



ESG Lab Validation Report™



Kazeon Information Server IS1200 eDiscovery & Storage Optimization

A validation study
by
ESG Lab
December 2007

Authors:
Tony Palmer, Brian Garrett

Table of Contents

Introduction	3
ESG Lab Validation	6
<i>Ease of Implementation</i>	8
<i>Ease of Management</i>	10
<i>Performance and Scalability</i>	14
<i>Cost of Ownership</i>	17
ESG Lab Validation Highlights	20
Issues to Consider.....	20
ESG Lab's View	20
Appendix	21

ESG Validation Reports

The goal of ESG Lab Validation reports is to educate customers about various storage and storage-related products, including storage systems, backup-to-disk solutions, storage management applications, backup and recovery software, storage virtualization platforms, etc. ESG Lab reports are not meant to replace the necessary evaluation process that end-users should conduct before making purchasing decisions, but rather to provide insight into these technologies. Our objective is to go over some of the highlighted features/functions of such products, show how they can be used to solve real customer problems, and identify any areas needing improvement. ESG Lab's expert third-party perspective is based on our own hands-on testing as well as on interviews with customers who use these products in production environments. This ESG Lab Validation report was commissioned by Kazeon Systems, Inc.

All trademark names are property of their respective companies. Information contained in this publication has been obtained by sources The Enterprise Strategy Group (ESG) considers to be reliable but is not warranted by ESG. This publication may contain opinions of ESG, which are subject to change from time to time. This publication is copyrighted by The Enterprise Strategy Group, Inc. Any reproduction or redistribution of this publication, in whole or in part, whether in hard-copy format, electronically, or otherwise to persons not authorized to receive it, without the express consent of the Enterprise Strategy Group, Inc., is in violation of U.S. Copyright law and will be subject to an action for civil damages and, if applicable, criminal prosecution. Should you have any questions, please contact ESG Client Relations at (508) 482.0188.

Introduction

Founded in 2003 and headquartered in Mountain View, California, Kazeon provides a wide spectrum of proactive and reactive eDiscovery solutions in response to litigation, information security & privacy, corporate investigations and regulatory compliance concerns. In addition, Kazeon provides a suite of Storage Optimization solutions, including data consolidation, migration & archival, file reporting and backup search.

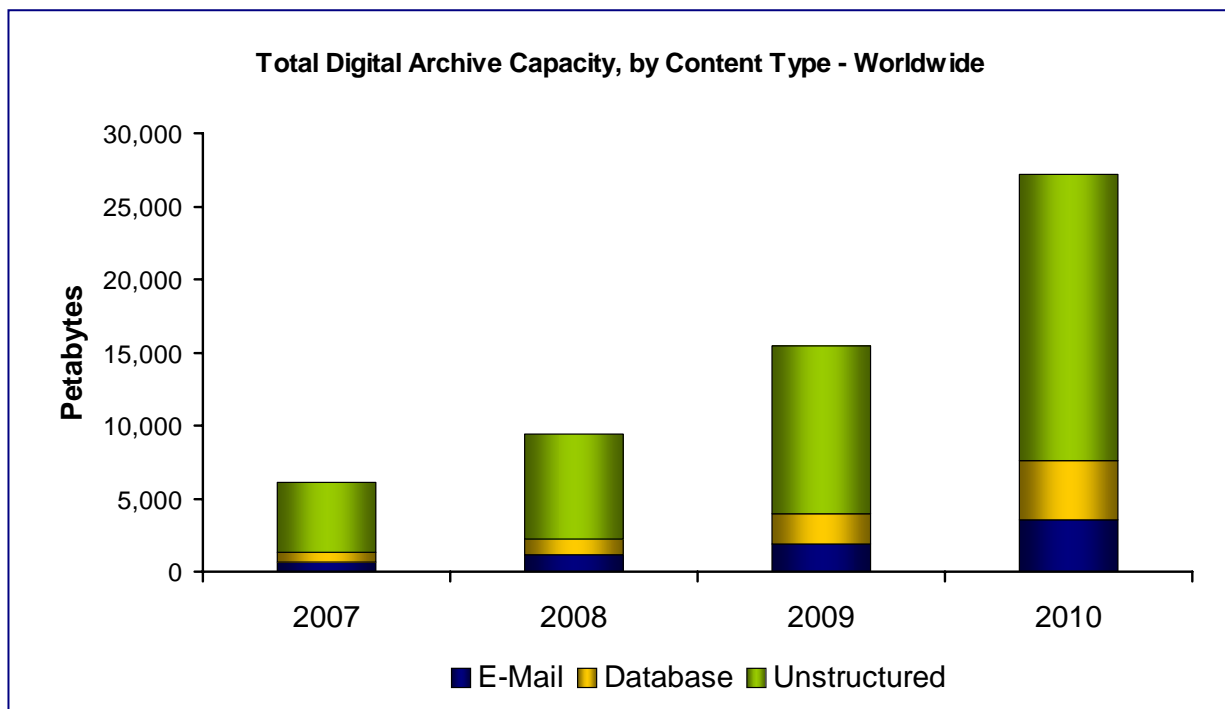
Launched in 2005, Kazeon’s flagship product, the Information Server IS1200, uses Information Access technology to integrate the eDiscovery functions of Identification, Collection, Preservation, Processing, Review and Analysis.

Companies are able to discover, index, classify, report, search and act on documents and e-mails distributed throughout the enterprise—within the data center(s), in branch offices, at desktops and on laptops. Any primary NFS or CIFS file system can be indexed, grouped and searched from a central location. Kazeon has integrated with—and built connectors to—dedicated archive devices and applications such as Network Appliance’s NearStore (with SnapLock), EMC’s Centera, Hitachi’s Content Archive Platform (HCAP), Plasmon’s UDO Archive Appliance and Symantec’s Enterprise Vault.

Background

ESG estimates that organizations will retain over 23,000 petabytes of unstructured data in digital archives over the next three years¹ in order to meet litigation and regulatory mandates (e.g., Federal Rules of Civil Procedure (FRCP), SOX, HIPAA, GLBA, PCI DSS) as well as to support business intelligence initiatives. To offer the reader perspective, 23,000 petabytes is enough capacity to store 2.3 million copies of the entire print collection of the library of Congress².

Figure One: Explosive Growth of Data



¹ ESG Research Report: *Digital Archiving: End-User Survey & Market Forecast 2006-2010*, January, 2006.

² Lyman, Peter and Hal R. Varian. "How Much Information", 2003, Retrieved from <http://www.sims.berkeley.edu/how-much-info-2003> on 7/1/2007

As the amount of information that organizations create and save grows, stress on IT departments also increases. The mandate to keep applications and their supporting systems available provides enough of a challenge to already over-burdened IT staffs that now face the additional tasks associated with harvesting newly generated and historical electronically stored information (ESI). Anyone who has needed to find anything on the Internet knows and understands the value of search. Users get fast, relevant search results culled from the 15 to 30 billion web pages on the net. This same capability is needed in the enterprise due to the amount of ESI retained and emerging eDiscovery requirements to, for instance, respond to litigation, address an audit request or handle a corporate investigation. All of the ESI in an organization—documents, presentations, spreadsheets, e-mail messages and attachments—should be searchable and reportable from the same interface, with the relevant results being actionable.

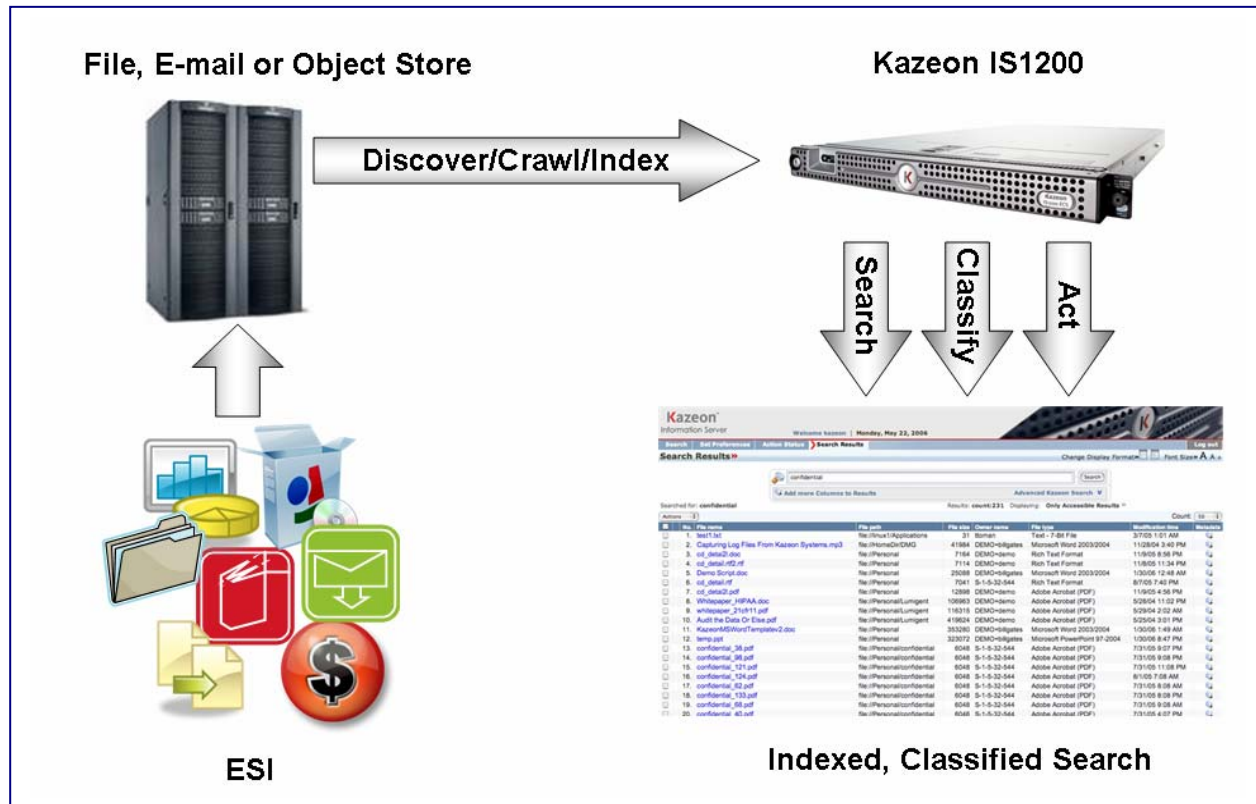
The corporate goals of improving eDiscovery processes and optimizing storage resources can be attainable if an organization can identify and produce specific subsets of data quickly and efficiently. Having realized that customers need to meet a variety of challenges and that an 'all-in-one' product approach may not meet specific requirements, Kazeon offers two enterprise solutions, each optimized for a specific business requirement: the Kazeon Information Server IS1200-ECS and the Kazeon Information Server IS1200-FRM.

The Kazeon Information Server IS1200-ECS offers a full suite of capabilities that are valuable for heads of litigation departments, paralegals, information security managers and compliance officers. During electronic discovery processes, customers need to reduce the amount of data being sorted in an effort to control processing and review costs. The IS1200-ECS can target specific pools of ESI or identify all common ESI locations so that information relevant to a compliance request or case can be quickly identified, secured to a single location and optionally preserved on compliant storage and set with the appropriate retention period. Customers are then able to apply Kazeon's processing, review and analysis capabilities to reduce the amount of data being sorted in an effort to control costs (typically from an outside service provider). The IS1200-ECS includes sophisticated reporting capabilities, enabling customers to identify duplicate files (even if the files have different names and extensions), thereby reducing the amount of data to be reviewed. In addition, customers can immediately identify irrelevant ESI by type, origin and date range. Kazeon's advanced search interface allows legal professionals to search for, review and analyze content for subject matter, important people or any specific vocabulary or jargon that might give further insight. Customers are able to easily label content with tags such as "relevant," "nonrelevant" or "privileged" to narrow the scope of the matter. At the appropriate time, all relevant and responsive information can be exported to 3rd party legal production tools via .zip or .pst containers. The entire process is auditable to ensure the process is defensible.

Ideal for storage administrators, The Kazeon Information Server IS1200-FRM enables customers to classify content based on file attributes (metadata) such as last access date, which facilitates the moving of any aged information to a lower cost storage system. Activities such as tracking trends for better planning, reviewing utilization by users and groups or attributing storage consumption to different departments for chargeback are also supported.

Most organizations store ESI on File Servers, NAS Appliances and/or CAS devices as illustrated in Figure Two. Kazeon's solution scans those systems and creates an index of metadata, which includes descriptive information about the file itself such as file name, extension, application, location, creation date, last modification date and all file system attributes. Using this metadata, Kazeon can provide data classification, reporting, federated search and even take actions based on policies to archive data.

Figure Two: Managing Electronically Stored Information (ESI) with Kazeon



ESG Lab set out to answer the following end-user questions in this report:

- How long does Kazeon take to set up and install?
- How much time will it take to index all our files?
- How much data can I search and index with a Kazeon Information Server IS1200 Appliance?
- How well does search perform?
- How much will it cost to manage all my electronically stored information with Kazeon?

In this report, ESG Lab presents the results of hands-on testing designed to validate Kazeon's ability to provide organizations with a cost-effective search, index and classification solution that scales to meet the ESI indexing needs of the largest enterprises, while remaining extremely easy to deploy and use.

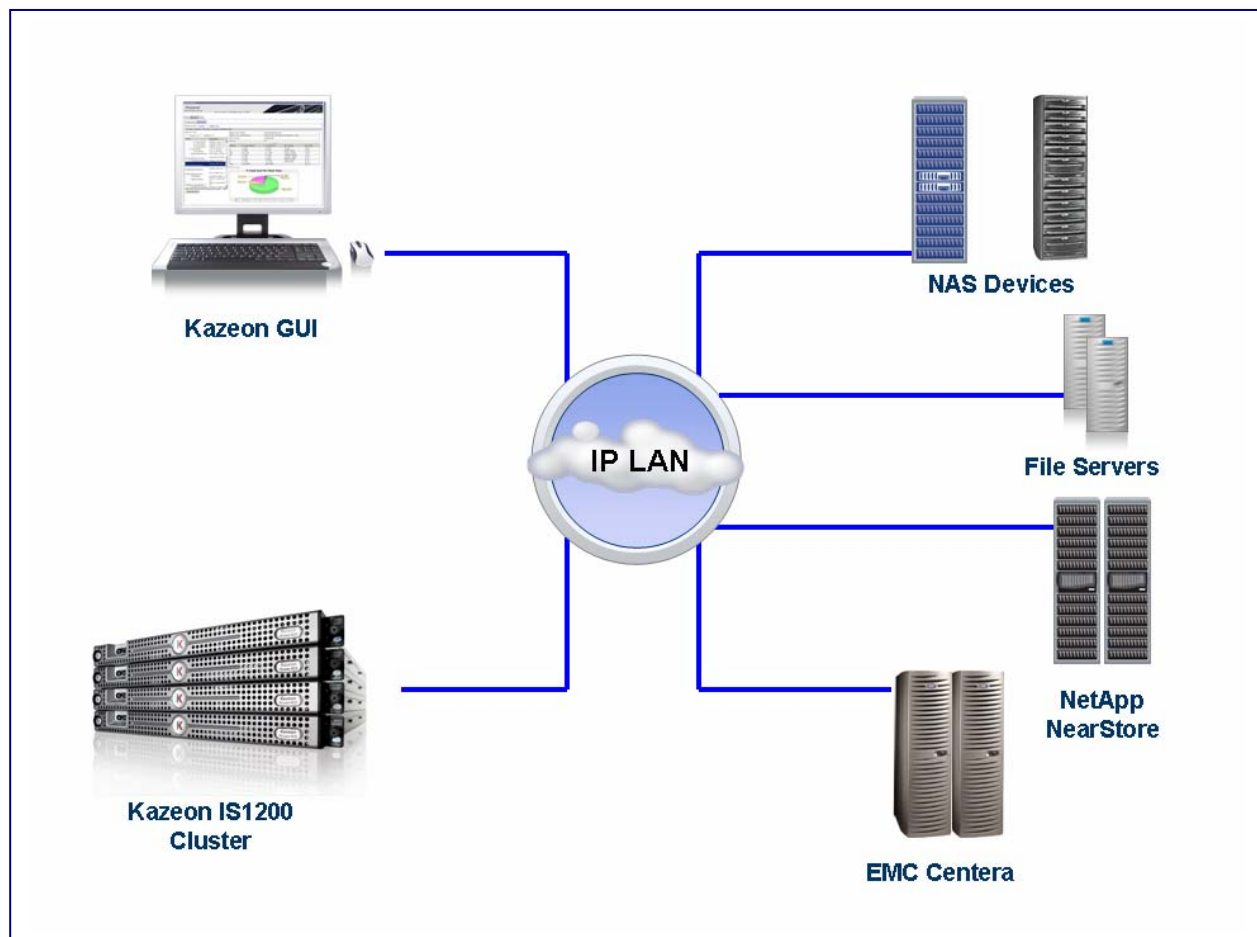
ESG Lab Validation

ESG Lab performed hands-on evaluation and testing of the Kazeon Information Server IS1200-ECS at Kazeon's Headquarters in Mountain View, California. Testing was designed to validate scalability of the Information Server platform and test its ability to scan, index and classify large volumes of electronically stored information as well as large numbers of individual objects. The IS1200's ease of use and total cost of ownership were also examined by ESG Lab to help customers balance the benefits with their investment.

Getting Started

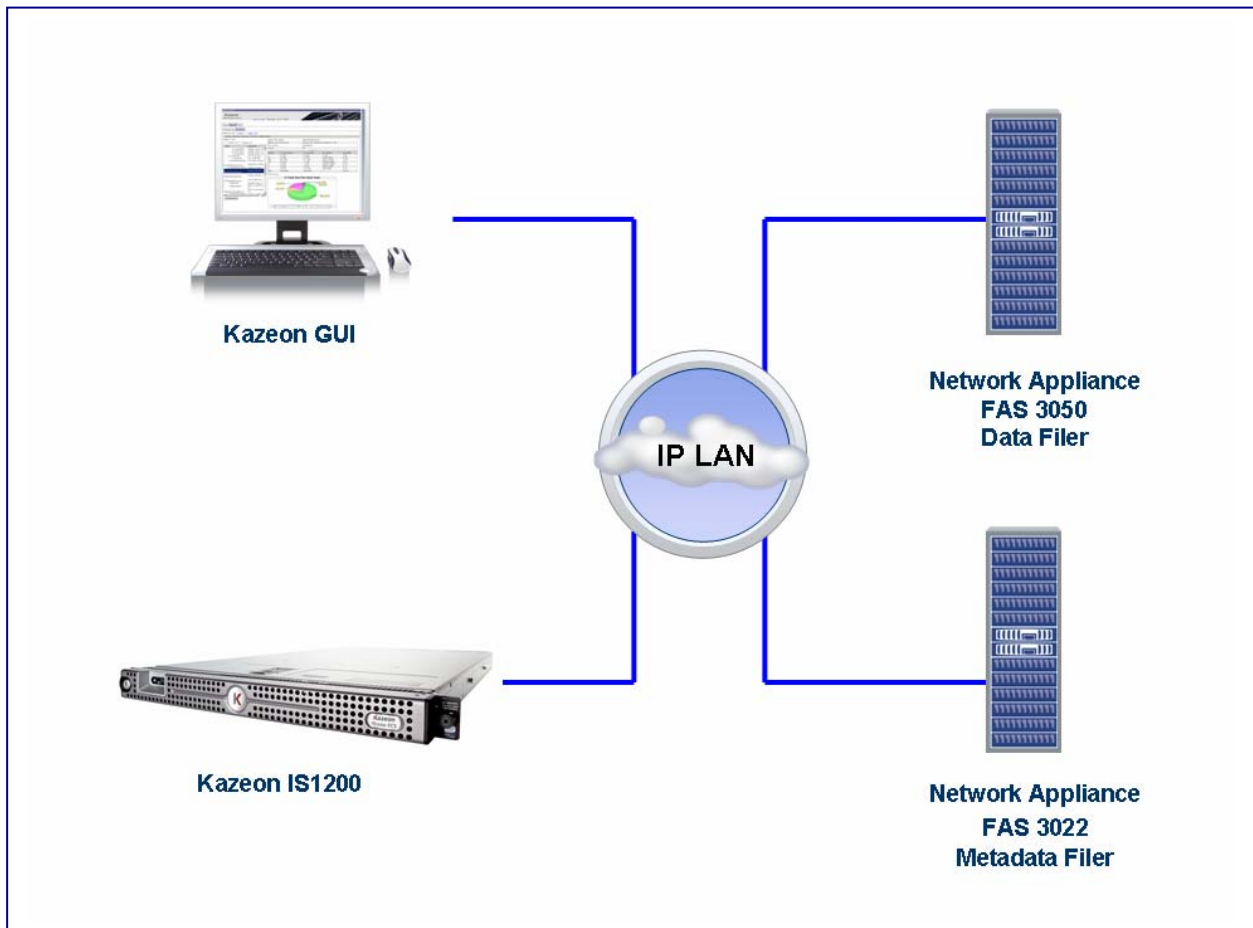
The Kazeon solution consists of Kazeon Information Server IS1200 appliances running Kazeon Information Server and Information Access Platform software. The IS1200 communicates with desktops, laptops, NAS systems, general purpose file servers, e-mail servers and archive storage applications and devices over an organization's corporate LAN. The Kazeon Java administration GUI and Command Line Interface (CLI) provide administrators with a central location for configuration and management. The IS1200 offers two web-based interfaces where authorized users can generate reports and conduct searches.

Figure Three: Kazeon at a Glance



ESG Lab's Test Bed for this report consisted of a single Kazeon IS1200 appliance and two Network Appliance filers³. One Network Appliance system contained the data sets to be tested against, while the second system stored the metadata and indices created by the Kazeon IS1200. The Kazeon solution does not store anything locally, which helps increase performance and allows customers to leverage the data protection inherent in the storage system. The data used for testing was gathered from three sources: An externally supplied data set consisting of a mix of productivity documents (Microsoft Office, Adobe PDF and other files), an internal data set consisting of the same types of files and a programmatically generated set of random text files (each file contains a randomly selected set of words and phrases).

Figure Four: The ESG Lab Test Bed



³ Detailed descriptions of hardware and software can be found in the appendix at the end of this report.

Ease of Implementation

ESG Lab's goals for this segment of testing were to validate the ease of configuration, monitoring and management from an administrator's perspective.

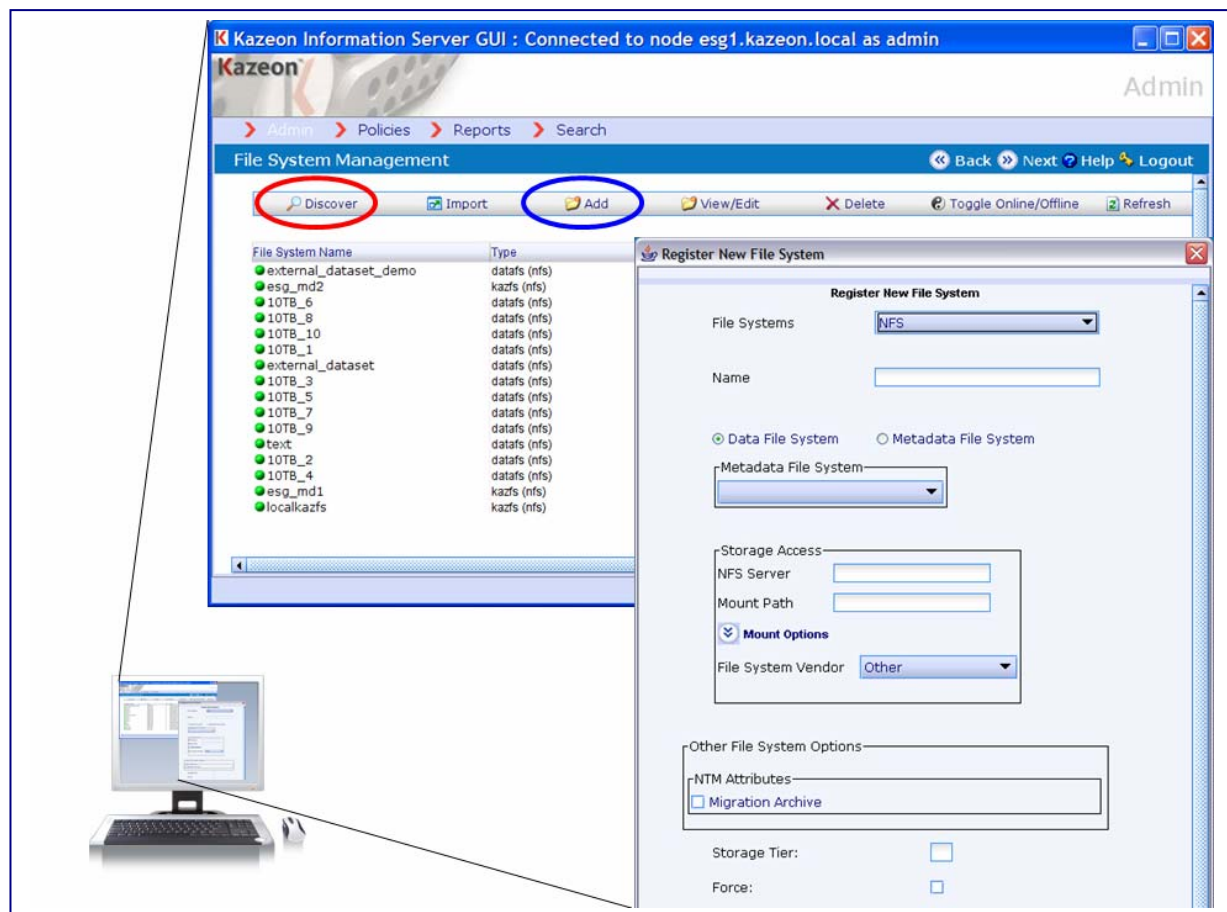
ESG Lab evaluated the ease of initial install by walking through a typical installation of a Kazeon IS1200 appliance. Installation was accomplished in the following sequence of steps:

- The appliance was plugged into the Ethernet network and an IP address was assigned.
- The metadata server was designated using the GUI.
- File systems were designated for scanning (crawl) via the GUI.
- Crawl progress was monitored via the GUI.

All of these steps, from power on to file system scan, were completed in twelve minutes.

The metadata file system lives on a NAS device or file server owned by the customer with the purpose of storing the metadata collected and generated during a crawl and index. To register a metadata file system, ESG Lab simply entered the name of the NAS server and the mount path of the share as seen in Figure Five.

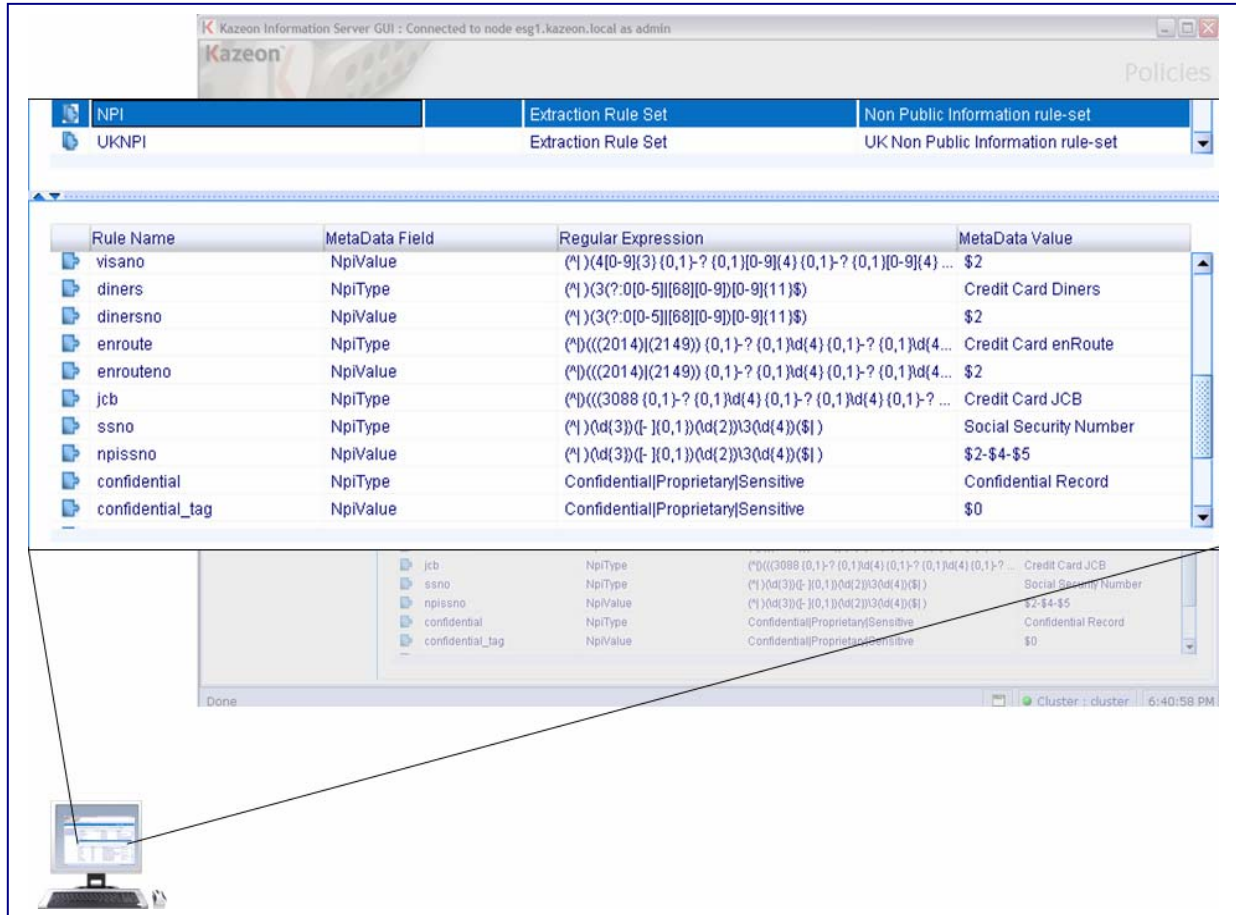
Figure Five: Registering File Systems



The next step was registration of the file systems to crawl and index. ESG Lab completed this by entering the names of the servers and mount paths of the file systems to crawl. Customers can also simply enter an IP address range and let Kazeon discover the filers, file servers and shares on the network. This is the 'Discover' option circled in red in Figure Five. The discovery method would be useful to an organization with a large number of NAS systems, file shares and/or servers that are not all managed by IT.

ESG Lab next examined policy-based management functionality, which enables organizations to define access, retention and logging policies and associate these policies with classes of data. Policies can also be used to take action based on a file’s attributes or content.

Figure Six: Configuring Policies



A policy group is made up of rules that define which files are affected by a particular policy. As seen in Figure Six, a rule can refer to a metadata field, content defined by a regular expression or both. The NPI policy in the diagram defines non-public information, such as credit card or Social Security numbers. This policy could be used to find all instances of non-public information—such as Word documents with employee social security numbers on a departmental share—in public file systems and move the files to a secure, private file system.

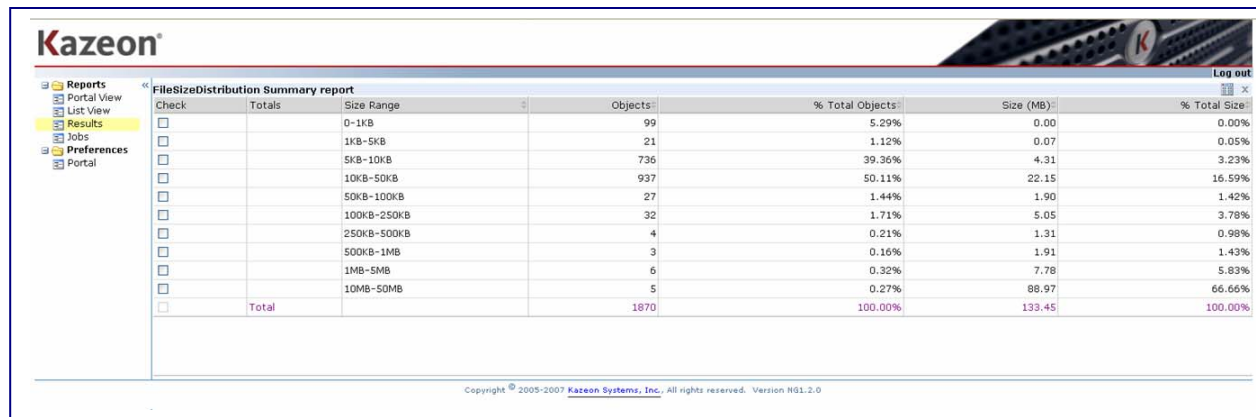
Why This Matters

With any emerging technology or solution, ease of implementation is of primary importance. A product may solve a business problem better than any other, but if it is difficult to install, organizations will tend to stick with existing, familiar methods, no matter how inefficient and costly. ESG Lab has tested and installed numerous storage systems and software packages and our goal is to perform an installation and have a solution fully functional in less than an hour. This was easily achieved with the Kazeon IS1200 appliance, with the entire process from power on to beginning a file system scan taking only 12 minutes. The ease of management and integration means that finding and acting on information stored throughout the enterprise becomes possible with minimal effort.

Ease of Management


ESG Lab's goals for this segment of testing were to evaluate the ease of search and discovery from a line of business perspective. ESG Lab first examined the reporting capabilities of the IS1200. Reports are a very important component of eDiscovery and Storage Optimization solutions, and Kazeon provides a wealth of predefined reports as well as the capability for the administrator to develop their own. ESG Lab found reports to be easy to use and understand, with a broad selection of available predefined reports based on both metadata and content. The file size distribution report seen in Figure Seven is based on collected metadata only and provides a summary of files by size.

Figure Seven: Summary File Size Distribution Report



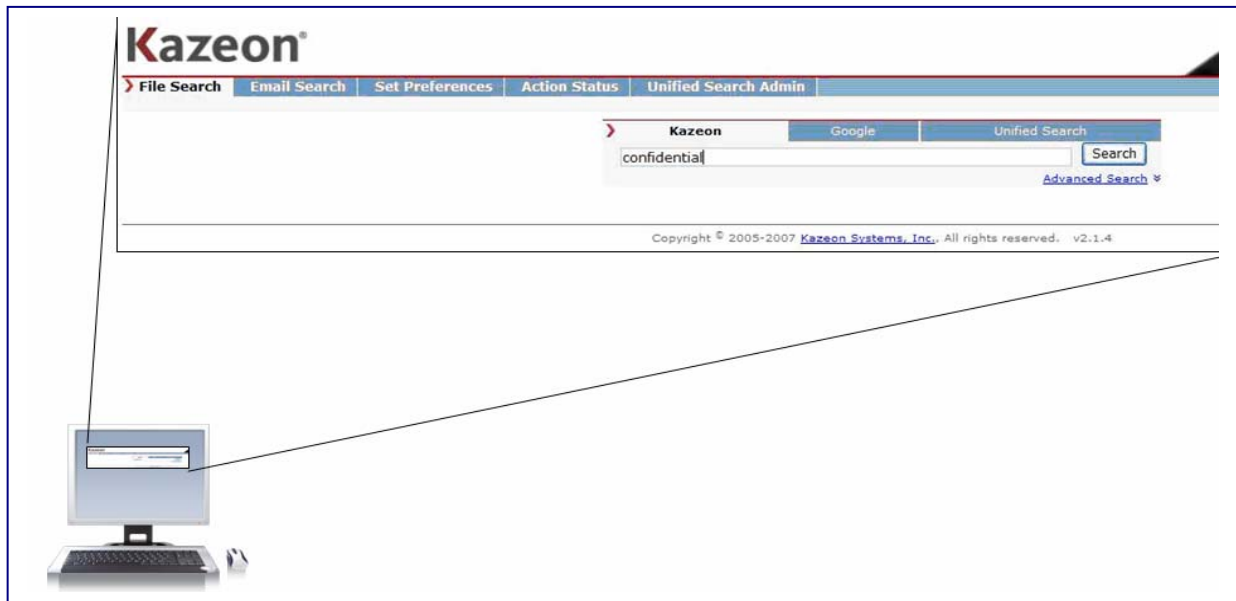
In contrast, the reports in Figure Eight draw on both attributes and content to provide detailed information on users, groups, compliance and other metadata or content-based criteria.

Figure Eight: Custom Reports



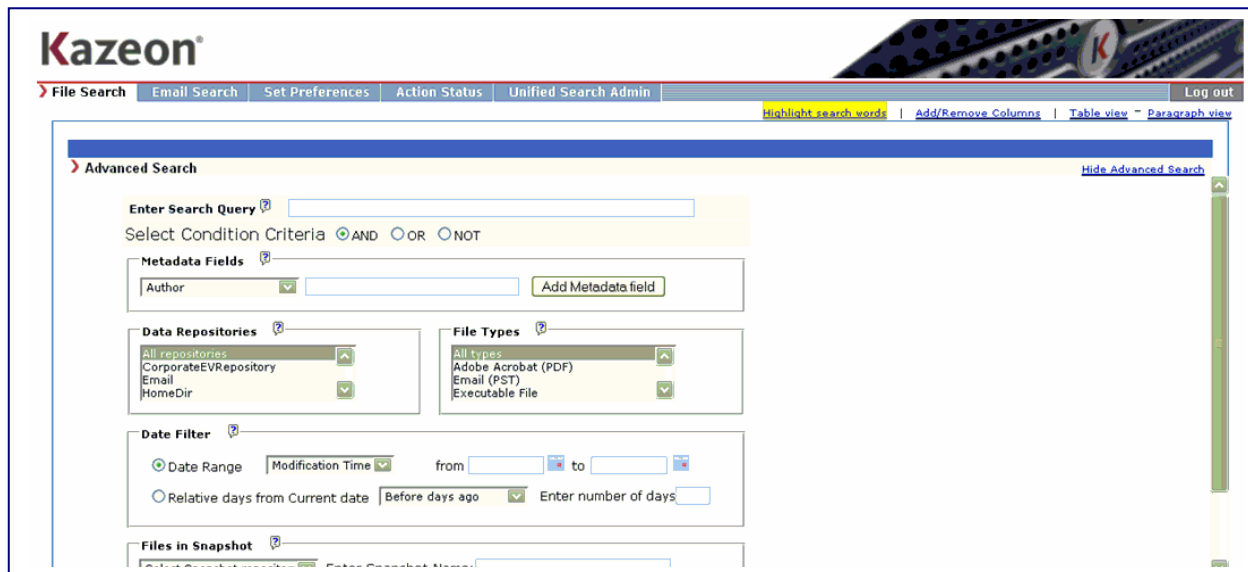
ESG Lab next explored a core functionality of the Kazeon solution: Actionable Search. It is notable that IT is able to configure hierarchical access controls so that a set of search results are displayed to a business user based on defined permissions. Access controls are especially important when litigation specialists, paralegals and information security managers are conducting eDiscovery queries. During the search process, counsel may reclassify certain files and messages as ‘privileged,’ preventing anyone from seeing them except for those working on the case. By securing search results, customers can ensure that employees only see information that is pertinent to them, which further ensures that sensitive information is only accessible to authorized individuals.

Figure Nine: The Basic Search



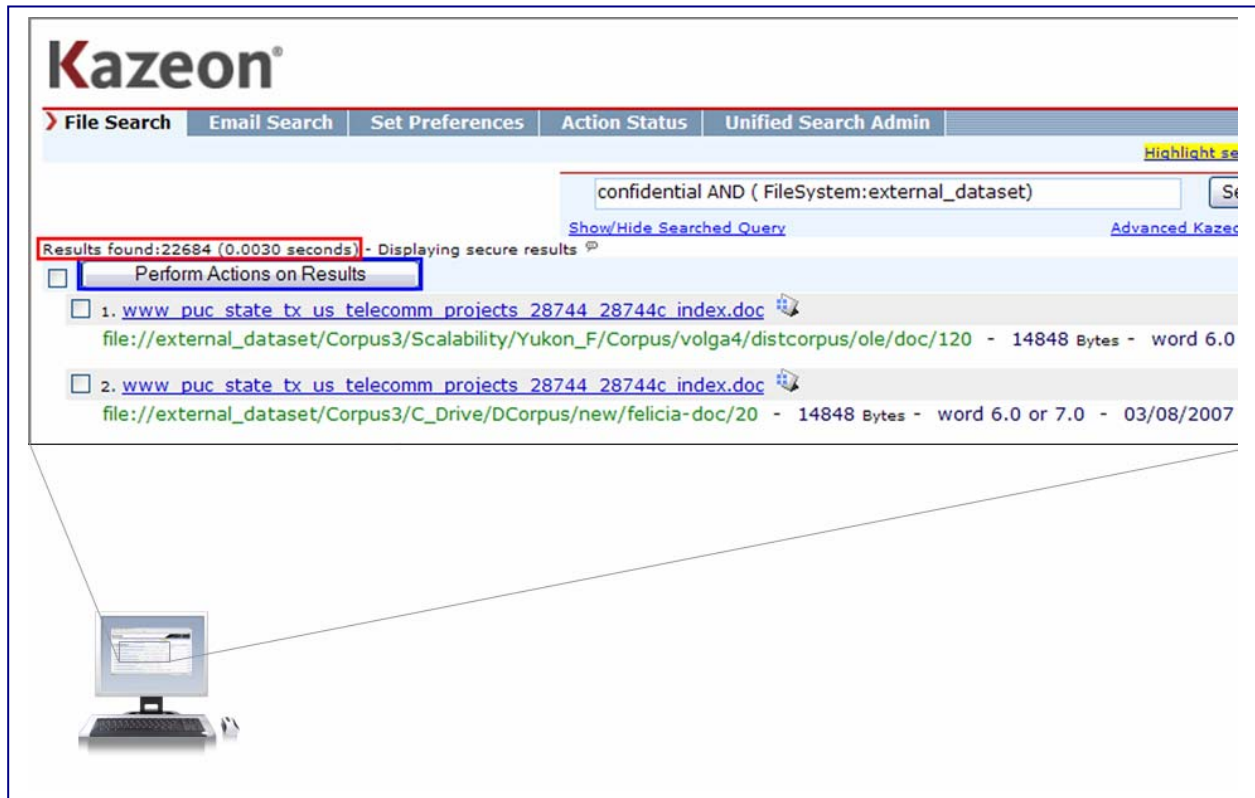
The basic Kazeon search interface—seen in Figure Nine—presents a familiar look and feel. There is a box to enter keywords and a link to access advanced search functions. The advanced search page—seen in Figure Ten—allows for search criteria such as Metadata Field, File System, File Type and a flexible set of Date Filters.

Figure Ten: Advanced Search Options



ESG Lab used the Advanced Search page and chose to search for the keyword ‘confidential’ in the ‘external_dataset’ file system. As seen in Figure Eleven, Kazeon returned 22684 files in .003 seconds. This is instantaneous from the point of view of an employee conducting a specific search. The search results are displayed as links, allowing a user to open any document by simply clicking on it.

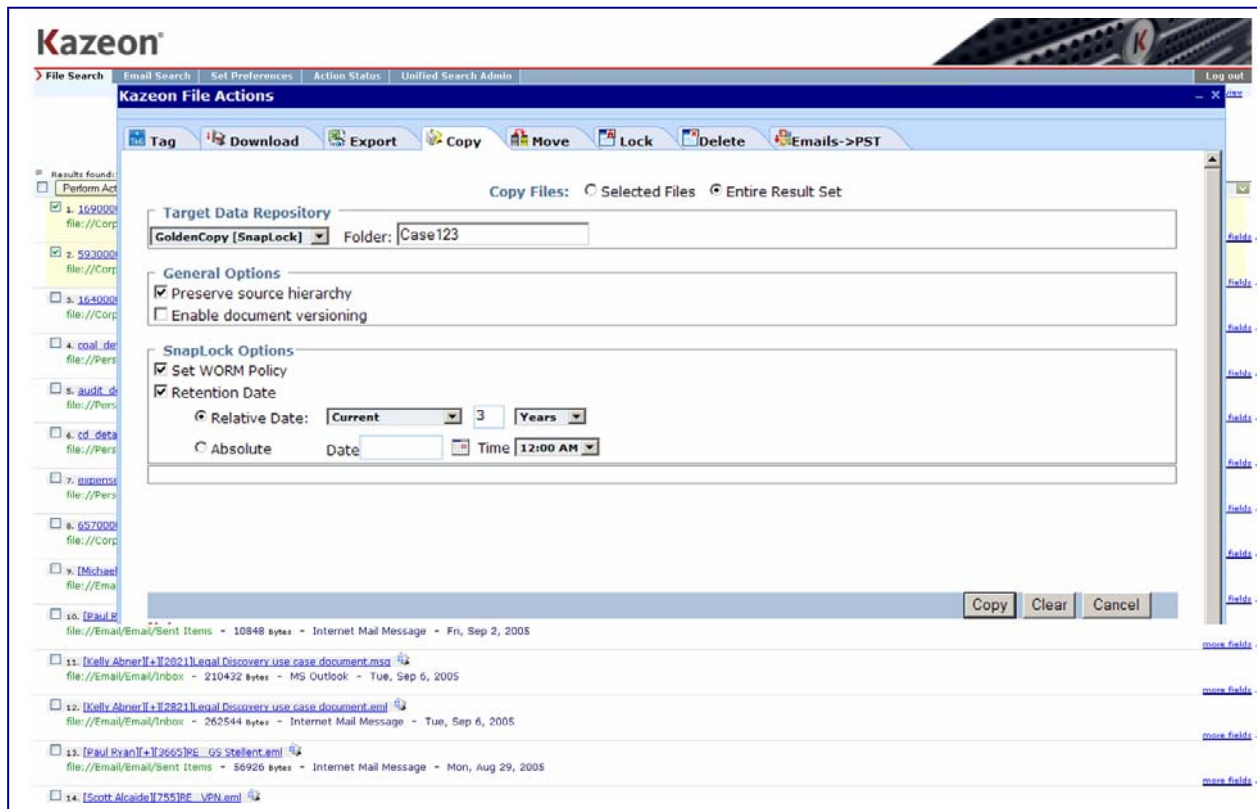
Figure Eleven: Search Results



Unlike traditional search engines, Kazeon offers an option to take action on any or all of the results. An employee conducting a search can elect to move files that are stored inappropriately (confidential files on an unsecured file server, for example) to a more suitable location. An attorney conducting an eDiscovery exercise might select specific or all files from a query and export them via .pst for litigation hold.

In Figure Twelve, the actions that can be taken on a file are shown in a series of tabs arrayed across the top of the window. The Copy option is highlighted in this instance, and the file(s) in question are being copied to a Network Appliance SnapLock volume. Kazeon is able to copy or move a file, designate it as WORM (Write Once, Read Many) and set the retention period—all in a single action.

Figure Twelve: Actionable Search



Why This Matters

The largest problem facing organizations looking to manage their ESI is the complexity of the environment in which this information lives. Traditional Enterprise Content Management software packages are complex systems that require high levels of expertise to use and manage. Administrators with this type and level of knowledge are hard to come by and therefore costly. ESG Lab found that the Kazeon solution is extremely straightforward to integrate into a complex environment and easy to use once in place. The Kazeon IS1200 presents a simple search interface with powerful content management capabilities. Ease of integration and use means that finding and acting on information stored throughout the enterprise becomes possible with minimal effort and expense.

Performance and Scalability

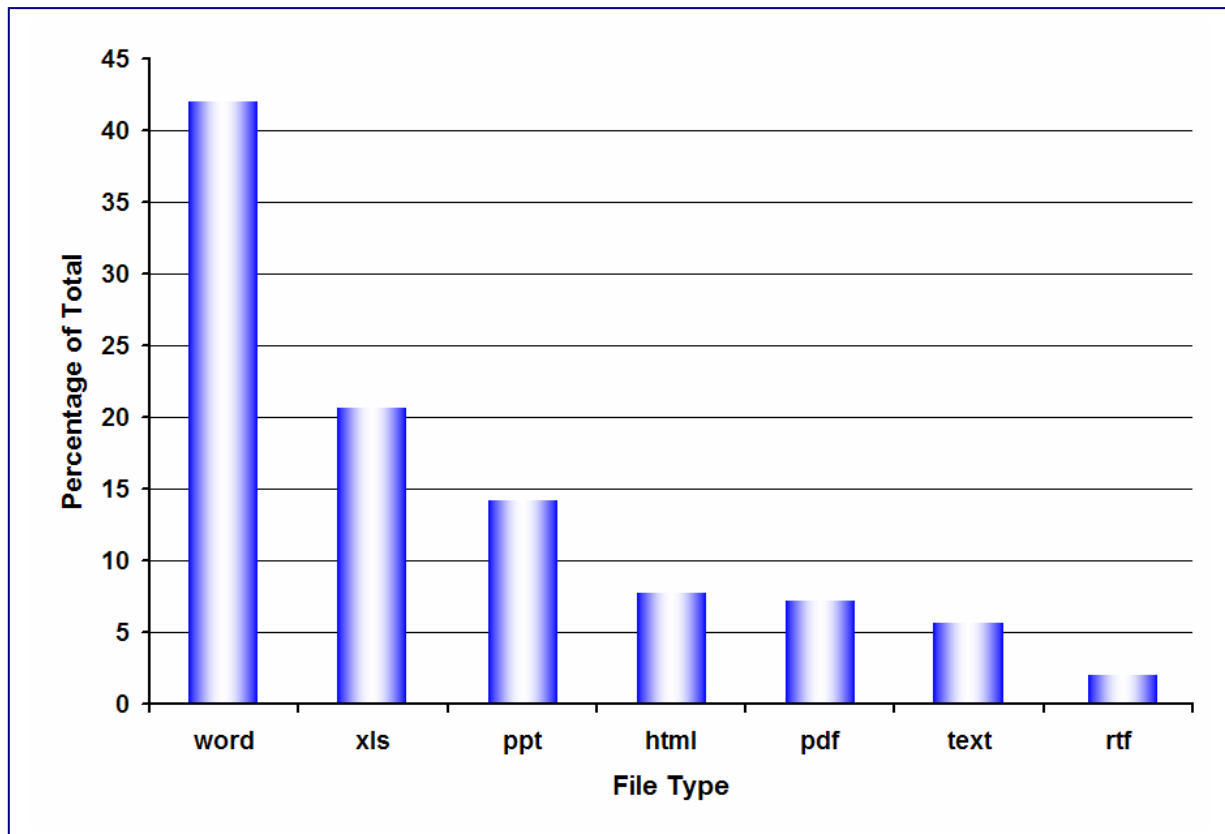
Electronically stored information, comprised of multiple types of data in heterogeneous file formats, typically lives on Networked Attached Storage or file servers—such as the ‘Data Filer’ in Figure Four. Data can also live on archive-specific devices and applications such as Network Appliance NearStore (with SnapLock), EMC Centera, Hitachi HCAP, Plasmon UDO Archive Appliance and Symantec Enterprise Vault.

To gather the necessary information, the first required step is a scan of the file systems—or ‘crawl’ in Kazeon terminology—and the creation of an index of the information. Following Kazeon-recommended best practices, the index is stored in a file system on a separate storage system, typically a Network Attached Storage device. The two types of crawl performed by Kazeon are ‘Deep,’ which scans metadata (information about the files being scanned) plus file content and ‘Basic,’ which scans metadata only.

ESG Lab Testing

ESG Lab first examined ‘Deep Crawl’ performance and scalability to determine the performance of an IS1200 as it scanned and indexed 10 terabytes of unstructured content. The distribution of file types used for this phase of testing is shown in Figure Thirteen. While the Kazeon appliance can open more than 360 types of files for indexing, ESG Lab tested it with seven of the most commonly found file types.

Figure Thirteen: Content Distribution Sorted by File Type



A Deep Crawl was performed on a data set containing approximately 10 TB of unstructured data files to validate the scalability of the Kazeon appliance. Since ‘Deep Crawl’ scans and indexes both metadata and file contents, ESG Lab chose to measure performance in MB/Sec—the amount of data processed per second.

Table One: Deep Crawl Performance Summary

Data Type	Size of Data Set	Elapsed Time (Hours)	Performance
Office Files	259 GB	3.5	22 MB/Sec
Text only	287 GB	3.7	21 MB/Sec
Office Files	~10 TB	63	46 MB Sec

ESG Lab examined system logs and validated that the IS1200 deep scanned and indexed a very large data set (10 TB) in just over 63 hours. In a clustered environment with 16 appliances, the Deep Crawl would be distributed across the appliances for faster performance.

ESG Lab tested the effectiveness of the Deep Crawl and index by searching keywords, phrases and file extensions. Table Two shows the results of four search queries entered through Kazeon’s search interface. The results appeared almost instantly for all queries. The file extension query returned with the longest latency—still under one second. This is extremely impressive given that the results represented fully 41% of the total number of files in the 10 TB dataset.

Table Two: Search Response Times

Query	Search Time in seconds
Googfroog	0.01 sec
“queries containing”	0.3 sec
Anurag	0.01 sec
type extension: pdf ,xls, ppt	< 1 sec

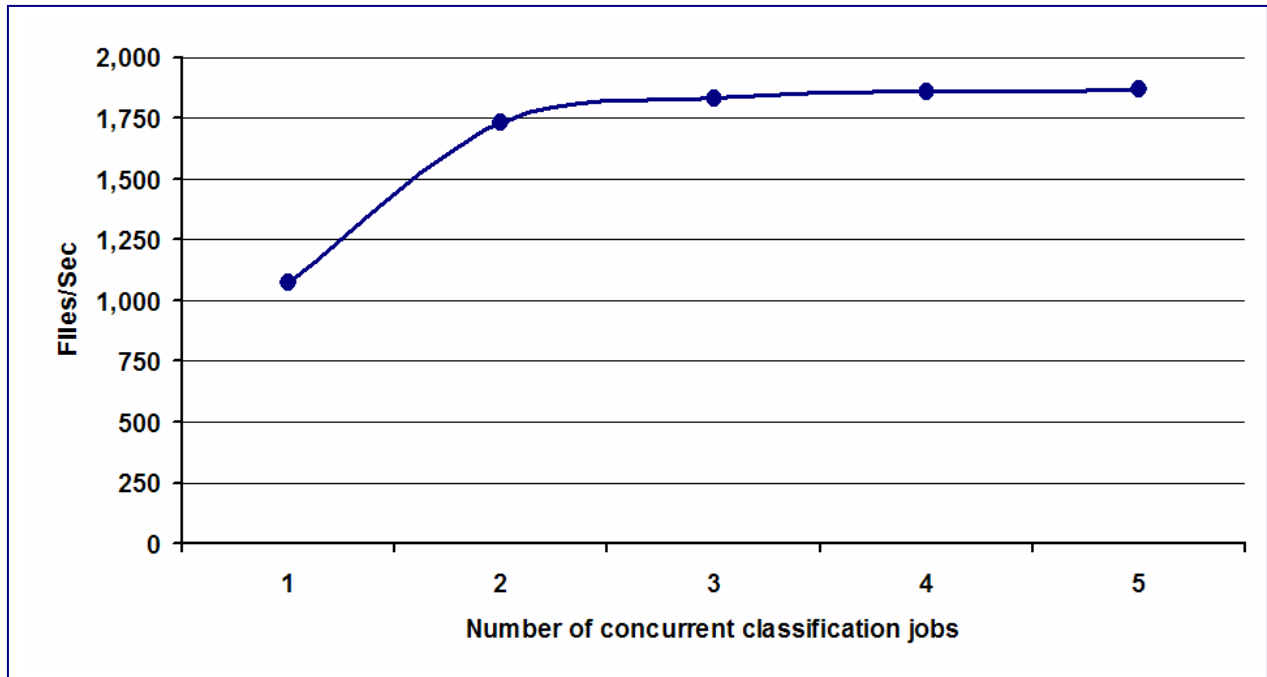
A Basic Crawl was performed on a data set containing approximately 100 million files to further validate the scalability of the Kazeon appliance. A Basic Crawl only indexes metadata, so ESG Lab chose to measure performance in Files/Sec—the number of files processed each second. The Basic Crawls were run by Kazeon prior to ESG Lab’s arrival on site. Examination of system logs yielded the data in Table Three.

Table Three: Basic Crawl Performance Summary

Number of Files	Search Index	Elapsed Time (Hours)	Files/Second
100 Million	No	18.5	1500
100 Million	Yes	39.7	700

ESG Lab also examined logs from multiple Basic Crawls to validate the maximum performance of a single IS1200 appliance. As illustrated in Figure Fourteen, Basic Classification scales sharply from one to two simultaneous jobs, but as the resources in the appliance begin to saturate, performance plateaus and remains steady. Resources were evenly distributed among the running jobs, which provided easily predictable performance. As with Deep Crawl, clustered environments would distribute Basic Crawls across multiple appliances for faster overall performance.

Figure Fourteen: Basic Classification Throughput vs. Number of Concurrent Jobs



What the Numbers Mean

- The Kazeon IS1200 sustained an impressive 47 MB/sec while performing a Deep Crawl on 10 TB of unstructured content.
- The Basic Crawl with search index disabled (for reporting only) ran in just under 20 hours, for an average scan rate of approximately 1,500 files per second.
- A Basic Crawl of the same data set with search indexing enabled ran in under 40 hours at a rate of about 700 files per second.
- Basic Crawl performance scaled well with multiple concurrent jobs per appliance.

These numbers prove that Kazeon can support large enterprise environments and provide consistent performance with very large data sets.

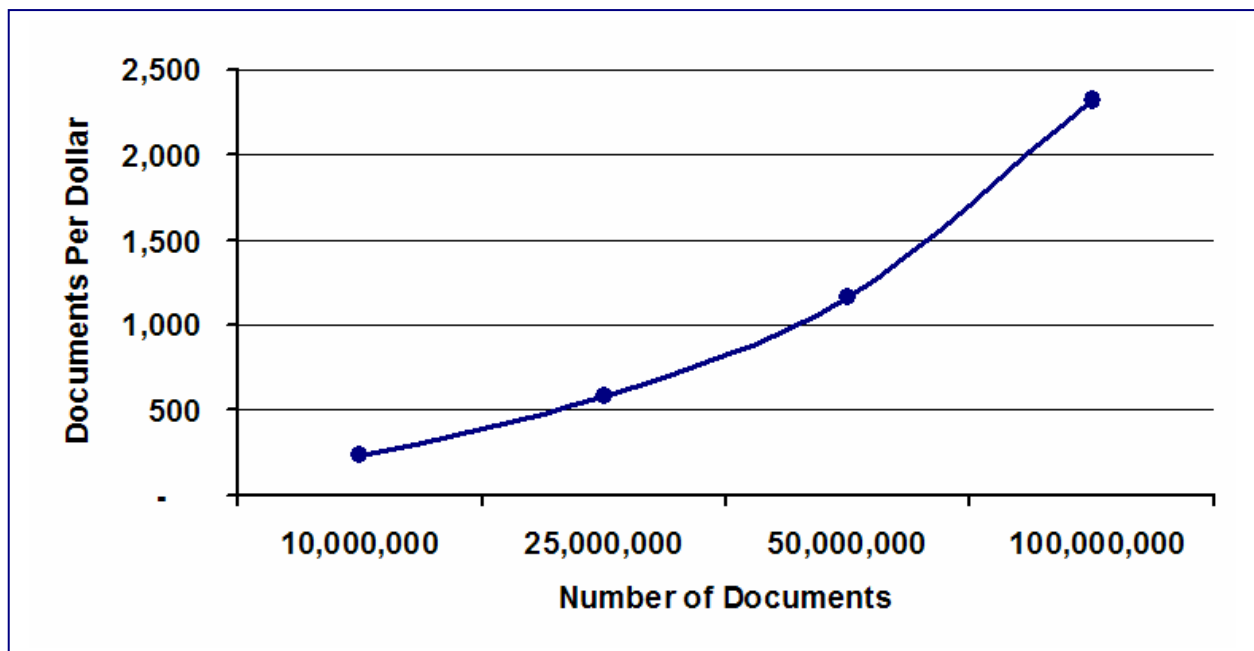
Why This Matters

ESG Research indicates that electronically stored content—such as office documents, e-mail, web pages, digital images, audio and video files—accounted for 78% of all archived capacity in 2007 and will continue to constitute the bulk of both primary and archived digital assets for the foreseeable future. Extracting useful information from that content manually is extremely labor- and time-intensive—and therefore expensive. Failure to produce data when required (as for electronic discovery) can be costly in the form of fines as well. ESG Lab validated that Kazeon’s IS1200 solution enables rapid scan and index of those mountains of data, along with very fast search and retrieval. With Kazeon business solutions, organizations can reduce liability and legal exposure while improving recovery time, lowering costs and improving operational efficiency.

Cost of Ownership

ESG Lab evaluated the cost of ownership of the Kazeon IS1200 enterprise search and index platform in the context of cost of acquisition, professional services and ongoing software maintenance. ESG Lab used two metrics to measure TCO in this context: ‘documents per dollar’ and ‘dollars per terabyte.’ Since Kazeon does not license their software or appliances by capacity or number of objects, recommended maximums for a single appliance and the list prices of Kazeon IS1200 appliances, professional services and maintenance over three years were used to calculate costs ⁴.

Figure Fifteen: Cost to Index 10 to 100 Million Documents Over Three Years

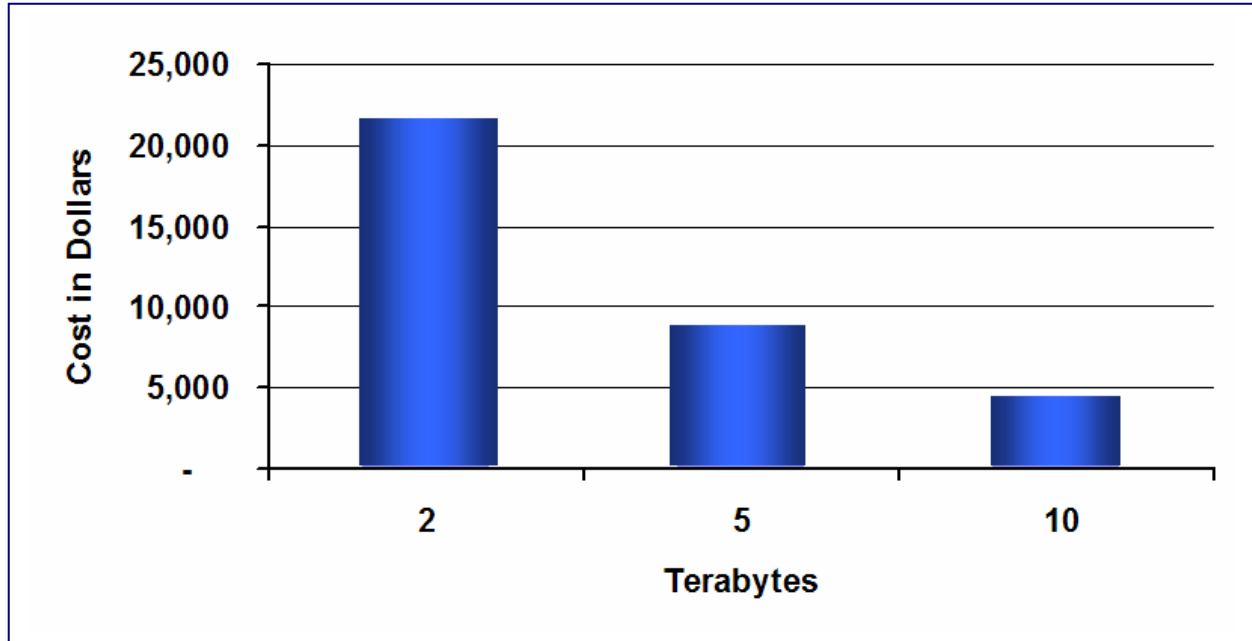


ESG Lab calculated that an enterprise can index and classify 100 Million files using just one appliance at a cost of \$0.0004 (.04 cents) per file per year. Stated differently, each dollar invested in a Kazeon solution can be used to index and classify up to 2,300 files per dollar, per year.

⁴ List pricing provided to ESG by Kazeon as of July, 2007. For more detail, please refer to Table Seven in the Appendix.

An analysis of an environment with Deep Crawl requirements for content indexing was equally impressive. Using the same list prices, the cost was just over \$4,300 per terabyte, per year for a 10 terabyte environment.

Figure Sixteen: Cost to Index and Search Up to 10 Terabytes of Content Over Three Years



ESG Lab next examined another component of the cost of search and index: Metadata storage requirements. Metadata for search and index should be kept on a reasonably robust storage system to maximize performance and availability of search. Kazeon’s metadata storage requirements vary with the type of crawl and index performed. An example of the metadata capacities generated with the data sets used by ESG Lab in this report can be seen in Table Four. ESG Lab found the capacity required to store metadata indices is negligible and does not have a significant impact on the overall cost of the solution. The column on the left lists the types of crawl and index performed, the two center columns indicate the quantity of content processed and the column on the right lists the total metadata generated for each data set.

Table Four: Metadata Efficiency

Crawl/Index	Storage Consumed	Number of Files	Metadata Generated
Basic, No Search	-	100 Million	116 GB
Basic, With Search	-	100 Million	237 GB
Deep	10 TB	-	73 GB

What the Numbers Mean

- At 2,322 documents per dollar, Kazeon's per file cost over three years is 0.04 **cents** per document, per year at list prices.
- The cost to manage unstructured content via Deep Crawl and index is \$4,300 per terabyte, per year.
- The Kazeon solution indexed 100 million files for reporting with a metadata footprint of about 100 bytes per file. With search enabled, the amount of metadata rose to only about 200 bytes per file.
- The Deep Crawl was equally impressive, requiring only 73 GB of metadata to index 10 terabytes of content.

Why This Matters

ESI is and will continue to be the single largest category of stored data well into the foreseeable future. ESG Research indicates that the worldwide capacity of digitally archived electronically stored content will exceed 27,000 petabytes by 2010. The 'dollars per document' pricing of existing solutions is unworkable with this quantity of information. A scalable solution that can bring these costs down to less than \$.01 per document will enable organizations to implement intelligent indexing, classification, search and actions across all their ESI. ESG Lab validated Kazeon's ability to provide a scalable, performant solution at a total cost of ownership of 1/25 of a cent per document and \$4,300 per terabyte. When examined in the context of an organization with multiple terabytes and hundreds of millions of files, this is a significant advantage.

ESG Lab Validation Highlights

- ☑ ESG Lab found the Kazeon solution very easy to implement and use, completing setup and installation in just twelve minutes.
- ☑ Scan and index were extremely fast and efficient, sustaining rates of 1,500 files per second for metadata only and 700 Files per second with search indexing enabled.
- ☑ Deep Crawl performance was also impressive at 47 MB/sec when deep scanning for content.
- ☑ A single Kazeon IS1200 appliance easily scanned, indexed and searched 100 million documents and 10 TB of ESI.
- ☑ The search functionality was familiar in operation and performed crisply, returning all results with sub-second response times.
- ☑ Kazeon is highly cost efficient, managing 2,300 documents per dollar per year. That's .04 cents per document.
- ☑ ESG Lab was particularly impressed with the integrated policy engine, enabling not only search and discovery, but actions against found objects.

Issues to Consider

- ☑ The ability to search within live structured databases like Oracle would greatly enhance organizations' ability to search. The Kazeon Information Server Software Version 3.0 added e-mail support for Exchange, PSTs, Journals, Notes and SMTP-based e-mail.
- ☑ Customers may already have existing migration tools, lessening their need for actionable search.
- ☑ Many customers already have existing eDiscovery processes and the move to a new process—no matter how easy—is perceived as a challenge.

ESG Lab's View

ESG Lab was quite impressed with Kazeon's integration of content-aware indexing, classification, search, reporting and migration together in one package to address legal discovery, compliance, data privacy and security challenges. The IS1200-ECS can extract information residing within hundreds of standard and non-standard file formats including a number of industry standard file formats such as Microsoft Outlook personal archive files (.PST) and the Digital Imaging and Communications in Medicine (DICOM) standard. Integration with the Google Search Appliance enables customers, for the first time, to issue a single search query and see combined intranet (Google) and unstructured content (Kazeon) search results in a single screen to address legal discovery demands.

Through innovative data indexing, classification and search technology, Kazeon has helped hundreds of customers drive out excessive cost, risk and productivity loss associated with the rapid growth of electronically stored information to tackle the challenges of Legal Discovery, Compliance and Data Privacy for less cost than traditional Enterprise Content Management software. Organizations across the globe in every industry sector are under increasing pressure and scrutiny to maintain the integrity of their data. Kazeon ensures that sensitive, business critical and confidential information residing on desktops, laptops and corporate networks is identified and managed—resulting in reduced liability and legal exposure, while delivering an immediate impact on business credibility. The largest challenge facing Kazeon today is staying ahead of the pack. Large players like EMC are preparing enterprise search offerings aimed squarely at Kazeon's core. In order to continue to succeed in the marketplace, Kazeon must continue to innovate and build their lead while the rest of the market plays catch up.

It is ESG Lab's conclusion that the Kazeon IS1200 platform scales very well, supporting 100 Million objects (or 10 terabytes) per appliance with the ability to scan and index tens of millions of objects per day while offering a clustering scheme that allows users to grow their environment as needed. ESG Lab verified impressive price/performance with a single appliance able to index 2,500 documents per dollar and a cost as low as \$4,300 dollars per terabyte. This is anywhere from half to an order of magnitude less than any competitive pricing ESG Lab was able to obtain as of this writing. ESG Lab found the system to be extremely easy to implement and operate, with an installation that completed in twelve minutes and full indexing and classification completed in less than three days.

Appendix

Hardware	Software
Kazeon IS 1200 Appliance	3.0
Network Appliance FAS 3022, FAS 3050	OnTap 7.0.6

Table Five: Data Set One File Type Distribution

File Type	Percent of Total	Percent by Size
word	41.54	16.72
xls	20.24	3.20
ppt	14.12	73.00
html	7.59	0.26
pdf	7.00	5.47
text	5.55	0.04
rtf	2.02	0.63

Table Six: List Pricing and Maintenance Costs

Component	List Price	Cost over Three Years
Kazeon IS1200	\$80,000	\$80,000
Professional Services	\$6,000	\$6,000
Maintenance and Support	\$14,400	\$43,200
	Total Cost	\$129,200
	Per Year Cost	\$43,067