


# i-drive ENHANCED STORAGE SOLUTION SYSTEM OVERVIEW

---



version 1.0

Title: i-drive ESS System Overview, release V1.0, [DATE].

Part#: 1001 1.0

Copyright ' [DATE], Anuvio/i-drive.com. All rights reserved.

i-drive, i-drive.com, and Sideload are registered trademarks of Anuvio i-drive.com.

Actuate and ReportCast are registered trademarks of Actuate Corporation. e.Report, e.Reporting, Live Report Extensions, ReportBlast, SmartSearch, Transporter, Virtual Report Distribution and XML Reports are trademarks of Actuate Corporation. All other products and corporate names referenced are trademarks of their respective companies. Cisco, and Cisco Systems are registered trademarks of Cisco Systems, Inc. or its affiliates in the U.S. or certain other countries. EMC<sup>+</sup> and EMC are registered trademarks and EMC Enterprise Storage. Celerra is a trademark of EMC Corporation. Exodus is a trademark of Exodus Communications, Inc. Foundry is a registered trademark of Foundry Networks, Inc. Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries. Linux is a registered trademark of Linus Torvalds. Nortel Networks and Alteon are trademarks of Nortel Networks. OpenLDAP is a trademark of the OpenLDAP Foundation. Oracle, Oracle 8i, SQL\*Plus, and SQL\*Loader are registered trademarks of the Oracle Corporation. SPARC is a registered trademark of SPARC International, Inc. Products bearing SPARC trademarks are based on an architecture developed by Sun Microsystems, Inc. Sun, Sun Microsystems, Ultra, Solaris, and Java are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All other company or product names are used for identification purposes only, and may be trademarks of their respective owners.

REVIEW DRAFT

Welcome to the *i-drive Enhanced Storage Solution (ESS) Overview*. This guide provides a technical overview of the i-drive system.

## i-drive ESS Documentation

*i-drive ESS System Overview*

*i-drive ESS Installation and Configuration Guide*

*i-drive ESS Application Deployment Guide* *i-*  
*drive ESS Operations and Maintenance Guide* *i-*  
*drive ESS Application Administration Tools Guide*

REVIEW DRAFT

---

Contents

REVIEW DRAFT

---

## Contents

Preface .....	3
i-drive ESS Documentation .....	3
Contents .....	4
Overview .....	7
Application Zone .....	7
Storage Zone .....	8
Network Zone .....	8
Application Zone .....	10
Web Tier .....	11
Overview .....	11
Load Balancing .....	11
Hardware and Software .....	11
Application Tier .....	12
Overview .....	12
Application and Integration Interfaces .....	12
File System .....	12
Subscriber Management .....	12
Directory .....	12
WebDAV .....	12
Functionality .....	13
Hardware and Software .....	13
Database .....	14
Overview .....	14
Logical Data .....	14
The Aardvark Name Server (ANS) .....	14
Universal Look-up Tables .....	14
Account Data .....	15
Metaphysical Data .....	15
Physical and Logical Data Mapping .....	15
Hardware and Software .....	16
Scheduler .....	17
Overview .....	17
The Scheduler Portal .....	17
Functionality .....	17

Hardware and Software .....	17
Storage Zone.....	17
Transfer Server .....	19
Overview .....	19
Features.....	19
Bandwidth Allocation .....	19
Custom HTTP/HTTPS Servers.....	19
Transfer Authentication .....	19
Transcoding .....	19
File Repository Access .....	19
Metadata Extractor .....	21
Overview .....	21
Physical File Data .....	21
Hardware and Software .....	21
GarbageMan .....	22
Overview .....	22
Functionality.....	22
Hardware and Software .....	22
Network-Attached Storage (NAS).....	23
Network Zone.....	24
Network.....	25
System Requirements and Tools.....	27
System Components and Hardware.....	28
Application Zones Diagram.....	29
Storage Zones: Logical Network .....	29
Storage Zone Diagram.....	30
Application Zones: Hardware and Network Overview.....	30
Application Zones: Hardware and Network Overview Diagram.....	32
System Administration Tools.....	33

## Overview

The i-drive Enhanced Storage Solution(ESS) is a robust, highly-scalable and extensible storage management platform (the Enhanced Storage Platform, ESP) and suite of file management and sharing applications. The i-drive ESS scales to millions of users, petabytes of logical storage, and hundreds of thousands concurrent file transfers.

This section provides an overview of the components that comprise the ESS, including the file storage system, and the network architecture. The i-drive ESS is organized into three zones:

Application Zone

Storage Zone

Network Zone

### Application Zone

Software in the application zone handles all application transactions except for physical file transfers. It is home to the core storage management platform and end-user applications.

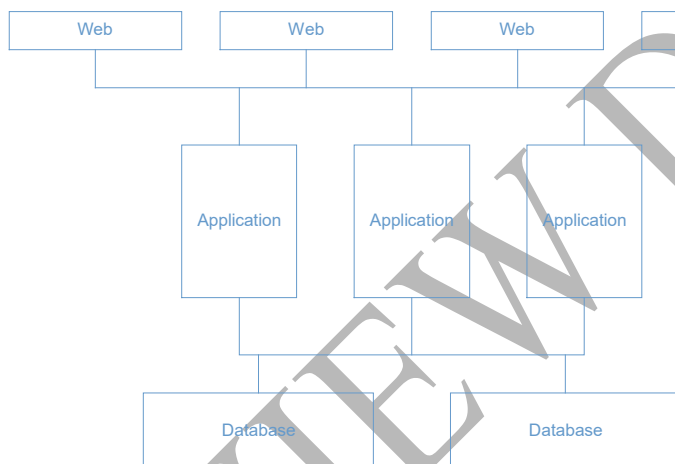


Diagram 1. Application Zone

The following table describes each main component of the application zone.

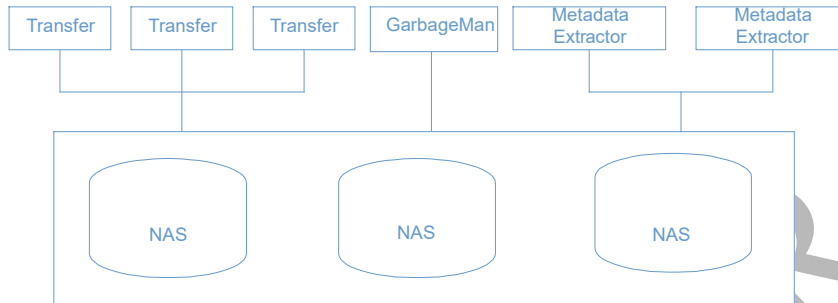
Component	Description
Web	Serves as the primary access point to the application. It manages all incoming requests and serves supporting static content. The Web server encrypts and decrypts all data using SSL. It also load balances across the application servers.
Application	Handles all requests forwarded by the Web servers. The core of the Enhanced Storage Solution. It is composed of two primary layers, the Enhanced Storage Platform (ESP) and the applications (PC and Mobile ESS) built on top of the ESP.

Database (Account, ANS, and metaphysical)	Stores all account, metadata, and file information managed by application instances.
Scheduler	Executes scheduled tasks, such as e-mail alerts or account updates.

## Storage Zone

The storage zone stores all the physical file data in the ESS. The storage zone handles file transfers, allocates bandwidth, transcodes and transforms content, and extracts metadata from files.

The following diagram shows each component of the storage zone.



*Diagram 2. Storage Zone*

The following table describes each main component of the storage zone.

Component	Description
Transfer Server	Serves all file uploads, downloads, and sideloads (file transfers from other servers). It authenticates transfers, allocates bandwidth, and transcodes files.
Metadata Extractor	Extracts structured information from files (such as keywords) when they enter the i-drive system, and stores the metadata in account databases.
GarbageMan	Manages the deletion of physical files and related information from the i-drive system.
Network-attached Storage (NAS)	Stores all physical file data. Typically accessed using NFS.

## Network Zone

An implementation of the ESS involves 3 networks. The following table describes these.

Component	Description
Content	Contains the load balancers, Web servers, application servers, and Transfer servers. For availability and scalability, there are typically multiple content networks, often referred to as application nodes .
Data	Contains the databases and handles all the traffic for database request processing.



Administrative	Contains the administrative and support systems needed to monitor and maintain the network.
----------------	---

REVIEW DRAFT

---

## Application Zone

This section overviews each main component of i-drive s application zone, including:

- Web Tier

- Application Tier

- Database Tier

- Scheduler

REVIEW DRAFT

---

## Web Tier

### Overview

The Web tier is the primary access point to an ESS installation. It handles all application requests, forwards requests for dynamic content to the application tier, and serves supporting static content. The Web tier performs SSL encryption and clustering/fail-over functions.

The Web tier forwards all requests for dynamic content to the Application tier. It does not serve dynamically generated content. The Web server requests dynamic content (via servlets) using the Apache JServ protocol.

The Web tier also does not serve any transfers (file downloads, uploads and sideloads). File transfer operations are managed by the Transfer Server.

**More Info** For more information, refer to the section titled Transfer Server.

### Load Balancing

Load balancing across content network nodes is implemented using DNS load balancing.

All Web servers in a single content network node share the same virtual IP address. Load-balancing at this level takes place at layer 4 of the network.

The Web tier then performs load balancing across the set of application tier instances in the same content network ( node ). Once a user has established a session, all traffic is routed to the same application tier instance.

### Hardware and Software

i-drive recommends the following web server configuration:

Intel PC hardware, running Red Hat 6.2 Linux OS

Apache 1.3.9 & Open SSL

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.

---

## Application Tier

### Overview

The i-drive application server is home to the core Enhanced Storage Platform services and is often the locus of end-user applications that locally invoke the platform and interact with users of a variety of devices supporting a variety of markup languages. i-drive s PC and Mobile ESS applications run within the same application server process. Other applications may invoke the platform s services locally or remotely using a standard RMI/CORBA Remote Procedure Call (RPC) mechanism.

The application is referred to by many names.

**Application instance** or instance is a single executing Java process executing the application.

An **application server** is a machine that is running at least one application instance.

An **app-tier** is a cluster of several application servers. (The app tier provides a superset of the functionality of a single application server).

A **VM** is the same as an application instance (short for Java Virtual Machine).

**Back end** is used because other front ends, or clients, can be built to use the ESP API.

i-drive s standard file management applications (the PC and Mobile ESS) reside in the application server and generate user interfaces using XML and XSL stylesheets. These stylesheets enable extensive customization and device optimization. Default support includes HTML, HDML/WML, and C-HTML. These applications invoke the ESP API using its local Java interfaces.

### Application and Integration Interfaces

The application server provides an object-oriented interface to all of i-drive s storage management capabilities as well as a number of integration points for interaction with existing software systems. Applications can run locally within the same process and directly invoke the Java API, or may run remotely using RMI/CORBA RPC and be implemented using a variety of languages and runtimes.

### File System

The application server provides an object-oriented interface to all of i-drive s storage management capabilities, roughly divided into Identity, Resource, and Information Management Services. Applications can run locally within the same process and directly invoke the Java API, or may run remotely using RMI/CORBA RPC.

### Subscriber Management

The Subscriber Management API enables integration with external provisioning, billing, and authentication systems.

### Directory

The ESP provides an abstraction for integration with external directories for access to personal address books as well as group management for addressing and access control. A standard LDAP adapter is available.

### WebDAV

The application server also supports the supports the WebDAV HTTP protocol (documented in RFC 2518). Version 3 of the ESP implemented three WebDAV s four primary features. The following table describes these:

*Note* The locking portion of the RFC 2518 is not implemented.

Feature	Description
Collections	Ability to create sets of documents and retrieve a hierarchical membership listing (like a directory listing in a file system).
Properties	Ability to create, remove, and query information about Web pages, such as their authors, creation dates, etc. Also, the ability to link pages of any media type to related pages.
Namespace operations	Ability to instruct the server to copy and move Web resources.

**More Info...** For more information on WebDAV, go to <http://www.webdav.org>.

## Functionality

The i-drive application server is a Java 2 servlet container and ORB. The servlet container is derived from Apache JServ. It uses the standard Sun-provided ORB for RMI connectivity. The application answers requests to invoke servlets via the AJPv11 protocol or answers remote procedure calls via RMI. Servlets decode the parameters in the request and form a response in some markup language (For example, HTML, WML, CHTML, XML, etc.). RMI requests proceed as normal remote procedure calls.

The following outlines the request to invoke a servlet:

The application server listens on a designated port for a servlet request (via the AJP protocol, Apache JServ Protocol).

Servlets delegate HTML/CHTML/WML page rendering to the i-drive Rendlets page composition library.

Rendlets invoke command objects to access and manipulate data.

Command objects encapsulate database transactions and use the ESP API to access and manipulate the account databases.

If a request accesses physical files, the application redirects file transfers to Transfer servers in the appropriate Storage Zone using encrypted tickets for authentication and identification. Otherwise, the page is returned to the user.

## Hardware and Software

The i-drive application server configuration is as follows:

Java 2 server application.

Three instances of the i-drive application run on a Sun Ultra 60. Each instance can handle 200 concurrent requests, thus each application server can handle 600 concurrent requests.

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.

## Database

### Overview

The i-drive ESP stores two types of information: logical data and physical data.

Logical data includes accounts, files, folders, and other information relevant to the user experience or operation of the system. The ESP stores logical data in a collection of Oracle databases.

Physical data is the contents of actual files. The ESP stores physical file data as discrete files using Network Attached Storage (NAS).

**More Info** Refer to the section titled Network-Attached Storage (NAS).

### Logical Data

Logical data includes the accounts, files, and folders as well as other information used by the system or relevant to the user experience. All logical data is stored in a collection of Oracle databases. The database system consists of four main logical components:

The Aardvark Name Server (ANS) (centralized data in one database).

Universal look-up tables (centralized data in one database).

Account data (distributed across multiple databases).

Metaphysical data (metadata about physical files).

Each logical component consists of a number of database tables, along with stored procedures for accessing or modifying the data.

### The Aardvark Name Server (ANS)

The Aardvark Name Server (ANS) contains centralized user account data that must be in one location (database) in order to preserve the uniqueness of the data.

The following table shows ANS account information:

Field Name	Description
Name	The unique user name.
Id	The number that uniquely identifies the account.
Domain id	A unique identifier for the domain the account is in. The domain id is mapped to the domain name by one of the look-up tables.
User id	The number that uniquely identifies the user. <i>Note...</i> In the future, there will be no distinction between user and account.
Database id	The identifier that determines the database that the account lives in.

### Universal Look-up Tables

Universal look-up tables contain data that is not part of an individual account. Tables reference the data in the account databases. Standardized values in short form are used to reduce the overall size of the database, which allows for more efficient queries. Data includes:

File mime types

Partnership information

Advertising information

**Note** This is centralized in order to preserve consistency of data across the site.

### Account Data

Account data is all the information specific to an individual i-drive account. Contents include:

Logical files

Folders

Friends

File metadata

Also, general account information (if not stored in an external directory), such as:

Password

E-mail address

Preferences

Each account is contained entirely in a single database, but as many account databases as necessary can be used.

### Metaphysical Data

Metaphysical data is metadata about physical files, such as:

ID3 tags for MP3 files

Author and title information from .doc files

Meta tag information from HTML files

Currently, metaphysical data is stored centrally. Future releases of the ESP will support distributed management of metaphysical data.

### Physical and Logical Data Mapping

When a file is copied from one i-drive to another, the physical data, or file, is not copied, rather, new logical data points to the single copy of physical data.

Data mapping requires that each logical file reference an identifier, which identifies the unique physical file. This identifier is called the `file_data_id`. Each `file_data_id` has a `location_id`. The `location_id` identifies a unique location (directory) in a physical filesystem. The combination of `location_id` and `file_data_id` is mapped to an absolute path by a PL/SQL package in the database. The application uses this package to identify the directory where new files can be added, or existing files can be downloaded. The database contains the following information:

Feature	Description
<code>nfs_server</code>	The 'home base' for a collection of file systems.
Filesystem	Each <code>nfs_server</code> exports a number of filesystems.
<code>host_server</code>	The application tiers and the Transfer servers.

mount_point	The mapping from a host_server/filesystem combination to an absolute path where the filesystem can be accessed from that server.
Filepath	The path within the filesystem where a file is located.

Each file lives in a filepath. To access a file from a particular host, the filesystem package takes the host\_server and mount\_point information and then adds the filesystem and filepath information to construct an absolute path.

The filesystem package also performs some accounting to track the space usage and number of files in a filesystem.

The following diagram shows the filesystem tables.

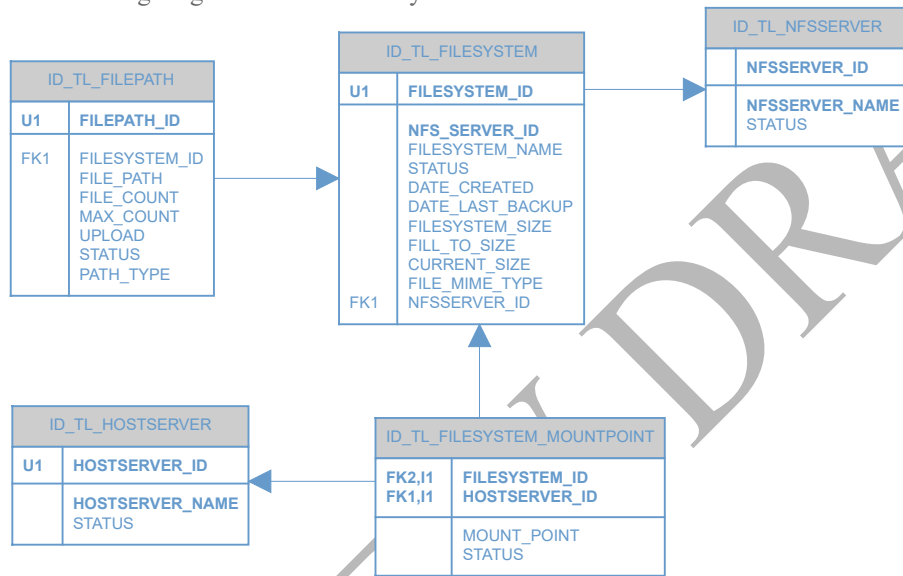


Diagram 3. Filesystem Tables

## Hardware and Software

Recommended database configuration is as follows:

Oracle 8i (8.1.6.3) Enterprise Edition and standard software as follows.

Sun Enterprise 220R or 420R systems running Oracle Parallel Server for the Aardvark Name Server (ANS), the look-up tables and the metaphysical data.

**Note...** It is possible to separate the look-up tables and metaphysical data into other database instances. Sun Enterprise 220R or 420R systems running Oracle Standard Edition for the individual account data.

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.



---

## Scheduler

### Overview

The Scheduler is an application that performs tasks scheduled by the i-drive application. The Scheduler executes tasks that are scheduled to take place in the future. Such tasks include:

- Automated reminders (account specific e-mails).
- Deleting expired, half-created accounts (from temporary i-drives).

To perform scheduled tasks, the Scheduler looks for tasks that are overdue, and then executes them.

### The Scheduler Portal

The Scheduler Portal is a Web-based interface that allows you view scheduled tasks. From a working application server, you can view the Scheduler Portal at: <http://www.idrive.com/mydrive/debugScheduler.jsp>. From here, you can:

- View all jobs currently scheduled.
- Change status of each job.
- Search for certain types of jobs by searching in the XML description.
- Search for jobs meeting certain time criteria (For example, expired or recently scheduled jobs).
- Delete or reschedule jobs.

**More Info...** Refer to the i-drive Administrator's Tools Guide for more information on how to use the Scheduler Portal.

### Functionality

The i-drive application places requests for scheduled work in the Scheduler queue on each account database (scheduled work for a given user is stored on his/her Account Database). Each request contains a 'firing time' (time to run the job) and an XML description of the data required to identify the class that will run the job, and describe the data required for the class that will run the job. The Scheduler acquires jobs from the queue whose firing time has passed, loads the class that will execute the job, and then executes the job. The job is then removed from the queue (or is rescheduled if it failed).

### Hardware and Software

The Scheduler configuration is as follows:

- Java 2 server application.

- The Scheduler runs on an application server.

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.

### Storage Zone

This section overviews each component of i-drive's storage zone, including:

- Transfer Server
- Metadata Extractor
- GarbageMan

REVIEW DRAFT

---

## Transfer Server

### Overview

i-drive uses the Transfer server to handle all file data transfers to and from users, and to communicate with the i-drive file system. The Transfer server is a logical peer to the application tier's Web and application tiers. The Transfer server is used to:

- Send file data to clients (download and streaming).

- Fetch data from third party servers (sideload).

- Retrieve files from clients (upload).

When users request downloads from the Application zone, the application sends an HTTP redirect to the client. The redirect tells the client browser to bounce back to the Transfer server and initiates the download. The redirect URL contains an encrypted string, called a download ticket, and has embedded information needed by the Transfer server. The download ticket is time stamped so that the URL cannot be captured and re-used. The download ticket also contains the service level at which the Transfer server is to deliver the file to the user.

Uploads occur similarly, except that the download ticket parameters are carried in HTTP POST data rather than a redirect URL.

### Features

The major features of the Transfer servers are:

#### Bandwidth Allocation

The Transfer server implements a bandwidth allocation policy. Every file transfer has allocated bandwidth from a fixed total based on priority supplied in the transfer ticket. i-drive can limit the transfer rate at which a customer is able to download or upload files, which gives i-drive a level of cost control not available with other servers. Bandwidth allocated to our customers can be changed at any time without service interruptions. Operators can also use bandwidth allocation to set higher transfers rates for preferred customers (or for streaming, etc.).

#### Custom HTTP/HTTPS Servers

Transfer servers are custom HTTP/HTTPS. They can multiplex large numbers of concurrent network connections but only use a small number of OS threads and other resources. For example, a single Transfer server can handle well in excess of 3000 simultaneous transfers on a Sun Ultra 60.

#### Transfer Authentication

Transfer authentication is provided by an external entity, typically the i-drive application, and generates a transfer ticket (encrypted identifier string), which allows the user to perform the transfer. The ticket also contains upper (and possibly lower) bounds for allocated resources. (For example, bandwidth). The Transfer server can also authenticate transfers without a ticket by contacting the application zone for authorization.

#### Transcoding

The transfer server also automatically transcodes content to users when requested, using an extensible transcoding plugin architecture. For example, a user might request a Microsoft Word document from his cell phone. The Transfer server can extract out the text from the document and send only the content, which the cell phone can view. Or a 4MB JPEG might be converted to a 200K GIF file for delivery to a mobile device.

#### File Repository Access

All files served by a farm of Transfer servers must be addressable as local file systems. Transfer tickets tell Transfer servers where a file is located.

## Hardware and Software

The Transfer Server configuration is as follows:

Java 2 server application.

The Transfer tier runs on Sun Ultra 60 servers, with a 50-60% bandwidth configuration for each server. Each Transfer server can support 60-70 megabits of bandwidth.

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.

REVIEW DRAFT

---

## Metadata Extractor

### Overview

The Metadata Extractor extracts information out of physical files when they enter the i-drive system. When a new file is uploaded, sideloaded, or mailed to the i-drive system, the Aardvark Name Server stores a record of the transfer. The extractor notices the new file and extracts its metadata. This metadata enables contextual advertising/promotion of content, community, services, and applications. It also enhances information navigation (For example, you can view by source, by artist, etc.). In addition, it is used for SingleSource technology to group files.

### Physical File Data

Analysis of the physical file data (and resulting metadata) is as follows:

- Extraction from structured files
- MP3 files: band, genre, song title, song length
- Images: size, statistics, comments
- Documents: keywords, author summary
- Inference from semi-structured content
- MD5 hashing

When the extractor extracts data, the resulting metadata is stored with the appropriate physical or logical file.

### Hardware and Software

The Metadata Extractor configuration is as follows:

Java 2 server application.

The Metadata Extractor runs on Sun Ultra 60 servers. One extractor is required per distinct grouping of file systems (For example, one per data center).

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.

---

## GarbageMan

### Overview

The GarbageMan is an application used to delete physical files and related information (physical data) from the i-drive system. The GarbageMan frees physical storage for reuse by new files. Since i-drive employs SingleSource technology, physical data can only be deleted when there are no logical files that reference some physical file.

**Note...** This process is similar to garbage collection (automatic memory allocation and de-allocation) in Java and many other modern languages.

### Functionality

When a user requests a file to be deleted from his/her i-drive account, the logical file is removed from the users account and placed in a shared (per Account Database) trash table. The GarbageMan acquires trash candidates by taking files out of the trash table. The GarbageMan then queries each account database to determine if the given logical file is not in use on all the account databases. If the logical file is not referenced it can be collected. The GarbageMan then deletes all the physical data, including metadata and the physical file itself.

### Hardware and Software

The GarbageMan configuration is as follows:

- Java 2 server application.

- The GarbageMan runs on a Sun Ultra 60.

**More Info** For detailed installation and configuration information, refer to the *i-drive ESS Installation and Configuration Guide*.

**More Info** For detailed tool and monitoring information, refer to the *i-drive ESS Application Administration Tools Guide*.

---

## Network-Attached Storage (NAS)

### Overview

i-drive requires network-mountable storage for the storage of physical file data. This typically means the use of the Network File System (NFS). Within a given Storage Zone, all transfer servers must mount all available storage. However, an installation may include any number of Storage Zones.

REVIEW DRAFT

---

## Network Zone

This section overviews each component of the ESP network zone, including:

- Content

- Data

- Administrative

REVIEW DRAFT



## Network

### Overview

An i-drive ESP network is designed as a scalable, reliable and redundant network. The ESP network includes the following two core pieces to the i-drive application:

**The content network:** Made up of Alteons, Catalyst, Foundries, load balancers, Web servers, application servers, and Transfer servers.

**The data network:** Made up of databases and Network-attached Storage (NAS) devices.

### Content Network

The content network is divided into several nodes. New nodes can be added when load dictates. Each node is a complete and independent vertical slice of i-drive ESP functionality and includes a load balancer for Web and download requests, Web servers, application servers, and Transfer servers. The nodes are load balanced at the DNS level by a 3DNS: when clients request <http://www.idrive.com/>, it returns an address of one of the load balancers.

Each node is a unit of scalability in i-drive. As traffic increases, new nodes can be added until we reach a maximum of about 8GB of traffic. Each node offers fault tolerance. If a load balancer or a portion of the network on a node fails, the other nodes continue to operate seamlessly to the end user. In addition, each node can be augmented with additional servers for horizontal scaling.

### Data Network

The data network handles all the traffic for database request processing and holds the NFS servers exported from the NAS. There is no routing between the content and data networks. Database requests initiated by the application servers go directly to the data network through an interface card on that network. Requests to read or write files through the NAS go directly onto this network.

The data network communication is point-to-point rather than routed through a single device, so the potential is not the same for network bottlenecks, therefore it is not segmented into different nodes.

### Administrative Network

The administration network contains the administrative and support systems needed to monitor and maintain the network. It is a single system with an interface connected to all networks (referred to as the admin machine) and has direct access to all machines.

### Network Systems

The following table lists the various network systems that exist, as well as their redundancy and access requirements:

Table	Description
Web servers	Redundant at the system level and application level inside and across nodes. Individual systems do not require redundant network access. They reside in content network only. Public access to ports 80 and 443.
Application servers	Redundant at the system level and application level inside and across nodes. Individual systems do not require redundant network access. They reside in content and data networks. No inbound public access, but servers to make requests to external systems.

Transfer servers	Redundant at system level and application level inside and across nodes. Individual systems do not require redundant network access, however, interfaces in multiple nodes maybe used to provide better fault tolerance. They reside in content and data networks. Public access to ports 80 and 443.
Database Servers	Redundant at system level and application level for the standby databases and the account database and through clustering for the Aardvark Name Server. Systems have redundant access paths. They reside in data network only.
NAS (NFS) Servers	Redundancy at system level and application level provided by NAS. Individual systems have redundant access paths. They reside in data network only.
Public DNS	Redundancy at system level and application level provided by DNS master/slave configuration. Redundancy at network level provided by different systems connected to different switches. Public access for DNS queries, no internal access for DNS.
Internal DNS	Redundancy at system level and application level provided by DNS master/slave configuration. These systems fall back on the public DNS for unhandled/unknown requests. No public access is allowed. All internal production systems relay on for <b>anem</b> lookups.
Mail servers	Redundancy at system level and application level, with different systems accessible through different network paths. Mail servers are load balanced for reliability and performance.
Metrics	No redundancy requirements
Database Backup	Redundancy on application level through backup software. Administrative access only. Backup is located on the backup network.
Internal Administration (ESS operations homebase):	Redundancy at system level and on application level. Administration access only.
Monitoring systems: OpenView, BigBrother, MRTG, etc.	Redundancy at system level. System needs to communicate with every other system in the data center.
QA Application Slice	Duplicates the network topology of a content network node.
Switch Administration	Dedicated connections for switch/router configuration. Administration access only.
APC Administration	Dedicated connections for APC use. Administration access only.

---

## System Requirements and Tools

This section overviews i-drive s system requirements and tools, including:

- System Components and Hardware

- System Administration Tools

REVIEW DRAFT

## System Components and Hardware

The following section overviews recommended i-drive ESP component and hardware configuration. These suggested requirements reflect the components and hardware needed to operate the i-drive system based on historical performance, storage capacity, and usage needs.

### Application Zones Overview

Component	Hardware	Description
Application server	Sun Ultra 60	UltraSparc II/450 Solaris 2.6 2 CPU 4 GB Executes Java process
Account database traffic fabric	Cisco Catalyst 6509	L2 Switching Handles brunt of account database traffic
Account databases	Sun E4500	Solaris 2.6 Oracle 8i 12 CPU UltraSparc II/450 12GB

## Application Zones Diagram



Diagram 4. Application Zones Diagram

## Storage Zones: Logical Network

Component	Hardware	Description
Load-balancing DNS	F5 3DNS	DNS Global load balancer Monitors load balanced services for availability
Network demarcation	Foundry BigIron 4000	Internet feed Filtering Conservative 2GB capacity
Node/Transfer traffic distribution	Alteon 180E	Virtual IP/L4 switching Transfer-tier load balancer, clustering Fail-over to other Alteon 180E
Transfer-tier traffic fabric	Cisco Catalyst 6509	L2 Switching Handles brunt of Transfer-tier traffic

Transfer server	Sun Ultra60	UltraSparc II/450 Solaris 2.6 2 CPU
		4 GB Transfer Server

### Storage Zone Diagram

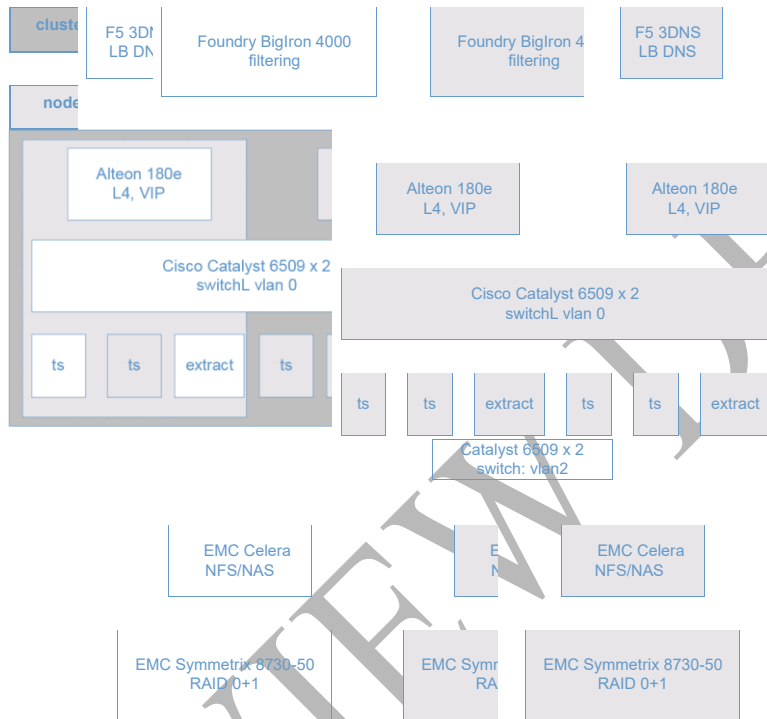


Diagram 5. Storage Zone Diagram

### Application Zones: Hardware and Network Overview

Component	Hardware	Description
Load balancing DNS	F5 3DNS	DNS Global load balancer Monitors load balanced services for availability
Network demarcation	Foundry BigIron 4000	Exodus feed entry point Filtering Conservative 2GB capacity

Node/Web traffic distribution	Alteon 180E	Virtual IP/L4 Switching Web-tier load balancer, clustering Fail-over to other Alteon 180E
Web-tier traffic fabric	Cisco Catalyst 6509	L2 Switching Handles brunt of Web-tier traffic
Web Server	VA Linux 2200	Pentium III/350 1 CPU 512 MB Apache 1.3.9 + Apache JServ
Application-tier fabric	Cisco Catalyst 6509	L2 Switching Handles brunt of application-tier traffic
Metadata extractor	Sun Ultra60	UltraSparc II/450 Solaris 2.6 2 CPU 4 GB Executes Java process
Physical file traffic fabric	Cisco Catalyst 6509	L2 Switching Handles brunt of physical-tier traffic

### Application Zones: Hardware and Network Overview Diagram

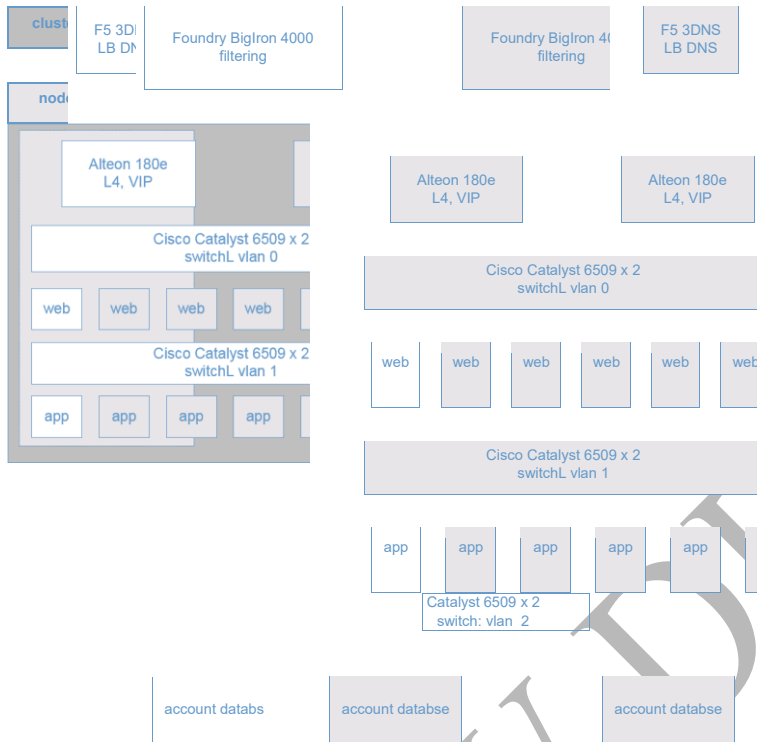


Diagram 6. Application Zones: Hardware and Network Overview Diagram



## System Administration Tools

**Overview** i-drive offers the following Web-based administration tools to manage i-drive system components.

Tool	Description
Ad Manager	Used for trafficking targeted advertising within the i-drive interface.
Domain Administration	Used to administer domain and account specific information. You can add new accounts or modify existing accounts.
Free Stuff Manager	Used for trafficking FreeStuff items for supporting content distribution. FreeStuff is a read/only community part of the PC-ESS UI that allows content distribution.
iConsole	Used as the main interface to access i-drive s application and system reporting tools.
Metadata Extractor	Used to administer the extractor and monitor its progress.
Partner Portal	Used to administer partner products.
Portal Transferserver	Used to display various internal statistics and configuration variables.
Scheduler Portal	Used to administer scheduled jobs. You can change job status, search for jobs, and make changes to jobs.

**More Info** For detailed tool information, refer to the i-drive ESS Application Administration Tools Guide.