation readiness and related processes.
2. Proactive litigation readiness and related processes.
3. Privacy data remediation.
4. Rapid turn-around on enterprise data risk analysis initiatives, programs and internal investigation.
5. Unprecedented speed of implementation of records management programs
6. M & A due diligence.

As a fundamentally important component of any data management initiative, the I/O speed variable that Blue Arc storage management technology brings to bear has the ability to provide a significant acceleration factor for any data management initiatives above. For example, privacy data remediation involving the identification of non-public information can be done twice as quickly with by organizations that implement a BlueArc infrastructure.

The other initiatives listed can be managed and completed from project initiation to in significantly less time and more cost effectively with Blue Arc than not. BlueArc exemplifies the philosophy that data management lifecycles are a series of interrelated processes and steps that can leverage the same storage management platform should be comforting news for CFO's with an eye on litigation cost containment, greater adherence to compliance requirements and broad based ROI on IT expenditures. When considering a BlueArc solution, CFO's can rely on the speed and power of BlueArc storage technology to help accelerate the information access and data classification process, both of which are absolutely essential for cost containment and risk management.

Take the BlueArc discovery challenge, we'll review your "legal spend" and provide you with an organization specific BlueArc ROI model based on existing legal spend without and with a BlueArc solution.

About BlueArc

BlueArc is a leading provider of high performance unified network storage systems to enterprise markets, as well as data intensive markets, such as electronic discovery, entertainment, federal government, higher education, Internet services, oil and gas and life sciences. Our products support both network attached storage, or NAS, and storage area network, or SAN, services on a converged network storage platform.

We enable companies to expand the ways they explore, discover, research, create, process and innovate in data-intensive environments. Our products replace complex and performance-limited products with high performance, scalable and easy to use systems capable of handling the most data intensive applications and environments. Further, we believe that our energy efficient design and our products' ability to consolidate legacy storage infrastructures, dramatically increases storage utilization rates and reduces our customers' total cost of ownership.

BLUEARC®

BlueArc Corporation
Corporate Headquarters
50 Rio Robles Drive
San Jose, CA 95134
t 408 576 6600
f 408 576 6601
www.bluearc.com

BlueArc UK Ltd.
European Headquarters
Queensgate House
Cookham Road
Bracknell RG12 1RB, United Kingdom
t +44 (0) 1344 408 200
f +44 (0) 1344 408 202