



SNOCAP Rights Management Service

Systems and Processes Enabling the Management and Distribution of Digital Content

Introduction

The SNOCAP Rights Management Service (RMS) allows any individual or organization to register digital content they own, establish terms of use for that content, and make that content available to third parties who wish to distribute it according to those terms. The RMS provides the basis for all other services offered by SNOCAP, including services where SNOCAP enables content owners to sell directly to consumers. This paper describes the capabilities, controls, and infrastructure associated with the SNOCAP RMS.

The following topics are presented:

- The capabilities offered by the SNOCAP RMS
- The systems that make up the SNOCAP RMS infrastructure
- How SNOCAP addresses scalability and performance requirements for the system
- How SNOCAP secures its interfaces, data and systems

The capabilities, interfaces and process flows associated with specific SNOCAP retail interface products are described more completely in separate whitepapers.

Capabilities Overview

The SNOCAP Rights Management Service is a repository of content and ownership rights accessed by a family of interfaces. These interfaces are built to serve the needs of parties that contribute content to the RMS and those who access that content. The capabilities offered by the service can be categorized based on who they serve—rights holders (content owners and the organizations that represent them) and distributors (retailers and parties that access rights information associated with content).

Capabilities for Rights Holders

- Rights holders register their content through a web-based application or supply content and metadata to SNOCAP for batch processing. Rights are registered with the system for the territories in which they are owned or controlled.
- SNOCAP uses specialized content identification technology to screen digital assets submitted by rights holders. This capability allows SNOCAP to detect when the same content is registered by multiple parties. Such content is flagged based on a variety of business rules that consider the type of organizations making the registration and the order in which the claims were made.
- Rights holders create and manage license information and apply those licenses to their catalog, indicating the terms under which the content can be distributed by retailers, or indicating that distribution is not allowed.
- Rights holders supply SNOCAP official master versions of digital assets (for example, audio files) for use by retailers licensed to access them.
- SNOCAP collects payments for purchased content from retailers and pays rights holders according to the wholesale price specified in the associated license terms.
- Per-retailer sales data is tracked and presented to rights holders.

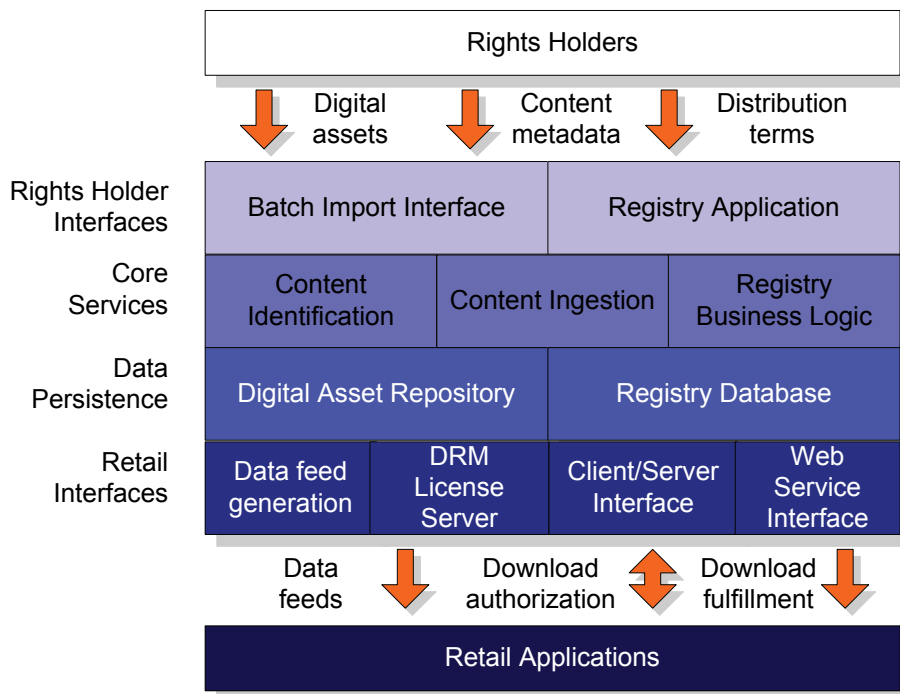
- When specified by the rights holder, SNOCAP manages licenses for audio files encrypted using the Windows Media format. DRM licenses are generated by SNOCAP according to the restrictions applied to the associated tracks.

Capabilities for Distributors

- SNOCAP enables any online service that wishes to identify or sell registered digital content the ability to do so according to the terms specified by the rights holder.
- SNOCAP provides application programming interfaces that allow applications (web-based or thick-client) to identify and obtain license information for registered content.
- SNOCAP fulfills downloads for authorized purchases using the official master files supplied by rights holders.
- SNOCAP provides client-side libraries to identify local content and correlate it to registered content in the RMS. Rights information is provided to the client application so that the content can be distributed or filtered according to the license terms.
- SNOCAP mediates the authorization and purchase of licensed content and maintains transaction information provided in reports to both the retailer and the rights holders.
- SNOCAP acts as a single payment interface between a retailer and all rights holders whose content was sold by the retailer.

Functional Components

The SNOCAP RMS can be described as a multi-layer set of interfaces, services, and persistent data stores. Rights holders add digital assets, content metadata, and distribution terms to the system through a set of rights holder interfaces. A series of core services support those interfaces and implement the business logic of the system. Both digital assets (for example, audio files) and rights information are maintained in databases and file stores, accessed by the core services. A set of retail interfaces also access the core services and expose functionality to web-based and thick-client retail applications. Each of these systems plays a role in delivering the capabilities of the RMS to rights holders and distributors.



Rights Holder Interfaces

Rights holder organizations access the RMS using either a batch interface (for delivering digital assets and associated metadata in bulk) or by using the Registry web application. These interfaces are described below.

Batch Import Interface

Rights holders with the appropriate SNOCAP account level that wish to register a large number of assets may use the Batch Import Interface. This interface allows rights holders to submit all the necessary content and ownership information to SNOCAP in an automated process. Typically, this interface accepts digitized content files along with a corresponding metadata feed in XML format. While SNOCAP prefers data to be delivered according to a standard schema, alternate formats can be accommodated by the plug-in architecture of the Content Ingestion System.

Registry Application

Rights holders access the RMS using the Registry web application. Retailers also have access to some functionality provided by this application, as described below. The following functions are available through this interface:

- Catalog management—rights holders upload new files into the system and specify associated metadata (for example, artist, title and album information, for audio content).
- Distribution term specification—rights holders define the terms that dictate how retailers can distribute content, including wholesale price. Rights holders with the appropriate SNOCAP account level specify sets of content that can be self-distributed through embeddable SNOCAP MyStore web applications.
- Opt-in to retail offers—retailers are able to define proposed business models in which rights holders may participate. Rights holders review retail offers and determine which retailers are licensed to distribute content according to specific distribution terms.
- View registration disputes—content registered by multiple rights holders is flagged. The RMS identifies the parties involved, allowing them to resolve the dispute and assign ownership appropriately.
- Access transaction reports—retailers and rights holders access month-to-date reports showing transaction details.

Core Services

The RMS business logic is implemented as a series of core services that are accessed by the rights holder and retail interfaces. These core services may invoke each other, and also access the persistence layer.

Content Ingestion

New content introduced into the RMS by rights holders is processed by the Content Ingestion System. This system accepts digital assets and metadata describing the content, either as a batch data drop delivered by rights holding organizations through the Batch Import Interface, or as small sets of content uploaded by individual rights holders through the Registry web interface. In either case, content delivered for ingestion is queued and processed by this system, which is responsible for identifying the content, checking for prior registration, storing digital assets in the Digital Asset Repository, and writing the metadata and distribution terms associated with the content to the Registry database.

Content Identification

The Content Ingestion System invokes the Content Identification System to determine whether a submitted piece of content has been previously identified, and if not, to record the content's fingerprint in a database for future access. The Content Identification System provides the ability

to identify audio content based on its acoustic data, avoiding misidentification due to erroneous metadata or changes to file format, filename or bit rate. The database of acoustic fingerprints is implemented so that inserts and lookup response time is linear even as the size of the database becomes very large.

Content identification can also be done by retail client applications. Applications that receive or share user-supplied content invoke a local SNOCAP content identification library to generate a fingerprint and look it up in the centralized fingerprint database. This allows local content to be correlated to rights information maintained in the RMS using a key based on properties of the content file that cannot be easily altered (for example, the acoustic properties of an audio file).

Registry Business Logic

The business rules that govern the behavior of the RMS are implemented as a set of services and programmatic logic that interact with the Registry application presentation layer, the Content Ingestion System, the retail interfaces, and the data persistence layer. This business logic includes:

- Rules and constraints that apply to the licensing of content in the system
- Actions, such as the generation of data feeds through the retail interfaces, that result when content is licensed for use by a retailer
- Access control at the content and application level
- Data access and presentation logic invoked by the Registry application
- Report generation logic and scheduling
- Accounting logic

Data Persistence

The RMS stores two principal types of information:

- Digital assets, such as audio files containing full length music tracks, and corresponding preview samples
- Relational data describing digital assets, organizations, and the rights information associated with them

These two persistence systems are described below.

Digital Asset Repository

In order to provide retailers access to content they are licensed to distribute, SNOCAP hosts copies of digital assets received from rights holders in a file repository. These files are typically cached using industry standard edge-caching services, and are correlated to entries in the Registry database where metadata and license terms are maintained. Access to these files are restricted so that only authorized downloads may take place.

Registry Database

The Registry Database contains all information about retailer and rights holder organizations, items of content, and the distribution terms that apply to the content. Additionally, transaction logs, report data, and accounting parameters are stored here. This database has strict access control, and uses standard replication technology to assure availability and scalability.

Retail Interfaces

SNOCAP has two interface products that allow third party applications to access and distribute content managed by the RMS. SNOCAP Linx is an interface for web-based applications. It is comprised of a periodic data feed sent from SNOCAP to the distributor, and a Web Service application programming interface (API) used by the retail application at run-time to conduct purchase and download transactions. The data feed contains metadata and pricing information about the content that the retailer is licensed to sell, and the Web Service API records the purchase and authorizes the download.

SNOCAP also has an interface product for thick-client applications. This product consists of a client-side plug-in that connects the application to the SNOCAP Client/Server Interface and performs local content identification, rights information access, and download authorization.

Data Feed Generation

As part of the SNOCAP Linx product, retailers may receive periodic data feeds describing the content they are licensed to sell. These data feeds are generated on a per-retailer basis, and contain information describing the digital assets the retailer can offer and the distribution terms that apply to them. Retail systems download these data feeds and update their own databases from which they generate their consumer-facing ecommerce web sites. When consumers select items for purchase, the retail system invokes the SNOCAP Web Service Interface to authorize the transaction and fulfill the download. The Data Feed Generation System creates tables based on data in the Registry Database, using the Registry Business Logic to formulate the relationships and constraints of the data delivered to the retailer.

Web Service Interface

The real-time component of the SNOCAP Linx product is a Web Service Interface that is invoked by the retailer when the consumer purchases a digital asset. This API is implemented using the SOAP protocol, and a method is exposed that accepts a certificate containing purchase information generated by the retailer, and a certificate provided to the retailer by SNOCAP used for authentication purposes. The Web Service Interface validates the supplied data, processes the download request, records the purchase transaction information, and returns a signed, single-use URL that is used to download the purchased asset.

Client/Server Interface

The Client/Server Interface mediates interactions between third party systems using the SNOCAP client plug-in and the SNOCAP RMS. The Client/Server Interface exposes services through a network interface to which the SNOCAP client plug-in connects. The client plug-in, in turn, exposes those services to third party systems through a software interface to which the client application integrates. The following services are exposed by the Client/Server Interface:

- Identify files on the file system of the client computer
- Obtain licensing information for identified files
- Authorize the download or sharing of files based on rights holder distribution terms
- Search and download official master files from the SNOCAP Digital Asset Repository

DRM License Server

SNOCAP maintains a license server that provides keys to encrypted files in the Windows Media digital rights management format. Rights holders who wish to distribute files protected in this manner must provide SNOCAP unencrypted WMA files. The Content Ingestion System encrypts the files and stores them in the Digital Asset Repository. License data specifying the download and play restrictions enforced by WMA DRM are interpreted by the SNOCAP system and the corresponding license key is generated when the file is downloaded by the client.

Scalability and Performance Considerations

The SNOCAP RMS performance requirements encompass several dimensions:

- Accommodate high load on the Digital Registry web application
- Accommodate a large number of retail application instances accessing the system
- Concurrently deliver a large number of digital assets

Each of these areas is discussed below.

Web Application Load

The systems implementing the Registry web application scale linearly. Extra machines are added if excessive load is reached. Performance is measured on these servers by monitoring the average response over a rolling time window. When a predefined threshold is hit, SNOCAP operations staff are notified and additional servers are deployed. SNOCAP maintains a pool of reserve servers for this purpose.

Retail Application Load

The time-critical action taken by retail applications is the processing of download authorizations. The systems involved in this process vary depending on the nature of the retail application. Thick-client retail applications use a client plug-in to establish a connection to the SNOCAP Client/Server Interface. The servers making up the Client/Server Interface are clustered, and each server caches all data necessary to fulfill a download authorization request without involving other components. This allows the load to be linearly scaled by the addition of such servers.

Applications using the SNOCAP Linx product use a combination of a pre-generated data feed supplied by SNOCAP and the Web Service Interface. Since the data feed generation is a batch process, there is no time critical aspect to it. The systems implementing the SNOCAP Web Service Interface scale linearly. Extra machines are added if excessive load is reached. As with the systems serving the Registry Web Application, performance of these servers is measured by monitoring the average response to a SOAP invocation over a rolling time window. When a predefined threshold is hit, SNOCAP operations staff are notified and additional servers are deployed. SNOCAP maintains a pool of reserve servers for this purpose.

Content Download Processing

Content files are served from the SNOCAP Digital Asset Repository and are cached using industry standard edge server technology. This approach satisfies high load requirements by offloading content delivery to the edge caching platform, which provides optimal performance, 100% uptime, and tremendous scalability. The caching system uses the DNS-based network intelligence mechanism to direct end users to the cache for all requests for objects from the Digital Asset Repository, pulling content from the Repository only when it has not yet been cached or an updated version of the asset is available.

Security Considerations

SNOCAP RMS security considerations encompass the following:

- Fraudulent access to the Registry web application
- Fraudulent registration of digital content
- Unauthorized access to digital assets
- Physical security of the facilities housing SNOCAP computing systems

These considerations are discussed below.

Web Application Security

Access to the SNOCAP Registry web application requires user-supplied credentials, which must be greater than a defined minimum length and require periodic updates. Accounts in the SNOCAP system require registration and approval prior to their creation. Third-party identity verification services are used to validate information about the user. If insufficient verification data is obtained, a SNOCAP customer care representative reviews the application.

Content Registration Security

User accounts are tiered based on the type of organization with which the account is associated. Individual content providers have limited capabilities in the system relative to larger organizations (for example, independent and major record labels), and content they register must correspond to the artist names they represent at account creation time. Furthermore, audio content that is uploaded into the SNOCAP system is identified using the SNOCAP Content Identification System. This system provides the ability to identify audio content based on its acoustic data, and therefore is not susceptible to misidentification due to erroneous metadata or changes to file format, filename or bit rate. The fingerprinting infrastructure maintains a database of fingerprints which is updated as new content is received. This system allows for the detection of content registered by multiple parties, and flags such instances for resolution based on a variety of business rules that consider the type of organizations making the claim and the order in which the claims were made.

Digital Asset Access Security

Several security measures in the SNOCAP Retail Interfaces minimize exposure to fraudulent download requests.

- SNOCAP issues a signed service certificate to the retailer. Download authorizations made without a valid certificate are rejected. Service certificates may be revoked by SNOCAP at any time.
- The service certificate contains a public key provided by the retailer. This certificate is passed to SNOCAP as part of a download authorization request, along with a purchase certificate signed with the retailer's private key. Upon receipt, SNOCAP verifies the purchase certificate signature with the public key embedded in the retailer's service certificate.
- Purchase certificates used for download authorization contain signatures for the license and content being authorized. These signatures are time sensitive and are known only by the retailer through the data feed provided by SNOCAP.

When a download authorization has been granted to the retail application, a URL to the requested asset is returned. SNOCAP protects against unauthorized use of this URL using state of the art security capabilities offered by industry leading edge-caching services. SNOCAP uses a centralized authorization mechanism, passing data used to authorize the request in the query string. The query string used by the retail application to access the requested content contains signed information that, with a secret key shared between the requesting application and the

SNOCAP system, is used to validate that the information sent in the request is valid. If so, the RMS authorizes the edge caching system to deliver the digital asset.

Physical Security

SNOCAP maintains a high degree of security in our facilities, systems and processes to ensure that the content and services we provide are protected from unauthorized access and tampering.

System Location

SNOCAP's production environment is located in a SAS 70 certified co-location facility. The facility is monitored 24x7 by co-location facility personnel using CCTV systems. Network and system monitoring is done by SNOCAP.

Protection Against Unauthorized Local Access

Access to the co-location facility is limited to personnel on an authorized access list maintained by our co-location facility and SNOCAP's Director of Operations. SNOCAP's production environment is located within a locked cage. SNOCAP will limit direct hardware access to the production environment to SNOCAP's Director of Operations, Database Administration staff and SNOCAP IT staff. In addition, physical hardware access may be granted at the discretion of SNOCAP's Director of Operations to maintain racks, power systems, cabling and contain physical damage. All access to machines is logged over the network to a secure logging server.

Protection Against Unauthorized Network Access

SNOCAP's production environment is segmented into multiple VLANs to limit potential security breaches. The systems are behind a firewall that prevents direct access from the Internet except for public network services. Each machine has a small number of authorized user accounts. Machines are accessed using SSH only. All access to machines is logged over the network to a secure logging server.