

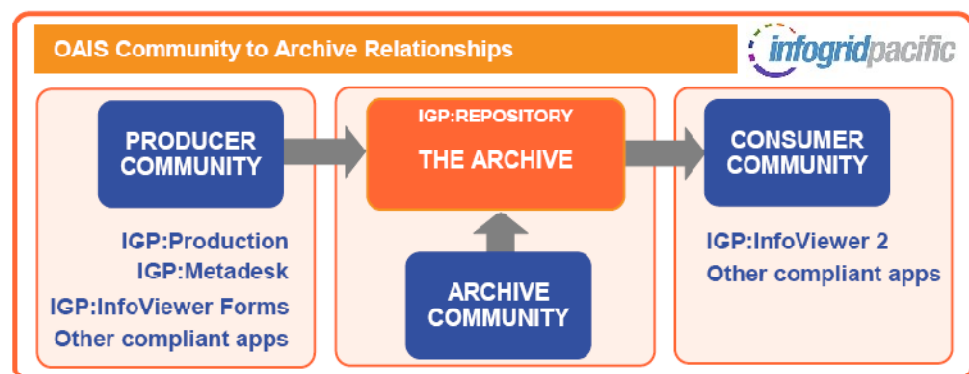
Overview

IGP:Repository is an advanced digital preservation archive designed for critical, demanding, long-term data archiving for a wide range of organization requirements. IGP:Repository successfully isolates content and content management from technology and technology obsolescence enabling the modern enterprise for a data-certain future.

It is purpose designed for:

- Document management including images, office documents, maps, etc.
- Asset management including images, audio and video
- Records management with statutory compliance requirements
- Archiving cultural artifacts (as digital surrogates) for museums and formal archives
- Maintaining large data sets, including mixed datasets

The design is a faithful execution of the OAIS Reference Model for digital archives. The benchmark for information system archives. IGP:Repository complies with a number of international standards for document and records management . It is designed specifically as a content management foundation to empower any organization to institute a best practices business model related to the preservation of data or content that is defined as valuable for the organization.



IGP:Repository can be deployed alone, in conjunction with other Infogrid Pacific solutions, or stand-alone as a Web Service integrating with other Web Services and applications.

The highly modular structure enables it to be scaled from small to large implementations. The unique Storage Zone architecture allows the addition of storage, migration of Information across and between technologies at any time.

IGP:Repository delivers replicated services up to hundreds of terabytes. It can be deployed on affordable commodity grade hardware, or work on enterprise class NAS and SAN storage devices.

About this Specification

IGP:Repository is a feature rich product, but unlike many other products these features are not visible in an interface. IGP:Repository is an extremely configurable digital preservation storage solution, this features specification is designed to allow the product to be evaluated for a wide range of applications. Therefore this document is narrative rather than bullet pointed to ensure the feature working can be understood.

Policy Driven

For a trust model to exist an archive must be able to operate without human intervention in the file management processes. IGP:Repository achieves these aims with a powerful policy based approach, (For those familiar with ISO 15489, policies are equivalent to mandates) which effectively means no human can easily know exactly where and how any specific file is stored. This is further abstracted through replication.

Policy Defined

Specifically in the IGP Information Product environment:

“A Policy is a set of controlling and limiting business rules that have been stated in a manner which allows them to be understood by, and enforced by an autonomous computer system to enable the management of Information Packages by computer systems over time with controlled and audited limits on the actions of competent and authorized people.”

By its nature a policy system is a constrained system and executes only the stated business rules exactly as stated and nothing more or less. It never says what cannot be done, only explicitly what can be done.

Authoring of policy statements and conversion of those statements into computer interpreted business rules are a specific requirement of all IGP information systems.

Policy Creation

Decisions on how and what the information products will do must be made from the business context before they can be usefully put to work in a technology context. Policies once set cannot be tampered or changed, therefore a formal approach to defining collection policies is useful and consideration of business purposes early in the information design development life-cycle is important, but possibly non-intuitive.

To simplify the policy authoring process IGP Information Design method has a number of standard template policy statements and their associated business rules available for a range of business ECM objectives. These documents can be useful as part of an ISO 9000 or other process quality accreditation program.

Standards Compliance

IGP:Repository has been created to conform to the requirements of a number of international standards. When used in an appropriate environment, with correct management and administration procedures, IGP:Repository is capable of fulfilling the highest trust obligations.

CCSDS 650.0-B-1 (also **ISO 14721:2003**) is not strictly a standard but a Reference Model created and maintained by CCSDS (Consultative Committee for Space Data Systems). The short title is the OAIS (Open Archival Information System). The IGP:Repository implementation is very close to the OAIS model to ensure all main requirements are included, especially with respect to life-cycle maintenance and trustability of stored data. We have deliberately maintained the OAIS vocabulary (as ugly as it is in places) to make it easy to relate our product implementation with the Reference Model.

METS is a US Library of Congress Standard. We use it for the standard Information Package within IGP:Repository. The value of METS is that Information Packages of any complexity can be created, those packages can easily replicate or migrate across systems and all Information Packages effectively become self-managing. METS is rapidly becoming the standard in national, educational and other significant archives.

ISO14589 Records Management. This is the definitive standard for any systems supporting records management. It defines not only the way data must be handled when it is regarded as a record, but also the environmental issues that establish the trustworthiness of digital records over time.

ISO 23081-1 Information and documentation — Records management processes — Metadata for records —Part 1:Principles. This standard sets a framework for creating, managing and using records management metadata and explains the principles that govern them. This defines how IGP:Repository manages life-cycle digital provenance metadata. This is a standard well worth owning as it not only relates to technology, but also organizational practices.

ISO/TR15801 Electronic imaging — Information stored electronically — Recommendations for trustworthiness and reliability. This Technical Report defines recommended practices for electronic storage of business or other information in image form. As such, complying with its recommendations is of value to organizations even when the trustworthiness of the stored information is not being challenged.

DoD 5015.2-STD. Design Criteria Standard for Electronics Record Management Software Applications. This Standard sets forth mandatory baseline functional requirements, and identifies non-mandatory features deemed desirable for Records Management Application (RMA) software. This is a certification program for companies wishing to supply solutions to the US military. We do not directly support all functions and have no intention of having our software certified, however the practices outlined are used in our application designs.

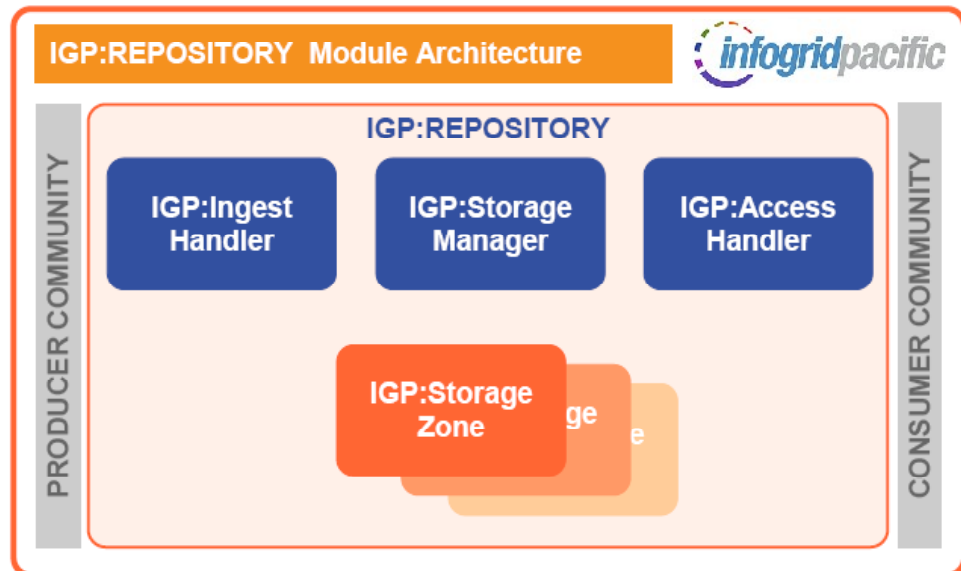
RFC 3161. Internet X.509 Public Key Infrastructure Time-Stamp Protocol (TSP). This document describes the format of a request sent to a Time Stamping Authority (TSA) and of the response that is returned. It also establishes several security-relevant requirements for TSA operation, with regards to processing requests to generate responses. There are several national and international standards such as ANSI X9.95:2005 and ISO/IEC 18014-1:2002 dealing with particular issues related to financial transactions. RFC 3161 is the foundation standard used for IGP:Evidence.

ISO 9564:1985 and **ISO 2766:1986** are both Guidelines for creating multi-lingual and mono-lingual thesaurii. We use these as part of the information design and directly guide the design of the IGP:Information Architect thesaurus builder. Where possible English terms are created as the "Exchange Language" and not a "Dominant Language"

Modules and Tools

IGP:Repository has five primary modules that can be deployed on the same or separate physical servers to address processing, storage and usage load balancing.

In addition there are a number of supporting applications to cover the widest possible requirements that may be faced by an organization.



Module	Explanation
<p>IGP:Distributor</p> <p>This is a stand-alone application that verifies applications trying to message between various IGP:Repository modules when deployed on separate machines or across the Internet.</p>	<p>IGP:Repository can be installed on a single physical server, modules can be deployed on different machines, and Storage Zones added at any time. To handle this flexibility IGP:Distributor along with the inbuilt IGP:Message Handler provides a number of services:</p> <ul style="list-style-type: none"> • Authentication of all participating applications • Authorization of all participating applications to communicate with each other and exchange data. This can include digital signature protocols. • Inter-module messaging and external gateway communication services. It also maintains the system-wide heart-beat monitors. • Choreography required to ensure system coherence and to balance resource heavy processes against lighter processes. It does this by enforcing a "Good Manners" protocol to ensure fast processes don't dominate processing queues over slow processes .
<p>IGP:Ingest Handler</p> <p>Arbitrates (accepts/rejects) and processes accepted SIPs to AIPs received from valid Producer communities</p>	<p>IGP:Ingest Handler is the only gateway into the Repository. It has the triple role of guarding the entrance, validating submissions to the archive and converting SIPs to AIPs for storage.</p> <p>To support its various validation and processing tasks there are a number of stand-alone configurable tools available. These are detailed separately.</p>

Module	Explanation
IGP:Storage Manager	<p>IGP:Storage Manager monitors Collection Storage, Retention and Disposal policies. It contains the main central database tables used to administer the entire system. In addition to its own data it maintains state data of other modules for fast and powerful reporting.</p> <p>All services must use Storage Manager to discover which Storage Zones AIPs are stored on. Storage Manager does not know the exact location of an AIP on storage Zone(s), only the Storage Zone(s) ID.</p>
IGP:Access Handler	<p>Provides the interfaces for external consumer applications (communities) to interface with the Repository. This includes authentication, authorization, and other constraints. Access Handler is the only method by which files can be delivered from the Repository. To do this it maintains specific parts of AIPs as DIPs. Only Distribution Information Packages are exposed outside the Repository.</p> <p>IGP:Access Handler Maps files and AIP Structure Maps to specific communities based on pre-defined Access Policies. It ensures consumer communities can only access the parts of an AIP to which they are authorized.</p> <p>IGP:Access Handler has Policy protocols with IGP:Ingest Handler for high-speed services from Ingest to Access when required or allowable.</p>
IGP:Storage Zone	<p>IGP:Storage Zones are the file management heart of the system. Any number of IGP:Storage Zones can be deployed for the ultimate disaster recovery system. It is specifically designed to take advantage of low-cost commodity hardware, but can be deployed on any level of hardware as all storage methods are independent of the hardware or Operating System technology.</p> <p>IGP:Storage Zone provides surface-refresh services by collection if configured. This means periodically files are moved and revalidated on the hard-disk surface. Storage Zones also deliver files through HTTP, FTP and Rsync services to authorized consumers.</p> <p>IGP:Repository is essentially a high-availability, spinning disk archive. However an IGP:Storage Zone could be deployed on a MAID server (Massive Array of Inactive Disks).</p>

Complimentary Applications

The following applications are provided as integral with IGP:Repository although they serve sub-functions to the primary Information Package / Community

Module	Explanation
IGP:Search	<p>A stand-alone application which indexes metadata and full text in a wide range of languages. It is the primary search and discover tool for consumer communities to access the IGP:Repository information.</p> <p>The implementation is based on the standard setting, open source Lucene Search Engine.</p> <p>IGP:Search indexes collection content based on the Discovery Policy. It can index any combination of metadata and/or full text search should it be required.</p> <p>The search engine can be externally exposed for consumer communities to discover content within the system.</p>

Module	Explanation
<p>IGP:Archive Administration</p> <p>A stand-alone client application that can be installed on Windows or Linux workstations. It is a fully featured administration appliance for the entire archive.</p>	<p>Provides the following features:</p> <ul style="list-style-type: none"> • Storage Zone to collection data space allocation • Collection policy generation for all policies • Control of Producers and Consumers accessing the system. • Administration of retention and disposal policies • Query, monitor and report system resource state and trends • General informational collection detail reports • System operation detail reports • AIP metadata maintenance (where permitted)
<p>IGP:Archive Administration Online</p> <p>This is a module provided with IGP:InfoViewer to allow Web Access to the Administration functions. It contains all the features and abilities of the client application</p>	<p>This is only available if IGP:InfoViewer is purchased and installed.</p>
<p>IGP:Information Architect</p> <p>Provides Localization and centrally controlled vocabulary creation and maintenance.</p>	<p>This tool allows the creation and application of controlled vocabularies. It is a Web Services application that can be installed on the same server as IGP:Distributor, or on its own server.</p> <p>Any number of multi-lingual thesaurii can be created and maintained. Because it is available as a Web Service remote producers can use the same data as the Repository, as can consumer search interfaces.</p> <p>Comes complete with dozens of useful generic “starter” thesaurii for a quick start in any specific deployment.</p>
<p>IGP:Process Content</p> <p>A powerful module based content processor that can be used on the Ingest Handler server, or deployed separately (for intensive processing requirements).</p> <p>It contains the following standard “Drop-In” processors:</p> <ul style="list-style-type: none"> • METS Packager • XML Validator • Component Checker • Vocabulary Checker • Format Convertor • SIP Report Generator 	<p>Current best practice is to ensure SIPs conform to the archive requirements at submission to reduce the programming and processing required to operate and maintain the archive. However it is not always possible to ensure this for all producer communities.</p> <p>Therefore a number of conformance, validation and processing tools are provided. However for any specific requirement these will have to be modified and customized on an installation by installation basis.</p>

Information Organization

Internally content is organized into collections (fonds). An Information Package can only belong to one collection to ensure consistency of policy enforcement. Information Packages can have relationships defined, including relationships between AIPs across collections.

As defined by the OAIS model. Information Packages can exist in three states during their association with the IGP:Repository system.

1. **SIP (Submission Information Package).** How information is presented to the Repository. The ingest handler inspects and conforms the SIP. If it passes all tests it is converted to an AIP else it is rejected.
2. **AIP (Archive Information Package).** How information is stored, maintained, replicated, migrated and life-cycle managed within the storage system.
3. **DIP (Dissemination Information Package).** How information is presented to consumer communities. A DIP may be only a part of the available information and different views may be constructed for different

Feature	Explanation
<p>Unlimited Collections (fonds)</p> <p>Collections are the “container” for AIPs. In addition collections have controlling policy rules which are applied to all AIPs within an collection.</p> <p>Collections are allocated storage space on one or more IGP:Storage Zones.</p>	<p>Collections can be created based on any information design requirement.</p> <p>EG. A Library may group by Document Genre, a Museum by Department, an Archive by Fonds and a business by simple organization groupings.</p> <p>Because collections are unlimited and policies apply to each collection, information diversity can be very broad.</p>
<p>Pre-defined collection policies</p> <p>This includes configuration of the following policies covering the gamut of Information archiving requirements:</p> <ul style="list-style-type: none"> • Ingest Policy • Retention Policy • Disposal Policy • Storage Policy • Access Policy • Rights Policy • Discovery Policy 	<p>Detail rich, pre-configured policy options cover the major operations of the archive</p> <p>Policies are created by filling in the appropriate forms in IGP:Repository Administration or IGP:Repository Administration Online.</p> <p>Policies must be configured for each collection. Policies completely automates the actions and operations of the archive (removes human intervention opportunities) to ensure it meets the trust criteria for records and cultural artifact storage.</p> <p>Standard Polices which are available for use are:</p> <ul style="list-style-type: none"> • Cache collection • Document collection • Records collection • Cultural collection • Media collection <p>Each of these specifies different policy parameters. These can be modified and customized on a per collection basis.</p>
<p>Metadata Support</p> <p>Supports any metadata schema at the information level, including mixed metadata schemas.</p>	<p>IGP:Repository is Metadata agnostic. It can handle any DTD/Schema that the deploying organization needs or wants.</p> <p>IGP:Repository does not try and understand metadata. Where required metadata key-values are extracted from metadata by the ingest processors and sent to the Search Engine for indexing.</p>

Feature	Explanation
	<p>IGP:Vanilla Metadata is also provided. This is a powerful cross-walk metadata schema that allows DC, TEI, EAD and similar cross-walks to be generated. It is suitable for closed-loop systems that do not require information exchange with other systems. It also provides optimal performance throughout the system.</p>
<p>METS Packaging</p> <ul style="list-style-type: none"> • This delivers • Standardization • Information Robustness • Information Exchange • Powerful technology neutral replication and migration strategies 	<p>METS is a Library of Congress Standard defined for the description, encoding and transmission of complex content items. IGP:Repository uses METS internally to store all Information Packages (AIP). This means dynamic components such as databases and new systems can be built from the stored item, increasing the robustness, survivability of core content.</p>
<p>Storage Referencing</p> <p>Constituent objects of an AIP (Data Objects) may be stored in the system (EG an IGP:Repository Storage Zone), reference any other URL (file or location in another system), or be configured to reference a physical object in a physical location.</p>	<p>The one information package can be used to locate digital surrogates, reference related or associated information or be a direct reference to the stored Data Objects.</p> <p>What this means is a file does not have to be stored in the Repository to be referenced by the Repository. This is useful when information is federated between cooperating repositories, or information is at a known stable URL. This means data is not un-necessarily duplicated.</p>
<p>Data Objects</p> <p>Stored files are referred to as Data Objects. A Data Object is a single file stored in the system therefore it is the most atomic object that is maintained.</p>	<p>An AIP is generally a list of data objects with associated metadata, stored in a METS file.</p>
<p>Complex Information Packages</p> <p>Using the properties of METS the system allows very complex packages to be created and maintained. Depending on the archive type and policies these can be appended and modified over time.</p>	<p>A METS driven Information Package can contain any number of data objects with enormous complexity of Structure Maps. This is very useful when managing Information Assets such as digitized books.</p> <p>These can contain hundreds and thousands of data objects which have to be recombined in the Consumer context to make a coherent document. All these are reliably stored and managed over time by a single AIP with suitable bibliographic metadata.</p>
<p>Different Archive Operations</p> <p>Supports all major archive type requirements. These can be mixed in a single installation by collection.</p> <ul style="list-style-type: none"> • Fixed archive • Accumulating archive • Caching archive • Dark Archive 	<p>Not all digital repositories are the same. Some need to allow changing datasets, the application of research information metadata to the core Information package and many other authorized operations.</p> <p>IGP:Repository ensures that modification policies are enforced by data object type and no inadvertent modifications are applied to any asset.</p>

Policy Features

Ingest Policy

The following Ingest Policy parameter can be configured by Collection. Low processing validations are done in Ingest Handler. Heavy processing requirements are passed to a processor.

Feature	Explanation
Minimum Mandatory Metadata (M3)	This is the SIP identification information that Ingest Handler evaluates to even consider further processing. If there is any error in the M3 the SIP will be rejected.
Action Request Checking	If the SIP or SIP message request an action on a collection that is not allowed, the SIP will be rejected and a System Alert raised.
Maximum Package Size	Package size checking can be turned on. In this case a maximum acceptable package size must be defined. If the package size is exceeded the SIP will be rejected.
Maximum AIP Count	The maximum number of AIPs allowed in a collection can be set. If a SIP exceeds this limit it will be rejected and a System Alert raised.
Maximum Files per AIP	A file count limit for an AIP can be set. If the number of files in a SIP exceeds this count the SIP will be rejected.
Mandatory Component Check <ul style="list-style-type: none"> • Yes/No • If yes - Processor Path and name 	A list of mandatory components can be specified and checked. A component condition list must be generated. This can contain terms such as at least one file, file type matches (eg Images to MIX metadata), etc. This check is carried out by a pre-defined IGP:Process Content Drop-In application.
Excluded Mime Types	If a Mime type on the excluded list is submitted then the SIP will be rejected.
Included Mime Types	If a Mime type of any other type than on this list is submitted then the SIP will be rejected.
Validate XML <ul style="list-style-type: none"> • Yes/No • If yes - Processor Path and name 	The specified XML files will be validated. The schema must be available to the validator. Validation is done from an IGP:Process Content Drop-In application
Vocabulary Checking <ul style="list-style-type: none"> • Yes/No • If yes - Processor Path and name 	Specified metadata fields are checked that their terms are the correct vocabulary and syntax. If any term is incorrect the SIP will be rejected. (This particularly protects against hand-keyed metadata fields). Checking is done from an IGP:Process Content Drop in connected to IGP:Information Architect.
Generate METS Package <ul style="list-style-type: none"> • Yes/No • If yes - Processor Path and name 	If the submission has passed all tests but is not an AIP conforming METS package the SIP will be packaged or re-packaged to conform to a package specification.
Pre-process <ul style="list-style-type: none"> • Yes/No • If yes - Processor Path and name 	This determines what processing may be required on a package before submission of the AIP to storage. This may include operations such as generation of thumbnails, image resizing and format conversion. Pre-processing is done from an IGP:Process Content Drop-In.
Synchronize IV2	Passes a message with parameters to Access Handler to notify

Feature	Explanation
	IV2 to request the DIP
Enable Discovery	Indexes the metadata and or full text in IGP:Search to enable discovery to be carried out. Note this only happens after the AIP is stored and ready for Access

Retention Policy Features

The following Retention Policy parameter can be configured by Collection

Feature	Explanation
Retention Type	<p>This field of the Retention Policy defines the type of retention to be assigned to an AIP when it is stored.</p> <p>The current allowable values for this field are:</p> <p>Permanent – The AIP will be retained within the Archive forever.</p> <p>Retain for Statutory Period – The AIP will be retained for the period specified as per the Statutory requirements and will be made available for disposal.</p> <p>Retain and Review (Period) – The AIP will be retained for the period specified and will be made available for review to the authorized person/committee role.</p>
Retention Period	<p>The value of this field is the actual period for which the AIP should be retained before it is submitted for disposal or review.</p> <p>The period can be specified in years, months, weeks, days, hours and minutes.</p>
Retention Period Commence	<p>It is not necessary that the retention period commences as soon as the AIP is stored within the Archive.</p> <p>The value of this field determines either the date or event after which the retention period should commence. The allowable values are:</p> <p>On Ingest Date – The retention period commences on the date the AIP is stored within the Archive.</p> <p>On Document Date – The retention period commences on the date the AIP was first produced. This date can differ from the date of Ingestion. If there is no document date the Ingest date can be used as the fallback</p> <p>On Defined Event – The retention period will commence only when a defined event occurs. For example:</p> <ul style="list-style-type: none"> • If the AIP is an record of a law suite, the defined event could be the date the case was closed. • If the AIP is an record of an employee, the defined event could be the from cessation of employment. • If the AIP is an record of an bank account, the defined event could be the closure of the account. <p>If the Retention Period Commencement is set to Defined Event, the policy author will have to specify the event and its value. The event/value must conform to the existing event/value that the Archive recognizes.</p>

Feature	Explanation
---------	-------------

On Manual Trigger – The retention period will commence when an authorized person (could be the Collection Administrator) manually imposes the retention commencement.

Disposal Policy Features

The following Disposal Policy parameter can be configured by Collection. Disposal Policy is triggered when Retention Policy conditions are fulfilled.

Feature	Explanation
---------	-------------

Disposal Type

This defines the disposal method and can be one of:

Permanent – The AIP is permanent in the Archive and cannot have any disposal action applied even if there is a retention period applied.

Auto Dispose – The AIP is removed from the Archive automatically without the intervention of a human being.

Single Review and Dispose – The AIP will be presented for review to an authorized role. The disposal action will be determined by this role.

The policy author specifies the authorized role so that the Archive knows who to contact at disposal action time.

Multi-review and Dispose – The AIP will be presented to multiple authorized roles or a committee for review. The disposal action will be determined by these authorized persons or committee.

The policy author will have to specify the authorized roles or the committee so that the Archive can provide the AIP for review.

As there could be different implication on the disposal of an AIP like business requirements, statutory obligation etc., it is possible that there could be different authorities /committees that could be assigned depending upon the contents of the AIP.

Multi-vote and Dispose – Once the retention period expires, the AIP will be presented to multiple authorized persons or a committee to vote on. If even one person disagrees to the disposal of the AIP, the AIP will be frozen until an appropriate disposal action is decided.

Disposal Authority

This field contains the names/IDs of the authorized persons or the committees that need to review or vote before the disposal of an AIP.

Disposal Report

If true a disposal report is created and stored in place of the original AIP. The AIP is retained in the system and its type is changed. This provides an evidential trail of how, why and when the disposal action took place.

Freeze Disposal

A disposal can be frozen by application of a Freeze Lock.

Administration Features

IGP:Repository encourages and enables a formal approach to digital archive management. In addition to the technology it comes complete with an extensive range of “soft features” such as pre-defined Management Policy templates, Producer agreement templates and Consumer Community agreement templates.

Administration is responsible for execution of management policies within the operation of the archive and to the limits prescribed by the software. To assist with high-level administration tasks there are two Administration toolkits. A Client application and a Web Services application. Either or both can be used simultaneously as appropriate.

The administration applications provide all the tools required to administer all Repository functions.

Feature	Explanation
Administer Storage Zones Resources	The application implementation can be assessed against the most powerful formal archiving reference model that exists. IGP:Repository uses most of the OAIIS grammar in its modules and methods to ensure conformance
Create Collection Policies	A collection policy for a formal institution is derived after considerable discussion and formal definition of the purpose, scope and resource usage of a collection. Once a formal Archive Management decision is made the Collection Policy interface allows the administrator to capture the decisions in software executed business rules.
Manage Retention / Disposal activities	Archive administrators review system generated lists of Information Packages
Discovery Tools	<p>Powerful discovery tools allow the administrator to view and evaluate the system based on digital provenance, fixity and other internal archive metadata, not only descriptive metadata and full text content.</p> <p>In addition to being able to explore content and metadata, the administrator can also view content by archive specific items such as storage dates, access frequency, life-cycle maintenance tasks and a wide range of other administrator specific tasks.</p>
Automated alerts and warnings	<p>The system has a number of self-monitoring points which generate both interface and email warnings and alerts to a notification list. This includes:</p> <ul style="list-style-type: none"> General module failure Multiple storage level alerts Unauthorized ingest or access events <p>All modules run heart-beat monitors from the central messaging service to ensure that any failure is detected within the pre-determined monitor period.</p>

Technology

IGP:Repository uses best of breed Open Source applications which are multi-vendor supported and running enterprise critical applications around the world. Our own code is not open source but is visual source and is subject to the product EULA. This combination of an Open Source foundation with a commercially supported application brings the maximum benefits to the end user delivering a digital preservation archive of exceptional reliability and flexibility at the lowest possible total cost of ownership.

A digital archive assumes that content must out-survive technology, therefore the selected technology must be appropriate for the task of getting to the next technology boundary. It also needs to allow organizations to optimize technology costs and ensure budget is allocated into essential information services rather than underlying technology.

Performance of IGP:Repository is defined by more hardware, not more expensive hardware. Performance enhancement is carried out with horizontal scaling, something IGP:Repository is designed to do at all points.

Feature	Explanation
<p>Preferred Server OS: Linux</p> <p>Optional Server OS: Windows Server</p>	<p>It is recommended that Linux is deployed over Windows because of licensing issues, code transparency and sustainable cost rather than any technical reason. A digital repository creates its own demands per seat, and CPU licenses need to be avoided for sustainability reasons.</p> <p>Once configured and operational an IGP:Repository needs very little System Administration support because of the essential policy driven self-managing functions.</p> <p>The viable alternatives for the Operating System are MS Windows and Linux. Because Windows brings absolutely no benefits for this application the costs of a commercial OS should be avoided.</p>
<p>Internet Server: Apache</p>	<p>Apache needs no introduction, defence or explanation. It is the Internet Server of choice.</p>
<p>Web Framework: Django</p>	<p>Django is a high performance open source Web appliance that provides the performance required of a scalable archive. It is programmed in Python and is therefore in alignment with the application programming language, Operating System and other software components.</p>
<p>Programming Language: Python</p>	<p>Python is a high performance byte-compiled language. The source code is visible and available to the Repository owners.</p> <p>The licensing allows Repository owners to modify and maintain the code for their own purposes (under the terms of the EULA).</p> <p>Code transparency is a recommendation for long-term archives to ensure confidence in the continuity of the archive in the event the code stops being maintained.</p>
<p>Database: PostgreSQL</p>	<p>PostgreSQL is a high reliability, well supported open source alternative to Oracle.</p>