INTEROPERABLE SYSTEMS AND METHODS FOR PEER-TO-PEER SERVICE ORCHESTRATION

RELATED APPLICATIONS

[001] This application claims priority from commonly-assigned U.S. Provisional Patent Application Nos. 60/476,357, entitled "Systems and Methods for Peer-to-Peer Service Orchestration," by William Bradley and David Maher, filed June 5, 2003, and 60/504,524, entitled "Digital Rights Management Engine Systems and Methods," by Gilles Boccon-Gibod, filed September 15, 2003, both of which are hereby incorporated by reference in their entirety; these two U.S. provisional patent applications form part of the instant description/specification and are attached hereto as Appendix A and Appendix B, respectively.

COPYRIGHT AUTHORIZATION

[002] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND

[003] Networks such as the Internet have become the predominant medium for the delivery of digital content and media related services. The emergence of standard web services protocols promises to accelerate this trend, enabling companies to provide services that can interoperate across multiple software platforms and support cooperation between business services and consumers via standardized mechanisms.

[004] Yet, significant barriers exist to the goal of an interoperable and secure world of media-related services. For example, multiple, overlapping de facto and formal standards can actually inhibit straightforward interoperability by forcing different implementations to choose between marginally standard, but otherwise incompatible, alternative technical approaches to addressing the same basic interoperability or interconnection problems. In some cases these incompatibilities are due to problems that arise from trying to integrate different generations of technologies, while in other cases the problems are due to market choices made by different parties operating at the same time but in different locales and with different requirements. Thus, despite standardization, it is often difficult to locate, connect to, and interact with devices that provide needed services. And there are frequently incompatibility issues between different trust and protection models.

[005] While emerging web service standards such as WSDL (Web Services Description Language) are beginning to address some of these issues for Internet-facing systems, such approaches are incomplete. They fail to address these issues across multiple network tiers spanning personal and local area networks; home, enterprise, and department gateways; and wide area networks. Nor do they adequately address the need for interoperability based on dynamic orchestration of both simple and complex services using a variety of service interface bindings (e.g., CORBA, WS-I, Java RMI, DCOM, C function invocation, .Net, etc.), thus limiting the ability to integrate many legacy applications. The advent of widely deployed and adopted peer-to-peer (P2P) applications and networks further compounds the challenges of creating interoperable media-related services, due in part to the fact that

there is no unified notion of how to represent and enforce usage rights on digital content.

SUMMARY

[006] Embodiments of the systems and methods described herein can be used to address some or all of the foregoing problems. In one embodiment, a services framework is provided that enables multiple types of stakeholders in the consumer or enterprise media space (e.g., consumers, content providers, device manufacturers, service providers) to find each other, establish a trusted relationship, and exchange value in rich and dynamic ways through exposed service interfaces. Embodiments of this framework—which will be referred to generally as the Network Environment for Media Orchestration (NEMO)—can provide a platform for enabling interoperable, secure, media-related e-commerce in a world of heterogeneous consumer devices, media formats, communication protocols, and security mechanisms. Distributed policy management of the service interfaces can be used to help provide trust and security, thereby facilitating commercial exchange of value.

[007] While emerging web service standards are beginning to address interoperability issues for Internet-facing services, embodiments of NEMO can be used to address interoperability across multiple network tiers spanning personal and local area networks; home, enterprise, and department gateways; and wide area networks. For example, NEMO can provide interoperability in one interconnected system using cell phones, game platforms, PDAs, PCs, web-based content services, discovery services, notification services, and update services. Embodiments of NEMO can further be used to provide dynamic, peer-to-peer orchestration of both simple and complex services using a variety of local and remote interface bindings

(e.g. WS-I [1], Java RMI, DCOM, C, .Net, etc.), thereby enabling the integration of legacy applications.

[008] In the media world, the systems and interfaces required or favored by the major sets of stakeholders (e.g., content publishers, distributors, retail services, consumer device providers, and consumers) often differ widely. Thus, it is desirable to unite the capabilities provided by these entities into integrated services that can rapidly evolve into optimal configurations meeting the needs of the participating entities.

[009] For example, diverse service discovery protocols and registries, such as Bluetooth, UPnP, Rendezvous, JINI, UDDI, and LDAP (among others) can coexist within the same service, enabling each node to use the discovery service(s) most appropriate for the device that hosts that node. Another service might support IP-based as well as wireless SMS notification, or various media formats (MP4, WMF, etc.).

[010] Embodiments of NEMO satisfy these goals using peer-to-peer (P2P) service orchestration. While the advantages of P2P frameworks have been seen for such things as music and video distribution, P2P technology can be used much more extensively.

[011] Most activity in web services has focused on machine-to-machine interaction with relatively static network configuration and client service interactions. NEMO is also capable of handling situations in which a person carries parts of their personal area network (PAN), moves into the proximity of a LAN or another PAN, and wants to reconfigure service access immediately, as well as connect to many additional services on a peer basis.

[012] Opportunities also exist in media and various other enterprise services, and especially in the interactions between two or more enterprises. While enterprises are most often organized hierarchically, and their information systems often reflect that organization, people from different enterprises will often interact more effectively through peer interfaces. For example, a receiving person/service in company A can solve problems or get useful information more directly by talking to the shipping person in company B. Traversing hierarchies or unnecessary interfaces generally is not useful. Shipping companies (such as FedEx and UPS) realize this and allow direct visibility into their processes, allowing events to be directly monitored by customers. Companies and municipalities are organizing their services through enterprise portals, allowing crude forms of self-service.

[013] However, existing peer-to-peer frameworks do not allow one enterprise to expose its various service interfaces to its customers and suppliers in such a way as to allow those entities to interact at natural peering levels, enabling those entities to orchestrate the enterprise's services in ways that best suit them. This would entail, for example, some form of trust management of those peer interfaces. Preferred embodiments of the present invention can be used to not only permit, but facilitate, this P2P exposure of service interfaces.

[014] In the context of particular applications such as DRM (Digital Rights Management), embodiments of NEMO can be used to provide a service-oriented architecture designed to address the deficiencies and limitations of closed, homogeneous DRM systems. Preferred embodiments can be used to provide interoperable, secure, media-related commerce and operations for disparate consumer devices, media formats, and security mechanisms.

[015] In contrast to many conventional DRM systems, which require relatively sophisticated and heavyweight client-side engines to handle protected content, preferred embodiments of the present invention enable client-side DRM engines to be relatively simple, enforcing the governance policies set by richer policy management systems operating at the service level. Preferred embodiments of the present invention can also provide increased flexibility in the choice of media formats and cryptographic protocols, and can facilitate interoperability between DRM systems.

[016] A simple, open, and flexible client-side DRM engine can be used to build powerful DRM-enabled applications. In one embodiment, the DRM engine is designed to integrate easily into a web services environment, and into virtually any host environment or software architecture.

[017] Service orchestration is used to overcome interoperability barriers. For example, when there is a query for content, the various services (e.g., discovery, search, matching, update, rights exchange, and notification) can be coordinated in order to fulfill the request. Preferred embodiments of the orchestration capability allow a user to view all home and Internet-based content caches from any device at any point in a dynamic, multi-tiered network. This capability can be extended to promote sharing of streams and playlists, making impromptu broadcasts and narrowcasts easy to discover and connect to, using many different devices, while ensuring that rights are respected. Preferred embodiments of NEMO provide an end-to-end interoperable media distribution system that does not rely on a single set of standards for media format, rights management, and fulfillment protocols.

[018] In the value chain that includes content originators, distributors, retailers, service providers, device manufacturers, and consumers, there are often a number of localized needs in each segment. This is especially true in the case of rights management, where content originators may express rights of use that apply differently in various contexts to different downstream value chain elements. A consumer gateway typically has a much more narrow set of concerns, and an end user device may have a yet simpler set of concerns, namely just playing the content. With a sufficiently automated system of dynamically self-configuring distribution services, content originators can produce and package content, express rights, and confidently rely on value added by other service providers to rapidly provide the content to interested consumers, regardless of where they are or what kind of device they are using.

[019] Preferred embodiments of NEMO fulfill this goal by providing means for multiple service providers to innovate and introduce new services that benefit both consumers and service providers without having to wait for or depend on a monolithic set of end-to-end standards. Policy management can limit the extent to which pirates can leverage those legitimate services. NEMO allows the network effect to encourage the evolution of a very rich set of legitimate services providing better value than pirates can provide.

[020] Some "best practice" techniques common to many of the NEMO embodiments discussed below include the following:

- Separation of complex device-oriented and service-oriented policies
- Composition of sophisticated services from simpler services
- Dynamic configuration and advertisement of services

- Dynamic discovery and invocation of various services in a heterogeneous environment
- Utilization of gateway services from simple devices

[021] A novel DRM engine and architecture is also presented that can be used with the NEMO framework. This DRM system can be used to achieve some or all of the following goals:

[022] *Simplicity.* In one embodiment, a DRM engine is provided that uses a minimalist stack-based Virtual Machine (VM) to execute control programs (e.g., programs that enforce governance policies). For example, the VM might consist of only a few pages of code.

[023] *Modularity.* In one embodiment, the DRM engine is designed to function as a single module integrated into a larger DRM-enabled application. Many of the functions that were once performed by monolithic DRM kernels (such as cryptography services) can be requested from the host environment, which may provide these services to other code modules. This allows designers to incorporate standard or proprietary technologies with relative ease.

[024] *Flexibility.* Because of its modular design, preferred embodiments of the DRM engine can be used in a wide variety of software environments, from embedded devices to general-purpose PCs.

[025] *Open.* Embodiments of the DRM engine are suitable for use as reference software, so that code modules and APIs can be implemented by users in virtually any programming language and in systems that they control completely. In one embodiment, the system does not force users to adopt particular content formats or restrict content encoding.

[026] *Semantically Agnostic.* In one embodiment, the DRM engine is based on a simple graph-based model that turns authorization requests into queries about the structure of the graph. The vertices in the graph represent entities in the system, and directed edges represent relationships between these entities. However, the DRM engine does not need to be aware of what these vertices and edges represent in any particular application.

[027] *Seamless Integration with Web Services.* The DRM client engine can use web services in several ways. For example, vertices and edges in the graph can be dynamically discovered through services. Content and content licenses may also be discovered and delivered to the DRM engine through sophisticated web services. Although one embodiment of the DRM engine can be configured to leverage web services in many places, its architecture is independent of web services, and can be used as a stand-alone client-side DRM kernel.

[028] *Simplified Key Management.* In one embodiment, the graph topology can be reused to simplify the derivation of content protection keys without requiring cryptographic retargeting. The key derivation method is an optional but powerful feature of the DRM engine—the system can also, or alternatively, be capable of integrating with other key management systems.

[029] *Separation of Governance, Encryption, and Content.* In one embodiment, the controls that govern content are logically distinct from the cryptographic information used to enforce the governance. Similarly, the controls and cryptographic information are logically distinct from content and content formats. Each of these elements can be delivered separately or in a unified package, thus allowing a high degree of flexibility in designing a content delivery system.

[029a] One aspect of the present invention provides a system for orchestrating services provided by network peers with sufficient interoperability to enable the exchange of value through participation in distributed applications, the system comprising: a. a first service provider having a first service adaptation layer that exposes a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider ; b. a first service consumer having a first service access point that accesses the first service interface exposed by the first service provider, discovers the first binding and invokes the first service by using that binding; and c. a first workflow collator that orchestrates the steps necessary to enable the first service provider to perform the first service for the first service consumer.

[029b] An alternate aspect provides a system for orchestrating services provided by network peers, the system comprising: a. a first service provider having a first service adaptation layer that exposes a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider; b. a first service consumer having a first service access point that accesses the first service interface exposed by the first service provider, discovers the first binding and invokes the first service by using that binding; c. a first workflow collator that orchestrates the steps necessary to enable the first service provider to perform the first service for the first service consumer; and d. a second service provider, having a second service adaptation layer that exposes a second service interface which enables network peers to access, via a second discoverable binding, a second service offered by the second service provider, the second service providing access to a service registry having a directory entry with information for locating and accessing the first service, wherein the first service access point accesses the director

[029c] A further aspect provides a system comprising: a. a first service provider having a first service adaptation layer that exposes a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider; b. a first service consumer having a first service access point that accesses the first service interface exposed by the first service provider, discovers the first binding and invokes the first service by using that binding; c. a first workflow collator that orchestrates the steps necessary to enable the first service provider to perform the first

service for the first service consumer; d., a trust management certificate for certifying the first service for access only by authorized service consumers; and e. a control program for providing limited access to the first service using the trust management certificate, wherein the first access point uses the trust management certificate to validate the first service consumer for authorized access to the first service.

[029d]Another aspect provides a method for orchestrating services provided by network peers with sufficient interoperability to enable the exchange of value through participation in distributed applications, the method comprising: a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider; b. accessing, from a first service access point of a service consumer, the first service interface exposed by the first service provider, discovering the first binding and invoking the first service by using that binding; and c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer.

[029e]Another alternate aspect provides a method for orchestrating services, the method comprising: a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding: a first service offered by the first service provider; b. accessing, from a first service access point of a service consumer, the first service interface exposed by the first service provider, discovering the first binding and invoking the first service by using that binding; c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer; and d. exposing, from a second adaptation layer of a second service provider, a second service interface which enables network peers to access, via a second discoverable binding, a second service offered by the second service provider, the second service providing access to a service registry having a directory entry with information for locating and accessing the first service, wherein the first service access point accesses the directory entry and uses the information to locate and invoke the first service.

[029f] Another alternate aspect provides a method comprising: a. exposing, from a first service adaptation layer of a first service provider, a first service interface which

enables network peers to access, via a first discoverable binding, a first service offered by the first service provider; b. accessing, from a first service access point of a service consumer, the first service interface exposed by the first service provider, discovering the first binding and invoking the first service by using that binding; c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer ; d. certifying the first service, using a trust management certificate, for access only by authorized service consumers; and e. executing a control program to provide limited access to the first service using the trust management certificate, wherein the first access point uses the trust management certificate to validate the first service consumer for authorized access to the first service.

[029g] Another alternate aspect provides a method for orchestrating services, the method comprising: a. accessing, from a first service access point of a service consumer, a first service interface exposed by a first service provider, the first service interface being operable to enable network peers to access, via a first discoverable binding, a first service offered by the first service provider; b. discovering the first discoverable binding; and c. invoking the first service by using that binding, wherein the steps necessary to perform the first service are orchestrated using a first workflow collator.

[029h]Another alternate aspect provides a method comprising: a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider; b. receiving a request for the first service from a first service access point of a first service consumer; and c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer.

[029i] Another alternate aspect provides a method comprising: a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the -first service provider; b. receiving a request for the first service from a first service access point of a first service consumer; c. certifying the first service, using a trust management certificate, for access only by authorized service consumers; d. executing a control program to provide limited access to the first service using the trust management

certificate; and e. orchestrating, using a first workflow collator, the steps necessary to, enable the first service provider to perform the first service for the first service consumer.

[030] Embodiments of the NEMO framework, its applications, and its component parts are described herein. It should be understood that the framework itself is novel, as are many of its components and applications. It should also be appreciated that the present inventions can be implemented in numerous ways, including as processes, apparatuses, systems, devices, methods, computer readable media, or a combination thereof. These and other features and advantages will be presented in more detail in the following detailed description and the accompanying drawings which illustrate by way of example the principles of the inventive body of work.

BRIEF DESCRIPTION OF THE DRAWINGS

[031] Embodiments of the inventive body of work will be readily understood by referring to the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

[032] Fig. 1 illustrates a sample embodiment of the system framework.

[033] Fig. 2a illustrates a conceptual network of system nodes.

[034] Fig. 2b illustrates system nodes in a P2P network.

[035] Fig. 2c illustrates system nodes operating across the Internet.

[036] Fig. 2d illustrates a system gateway node.

[037] Fig. 2e illustrates a system proxy node.

[038] Fig. 2f illustrates a system device adapter node.

[039] Fig. 3 illustrates a conceptual network of DRM devices.

[040] Fig. 4 illustrates a conceptual DRM node authorization graph.

[041] Fig. 5a illustrates a conceptual view of the architecture of a system node.

[042] Fig. 5b illustrates multiple service interface bindings supported by the service adaptation layer of a system node.

[043] Fig. 6a illustrates basic interaction between a service-providing system node and a service-consuming system node.

[044] Fig. 6b is another example of an interaction between a service-providing system node and a service-consuming system node.

[045] Fig. 7a illustrates a service access point involved in a client-side WSDL interaction.

[046] Fig. 7b illustrates a service access point involved in a client-side native interaction.

[047] Fig. 7c illustrates a service access point involved in a service-side point-to-point interaction pattern.

[048] Fig. 7d illustrates a service access point involved in a service-side point-to-multiple point interaction pattern.

[049] Fig. 7e illustrates a service access point involved in a service-side point-to-intermediary interaction pattern.

[050] Fig. 8 illustrates an embodiment of the architecture of the service adaptation layer.

[051] Fig. 9a illustrates an interaction pattern of a workflow collator relying upon external service providers.

[052] Fig. 9b illustrates an interaction pattern of a workflow collator involved in direct multi-phase communications with a client node.

[053] Fig. 9c illustrates a basic intra-node interaction pattern of a workflow collator.

[054] Fig. 9d illustrates a relatively complex interaction pattern of a workflow collator.

[055] Fig. 10 illustrates the system integration of a DRM engine.

[056] Fig. 11 illustrates an embodiment of the architecture of a DRM engine.

[057] Fig. 12a illustrates a DRM engine and related elements within a client-side system node.

[058] Fig. 12b illustrates a DRM engine and related elements within a service-side system node.

[059] Fig. 13 illustrates an embodiment of content protection and governance DRM objects.

[060] Fig. 14 illustrates an embodiment of node and link DRM objects.

[061] Fig. 15 illustrates an embodiment of DRM cryptographic key elements.

[062] Fig. 16 illustrates a basic interaction pattern between client and service-providing system nodes.

[063] Fig. 17a illustrates a set of notification processing nodes discovering a node that supports a notification handler service.

[064] Fig. 17b illustrates the process of notification delivery.

[065] Fig. 18a illustrates a client-driven service discovery scenario in which a requesting node makes a service discovery request to a targeted service providing node.

[066] Fig. 18b illustrates a peer registration service discovery scenario in which a requesting node seeks to register its description with a service providing node.

[067] Fig. 18c illustrates an event-based service discovery scenario in which an interested node receives a notification of a change in service availability (e.g., the existence of a service within a service-providing node).

[068] Fig. 19a illustrates the process of establishing trust using a service binding with an implicitly trusted channel.

[069] Fig. 19b illustrates the process of establishing trust based on a request/response model.

[070] Fig. 19c illustrates the process of establishing trust based on an explicit exchange of security credentials.

[071] Fig. 20 illustrates policy-managed access to a service.

[072] Fig. 21 illustrates a sample DRM node graph with membership and key access links.

[073] Fig. 22 illustrates an embodiment of the format of a DRM VM code module.

[074] Fig. 23 illustrates a system function profile hierarchy.

[075] Fig. 24 illustrates DRM music player application scenarios.

DETAILED DESCRIPTION

[076] A detailed description of the inventive body of work is provided below. While this description is provided in conjunction with several embodiments, it should be understood that the inventive body of work is not limited to any one embodiment, but instead encompasses numerous alternatives, modifications, and

equivalents. For example, while some embodiments are described in the context of consumer-oriented content and applications, those skilled in the art will recognize that the disclosed systems and methods are readily adaptable for broader application. For example, without limitation, these embodiments could be readily adapted and applied to the context of enterprise content and applications. In addition, while numerous specific details are set forth in the following description in order to provide a thorough understanding of the inventive body of work, some embodiments may be practiced without some or all of these details. Moreover, for the purpose of clarity, certain technical material that is known in the art has not been described in detail in order to avoid unnecessarily obscuring the inventive body of work.

1. CONCEPTS

1.1. Web Services

[077] The Web Services Architecture (WSA) is a specific instance of a Service Oriented Architecture (SOA). An SOA is itself a type of distributed system consisting of loosely coupled, cooperating software *agents*. The agents in an SOA may provide a service, request (consume) a service, or do both. A *service* can be seen as a well-defined, self-contained set of operations managed by an agent acting in a service provider role. The operations are invoked over the network at some network-addressable location, called an *endpoint*, using standard protocols and data formats. By self-contained, it is meant that the service does not depend directly on the state or context of another service or encompassing application.

[078] Examples of established technologies that support the concepts of an SOA include CORBA, DCOM, and J2EE. WSA is attractive because it is not tied to a specific platform, programming language, application protocol stack, or data format

convention. WSA uses standard formats based on XML for describing services and exchanging messages which promotes loose coupling and interoperability between providers and consumers, and supports multiple standard Internet protocols (notably HTTP), which facilitates deployment and participation in a potentially globally distributed system.

[079] An emerging trend is to view an SOA in the context of a "plug-and-play" service bus. The service bus approach provides for orchestration of services by leveraging description, messaging, and transport standards. The infrastructure may also incorporate standards for discovery, transformation, security, and perhaps others as well. Through the intrinsic qualities of the ubiquitous standards incorporated into the WSA, it is flexible, extensible, and scalable, and therefore provides the appropriate foundation for constructing an orchestrated service bus model. In this model, the fundamental unit of work (the service) is called a web service.

[080] There are a wide number of definitions for a web service. The following definition comes from the World Wide Web Consortium (W3C) Web Services Architecture working draft (August 8, 2003 – see www.w3.org/TR/ws-arch):

> *A Web service is a software system designed to support interoperable machine-to-machine interaction over a network. It has an interface described in a machine-processable format (specifically WSDL). Other systems interact with the Web service in a manner prescribed by its description using SOAP-messages, typically conveyed using HTTP with an XML serialization in conjunction with other Web-related standards.*

While the W3C definition provides a useful starting point, it should be understood that the term "web services" is used herein in a broader sense, without limitation, for

example, to the use of specific standards, formats, and protocols (e.g., WSDL, SOAP, XML, HTTP, etc.).

[081] A particular web service can be described as an abstract interface for a logically coherent set of operations that provides a basis for a (possibly transient) relationship between a service provider and a service requester.

[082] Of course, actual web services have concrete implementations. The provider's concrete implementation is sometimes referred to as the *service* (as distinguished from *web service)*. The software that actually implements the functionality for the service provider is the *provider agent* and for the service requester, the *requester agent*. The person or organization that owns the agent is referred to as the *provider entity* or *requester entity*, as appropriate. When used by itself, *requester* or *provider* may refer to either the respective entity or agent depending on context.

[083] A web service exists to fulfill a purpose, and how this is achieved is specified by the *mechanics* and *semantics* of the particular web service message exchange. The mechanics refers to the precise machine-processable technical specifications that allow the message exchange to occur over a network. While the mechanics are precisely defined, the semantics might not be. The semantics refers to the explicit or implicit "contract," in whatever form it exists, governing the understanding and overall expectations between the requester and provider entities for the web service.

[084] Web services are often modeled in terms of the interactions of three roles: (i) Service Provider; (ii) Service Requester; and (iii) Service Registry. In this model, a service provider "publishes" the information describing its web service to a

service registry. A service requester "finds" this information via some discovery mechanism, and then uses this information to "bind" to the service provider to utilize the service. *Binding* simply means that the requester will invoke the operations made available by the provider using the message formatting, data mapping, and transport protocol conventions specified by the provider in the published service description. The XML-based language used to describe this information is called Web Services Description Language (WSDL).

[085] A service provider offers access to some set of operations for a particular purpose described by a WSDL service description; this service description is published to a registry by any of a number of means so that the service may be discovered. A registry may be public or private within a specific domain.

[086] A service registry is software that responds to service search requests by returning a previously published service description. A service requester is software that invokes various operations offered by a provider according to the binding information specified in the WSDL obtained from a registry.

[087] The service registry may exist only conceptually or may in fact exist as real software providing a database of service descriptions used to query, locate, and bind to a particular service. But whether a requester actually conducts an active search for a service or whether a service description is statically or dynamically provided, the registry is a logically distinct aspect of the web services model. It is interesting to note that in a real world implementation, a service registry may be a part of the service requester platform, the service provider platform, or may reside at another location entirely identified by some well-known address or an address supplied by some other means.

[088] The WSDL service description supports loose coupling, often a central theme behind an SOA. While ultimately a service requester will understand the semantics of the interface of the service it is consuming for the purpose of achieving some desired result, the service description isolates a service interface from specific service binding information and supports a highly dynamic web services model.

[089] A service-oriented architecture can be built on top of many possible technology layers. As currently practiced, web services typically incorporate or involve aspects of the following technologies:

[090] *HTTP* – a standard application protocol for most web services communications. Although web services can be deployed over various network protocols (e.g., SMTP, FTP, etc), HTTP is the most ubiquitous, firewall-friendly transport in use. For certain applications, especially within an intranet, other network protocols may make sense depending on requirements; nevertheless, HTTP is a part of almost any web services platform built today.

[091] *XML* – a standard for formatting and accessing the content (and information about the content) of structured information. XML is a text-based standard for communicating information between web services agents. Note that the use of XML does not mean that message payloads for web services may not contain any binary data; but it does mean that this data will be formatted according to XML conventions. Most web services architectures do not necessarily dictate that messages and data be serialized to a character stream – they may just as likely be serialized to a binary stream where that makes sense – but if XML is being used, these streams will represent XML documents. That is, above the level of the transport mechanism, web service messaging will often be conducted using XML documents.

[092] Two XML subset technologies that are particularly important to many web services are XML Namespaces and XML Schema. XML-Namespaces are used to resolve naming conflicts and assert specific meanings to elements contained with XML documents. XML-Schema are used to define and constrain various information items contained within an XML document. Although it is possible (and optional) to accomplish these objectives by other means, the use of XML is probably the most common technique used today. The XML document format descriptions for web service documents themselves are defined using XML-Schema, and most real world web services operations and messages themselves will be further defined incorporating XML-Schema.

[093] *SOAP* – an XML-based standard for encapsulating instructions and information into a specially formatted package for transmission to and handling by other receivers. SOAP (Simple Object Access Protocol) is a standard mechanism for packaging web services messages for transmission between agents. Somewhat of a misnomer, its legacy is as a means of invoking distributed objects and in that respect it is indeed "simpler" than other alternatives; but the recent trend is to consider SOAP as an XML-based wire protocol for purposes that have transcended the original meaning of the acronym.

[094] SOAP defines a relatively lightweight convention for structuring messages and providing information about content. Each SOAP document contains an envelope that is divided into a header and a body. Although structurally similar, the header is generally used for meta-information or instructions for receivers related to the handling of the content contained in the body.

[095] SOAP also specifies a means of identifying features and the processing needed to fulfill the features' obligations. A *Message Exchange Pattern* (MEP) is a feature that defines a pattern for how messages are exchanged between nodes. A common MEP is request-response, which establishes a single, complete message transaction between a requesting and a responding node (see http://www.w3.org/TR/2003/REC-soap12-part2-20030624/#soapsupmep.).

[096] *WSDL* – an XML-based standard for describing how to use a web service. From a WSDL perspective, a service is related to a set of messages exchanged between service requesters and providers. Messages are described in an abstract manner that can be mapped to specific protocols. The exchange of messages that invokes some functionality is called an *operation*. A specific set of operations defines an *interface*. An interface is tied to a concrete message format and protocol by a named *binding*. The binding (the mapping of an interface to a concrete protocol) is associated with a URI appropriate to the protocol, resulting in an *endpoint*. A collection of one or more related endpoints (mapping an interface to concrete protocols at specific URIs) comprises a service.

[097] These definitions map to specific WSDL elements:

Types	**container element for type definitions**
Message	**an abstract definition of the type of data being sent**
Operation	**an abstract description of an action based on a combination of input, output , and fault messages**
portType	**an abstract set of operations – an *interface***
binding	**specification of a concrete protocol and data format for an interface (portType)**
port	**the combination of a binding and an actual network address – an *endpoint***

service	a collection of related ports (endpoints)

[098] WSDL defines a common binding mechanism and then defines specific binding extensions for SOAP, HTTP GET/POST, and MIME. Thus, binding does not necessarily mean binding to a transport protocol directly, but to a specific wire format. The most common binding for web services is SOAP, although actual SOAP message exchanges generally occur over HTTP on port 80 (via an http:// URI). However, an interface *can* be directly bound to HTTP; alternatively, for example, a binding for SOAP can use SMTP (via a mailto:// URI). An implementation can even define its own wire format and use a custom binding extension.

[099] WSDL encourages maintainability and reusability by providing support for an <import> element. Using import, a WSDL document can be divided into separate pieces in ways that make sense to an organization. For a cohesive web services environment desiring some degree of separation between an interface definition and an implementation definition, the following separation into three documents is reasonable:

A schema (.xsd) document – the root node is <schema> and the namespace is "http://www.w3.org/2001/XMLSchema."

A service interface description containing what is considered the *reusable* portion

<message>

<portType>

<binding>

A service implementation definition containing the specific service endpoint

<service>

[0100] WSDL interfaces are not exactly like Java (or IDL, or some other programming language) interfaces. For example, a Java interface declaration specifies a set of methods that must match at least a subset of the methods of a class

claiming to implement that interface. More than one class can implement an interface, and each implementation can be different; but the method signatures (method name and any input or output types) generally must be identical. This is mandated by the language and enforced at compile time, runtime, or both.

[0101] A WSDL interface is different, and more like an actual abstract class that alone is not fully useful. Various WSDL interfaces, or portTypes, of a single web service are logically related in the sense that the set of operation names should be identical – as if the portType *did*, in fact, implement a specific contract defined somewhere else – but no such element actually exists and there is no mechanism for enforcing portType symmetry. Each portType is generally named to identify the type of binding it supports – even though a portType alone does not create a binding. The portType operations for related portTypes are named the same, but the input, output, and fault messages (if present) are mapped to specific messages that contain named parts also necessary for supporting a specific binding. This raises the point that messages themselves are not completely abstract. A web service may and often does need to define similar but distinct messages for the various bindings required.

[0102] As will be illustrated below, by leveraging emerging web service and related standards, a system architecture can be developed that facilitates the creation of networked interoperable media-related services that utilize a variety of different protocols and interfaces across a wide range of hardware and software platforms and operating environments.

1.2. Roles

[0103] Preferred embodiments of the present invention seek to enable, promote, and/or actively support a peer-to-peer environment in which peers can

spontaneously offer a variety of functionality by exposing services. One embodiment of the framework discourages viewing peers as having a fixed set of capabilities; and instead encourages a model where a peer at any point in time is a participant in one or more roles.

[0104] A role can be defined by a set of services that a given peer exposes in combination with a specific behavior pattern. At any given moment a NEMO-enabled node may act in multiple roles based on a variety of factors: its actual implementation footprint providing the functionality for supporting a given set of services, administrative configuration, information declaring the service(s) the peer is capable of exposing, and load and runtime policy on service interfaces.

[0105] An explicit set of roles could be defined based on various different types of services. Over time, as common patterns of participation are determined and as new services are introduced, a more formal role categorization scheme could be defined. A preliminary set of roles that may be formalized over time could include the following:

[0106] *Client* – a relatively simple role in which no services are exposed, and the peer simply uses services of other peers.

[0107] *Authorizer* - this role denotes a peer acting as a Policy Decision Point (PDP), determining if a requesting principal has access to a specified resource with a given set of pre-conditions and post-conditions.

[0108] *Gateway* - in certain situations a peer may not be able to directly discover or interact with other service providers, for reasons including: transport protocol incompatibility, inability to negotiate a trusted context, or lack of the processing capability to create and process the necessary messages associated with a

given service. A gateway is a peer acting as a bridge to another peer in order to allow the peer to interact with a service provider. From the perspective of identity and establishing an authorized and trusted context for operation, the requesting peer may actually delegate to the gateway peer its identity and allow that peer to negotiate and make decisions on its behalf. Alternatively, the gateway peer may act as a simple relay point, forwarding or routing requests and responses.

[0109] *Orchestrator* - in situations where interaction with a set of service providers involves nontrivial coordination of services (possibly including transactions, distributed state management, etc.), it may be beyond a peer's capability to participate. An orchestrator is a specialization of the gateway role. A peer may request an orchestrator to act on its behalf, intervening to provide one or more services. The orchestrating peer may use certain additional NEMO components, such as an appropriately configured Workflow Collator in order to satisfy the orchestration requirements.

[0110] Given the goal of "providing instant gratification by satisfying a request for any media, in any format, from any source, at any place, at any time, on any device complying with any agreeable set of usage rules," the following informal model illustrates how this goal can be achieved using embodiments of the NEMO framework. It will become apparent from the highest level of the model (without enumerating every aspect of how NEMO enables all of the media services that one can imagine) how NEMO enables lower-level services from different tiers in the model to be assembled into richer end-to-end media services.

[0111] In one embodiment of this model there are four tiers of service components: 1) Content Authoring, Assembly, and Packaging services, 2) Web-based

Content Aggregation and Distribution services, 3) Home Gateway services, and 4) Consumer Electronics devices.

[0112] Each of these four tiers typically has different requirements for security, rights management, service discovery, service orchestration, user interface complexity, and other service attributes. The first two tiers fit very roughly into the models that we see for "traditional" web services, while the last two tiers fit more into what we might call a personal logical network model, with certain services of the home gateway being at the nexus between the two types of models. However, services for CE devices could occasionally appear in any of the tiers.

[0113] One dilemma lies in the desire to specialize parts of the framework for efficiency of implementation, while being general enough to encompass an end-to-end solution. For example, a UDDI directory and discovery approach may work well for relatively static and centralized web services, but for a more dynamic transient merging of personal networks, discovery models such as those found in UPnP and Rendezvous may be more appropriate. Thus, in some embodiments multiple discovery standards are accommodated within the framework

[0114] Similarly, when rights management is applied to media distribution through wholesale, aggregator, and retail distribution sub-tiers, there can be many different types of complex rights and obligations that need to be expressed and tracked, suggesting the need for a highly expressive and complex rights language, sophisticated content governance and clearing services, and a global trust model. However, rights management and content governance for the home gateway and CE device tiers may entail a different trust model that emphasizes fair use rights that are relatively straightforward from the consumer's point of view. Peer devices in a

personal logical network may want to interact using the relatively simple trust model of that network, and with the ability to interact with peers across a wide area network using a global trust model, perhaps through proxy gateway services. At the consumer end, complexity arises from automated management of content availability across devices, some of which are mobile and intermittently intersect multiple networks. Thus, an effective approach to rights management, while enabling end-to-end distribution, might also be heterogeneous, supporting a variety of rights management services, including services that interpret expressions of distribution rights and translate them, in context, to individual consumer use rights in a transaction that is orchestrated with a sales transaction, or perhaps another event where a subscription right is exercised.

1.3. Logical Model

[0115] In one embodiment, the system framework consists of a logically connected set of nodes that interact in a peer-to-peer (P2P) fashion. Peer-to-peer computing is often defined as the sharing of resources (such as hard drives and processing cycles) among computers and other intelligent devices. *See* http://www.intel.com/cure/peer.htm. Here, P2P may be viewed as a communication model allowing network nodes to symmetrically consume and provide services of all sorts. P2P messaging and workflow collation allow rich services to be dynamically created from a heterogeneous set of more primitive services. This enables examination of the possibilities of P2P computing when the shared resources are services of many different types, even using different service bindings.

[0116] Different embodiments can provide a media services framework enabling stakeholders (e.g., consumers, content providers, device manufacturers, and

service providers) to find one another, to interact, exchange value, and to cooperate in rich and dynamic ways. These different types of services range from the basic (discovery, notification, search, and file sharing) to more complex higher level services (such as lockers, licensing, matching, authorization, payment transaction, and update), and combinations of any or all of these.

[0117] Services can be distributed across peer-to-peer communicating nodes, each providing message routing and orchestration using a message pump and workflow collator (described in greater detail below) designed for this framework.

[0118] Nodes interact by making service invocation requests and receiving responses. The format and payload of the request and response messages are preferably defined in a standard XML schema-based web service description language (e.g., WSDL) that embodies an extensible set of data types enabling the description and composition of services and their associated interface bindings. Many of the object types in WSDL are polymorphic and can be extended to support new functionality. The system framework supports the construction of diverse communication patterns, ranging from direct interaction with a single service provider to a complex aggregation of a choreographed set of services from multiple service providers. In one embodiment, the framework supports the basic mechanisms for using existing service choreography standards (WSCI, BPEL, etc.), and also allows service providers to use their own conventions.

[0119] The syntax of messages associated with service invocation are preferably described in a relatively flexible and portable manner, as are the core data types used within the system framework. In one embodiment, this is accomplished

using WSDL to provide relatively simple ways for referencing semantic descriptions associated with described services.

[0120] A service interface may have one or more service bindings. In such an embodiment, a node may invoke the interface of another node as long as that node's interface binding can be expressed in, e.g., WSDL, and as long as the requesting node can support the conventions and protocols associated with the binding. For example, if a node supports a web service interface, a requesting node may be required to support SOAP, HTTP, WS-Security, etc.

[0121] Any service interface may be controlled (e.g., rights managed) in a standardized fashion directly providing aspects of rights management. Interactions between nodes can be viewed as governed operations.

[0122] Virtually any type of device (physical or virtual) can be viewed as potentially NEMO-enabled, and able to implement key aspects of the NEMO framework. Device types include, for example, consumer electronics equipment, networked services, and software clients. In a preferred embodiment, a NEMO-enabled device (node) typically includes some or all of the following logical modules (discussed in greater detail below):

[0123] *Native Services API* – the set of one or more services that the device implements. There is no requirement that a NEMO node expose any service directly or indirectly in the NEMO framework.

[0124] *Native Service Implementation* – the corresponding set of implementations for the native services API.

[0125] *Service Adaptation Layer* – the logical layer through which an exposed subset of an entity's native services is accessed using one or more discoverable bindings described in, e.g., WSDL.

[0126] *Framework Support Library* – components that provide support functionality for working with the NEMO Framework including support for invoking service interfaces, message processing, service orchestration, etc.

1.4. Terminology

[0127] In one embodiment, a basic WSDL profile defines a minimum "core" set of data types and messages for supporting interaction patterns and infrastructural functionality. Users may either directly, in an ad-hoc manner, or through some form of standardization process, define other profiles built on top of this core, adding new data and service types and extending existing ones. In one embodiment, this core profile includes definitions for some or all of the following major basic data types:

[0128] *Node* – a representation of a participant in the system framework. A node may act in multiple roles including that of a service consumer and/or a service provider. Nodes may be implemented in a variety of forms including consumer electronic devices, software agents such as media players, or virtual service providers such as content search engines, DRM license providers, or content lockers.

[0129] *Device* – encapsulates the representation of a virtual or physical device.

[0130] *User* – encapsulates the representation of a client user.

[0131] *Request* – encapsulates a request for a service to a set of targeted Nodes.

[0132] *Request Input* – encapsulates the input for a Request.

[0133] *Response* – encapsulates a Response associated with a Request.

[0134] *Request Result* – encapsulates the Results within a Response associated with some Request.

[0135] *Service* – encapsulates the representation of a set of well-defined functionality exposed or offered by a provider Node. This could be, for example, low-level functionality offered within a device such as a cell phone (e.g. a voice recognition service), or multi-faceted functionality offered over the world-wide web (e.g. a shopping service). Services could cover a wide variety of applications, including DRM-related services such as client personalization and license acquisition.

[0136] *Service Provider* – an entity (e.g., a Node or Device) that exposes some set of Services. Potential Service Providers include consumer electronics devices, such as cell phones, PDAs, portable media players and home gateways, as well as network operators (such as cable head-ends), cellular network providers, web-based retailers and content license providers.

[0137] *Service Interface* – a well-defined way of interacting with one or more Services.

[0138] *Service Binding* – encapsulates a specific way to communicate with a Service, including the conventions and protocols used to invoke a Service Interface. These may be represented in a variety of well-defined ways, such as the WS-I standard XML protocol, RPC based on the WSDL definition, or a function invocation from a DLL.

[0139] *Service Access Point (SAP)* – encapsulates the functionality necessary for allowing a Node to make a Service Invocation Request to a targeted set of Service Providing Nodes, and receive a set of Responses.

[0140] *Workflow Collator (WFC)* – a Service Orchestration mechanism that provides a common interface allowing a Node to manage and process collections of Requests and Responses related to Service invocations. This interface provides the basic building blocks to orchestrate Services through management of the Messages associated with the Services.

[0141] In the context of a particular application, such as digital rights management (DRM), a typical profile might include various DRM-related services (described below) for the following set of content protection and governance objects, which represent entities in the system, protect content, associate usage rules with the content, and determine if access can be granted when requested:

[0142] *Content Reference* – encapsulates the representation of a reference or pointer to a content item. Such a reference will typically leverage other standardized ways of describing content format, location, etc.

[0143] *DRM Reference* – encapsulates the representation of a reference or pointer to a description of a digital rights management format.

[0144] *Link* – links between entities (e.g., Nodes).

[0145] *Content* – represents media or other content.

[0146] *Content Key* – represents encryption keys used to encrypt Content.

[0147] *Control* – represents usage or other rules that govern interaction with Content.

[0148] *Controller* – represent associations between Control and ContentKey objects

[0149] *Projector* – represent associations between Content and ContentKey objects

[0150] In one embodiment, a core profile includes definitions for some or all of the following basic Services:

[0151] *Authorization* – a request or response to authorize some participant to access a Service.

[0152] *Governance* – The process of exercising authoritative or dominating influence over some item (e.g., a music file, a document, or a Service operation), such as the ability to download and install a software upgrade. Governance typically interacts with Services providing functionality such as trust management, policy management, and content protection.

[0153] *Message Routing* – a Request or Response to provide message routing functionality, including the ability to have the Service Providing Node forward the message or collect and assemble messages.

[0154] *Node Registration* – a Request or Response to perform registration operations for a Node, thereby allowing the Node to be discovered through an Intermediate Node.

[0155] *Node Discovery (Query)* – a Request or Response related to the discovery of Nodes.

[0156] *Notification* – a Request or Response to send or deliver targeted Notification messages to a given set of Nodes.

[0157] *Security Credential Exchange* – a Request or Response related to allowing Nodes to exchange security related information, such as key pairs, certificates, or the like.

[0158] *Service Discovery (Query)* – a Request or Response related to the discovery of Services provided by some set of one or more Nodes.

[0159] *Service Orchestration* – The assembly and coordination of Services into manageable, coarser-grained Services, reusable components, or full applications that adhere to rules specified by a service provider. Examples include rules based on provider identity, type of Service, method by which Services are accessed, order in which Services are composed, etc.

[0160] *Trust Management* – provides a common set of conventions and protocols for creating authorized and trusted contexts for interactions between Nodes. In some embodiments, NEMO Trust Management may leverage and/or extend existing security specifications and mechanisms, including WS-Security and WS-Policy in the web services domain.

[0161] *Upgrade* – represents a Request or Response related to receiving a functionality upgrade. In one embodiment, this service is purely abstract, with other profiles providing concrete representations.

1.5. Illustrative Interaction Between Nodes

[0162] As will be discussed in greater detail below, the basic logical interaction between two system nodes, a service requester and a service provider, typically includes the following sequence of events. From the perspective of the service requesting node:

[0163] The service requesting node makes a service discovery request to locate any NEMO-enabled nodes that can provide the necessary service using the specified service bindings. A node may choose to cache information about discovered services. The interface/mechanism for service discovery between nodes can be just another service that a NEMO node chooses to implement.

[0164] Once candidate service providing nodes are found, the requesting node may choose to dispatch a request to one or more of the service providing nodes based on a specific service binding.

[0165] In one embodiment, two nodes that wish to communicate securely with each other will establish a trusted relationship for the purpose of exchanging WSDL messages. For example, they may negotiate a set of compatible trust credentials (e.g., X.500 certificates, device keys, etc.) that may be used in determining identity, verifying authorization, establishing a secure channel, etc. In some cases, the negotiation of these credentials may be an implicit property of the service interface binding (e.g., WS-Security if WS-I XML Protocol is used, or an SSL request between two well-known nodes). In other cases, the negotiation of trust credentials may be an explicitly separate step. In one embodiment, it is up to a given node to determine which credentials are sufficient for interacting with another node, and to make the decision that it can trust a given node.

[0166] The requesting node creates the appropriate WSDL request message(s) that correspond to the requested service.

[0167] Once the messages are created, they are dispatched to the targeted service providing node(s). The communication style of the request may, for example, be synchronous or asynchronous RPC style, or message-oriented based on the service binding. Dispatching of service requests and receiving of responses may be done directly by the device or through the NEMO Service Proxy. The service proxy (described below) provides an abstraction and interface for sending messages to other participants, and may hide certain service binding issues, such as compatible message formats, transport mechanisms, message routing issues, etc.

[0168] After dispatching a request, the requesting node will typically receive one or more responses. Depending on the specifics of the service interface binding and the requesting node's preferences, the response(s) may be returned in a variety of ways, including, for example, an RPC-style response or a notification message. The response, en-route to the targeted node(s), may pass through other intermediate nodes that may provide a number of relevant services, including, e.g., routing, trust negotiation, collation and correlation functions, etc.

[0169] The requesting node validates the response(s) to ensure it adheres to the negotiated trust semantics between it and the service providing node.

[0170] Appropriate processing is then applied based on the message payload type and contents.

[0171] From the perspective of the service providing node, the sequence of events typically would include the following:

[0172] Determine if the requested service is supported. In one embodiment, the NEMO framework does not mandate the style or granularity of how a service interface maps as an entry point to a service. In the simplest case, a service interface may map unambiguously to a given service and the act of binding to and invoking it may constitute support for the service. However, in some embodiments a single service interface may handle multiple types of requests; and a given service type may contain additional attributes that need to be examined before a determination can be made that the node supports the specifically desired functionality.

[0173] In some cases it may be necessary for the service provider to determine if it trusts the requesting node and to negotiate a set of compatible trust credentials. In

one embodiment, regardless of whether the service provider determines trust, any policy associated with the service interface will still apply.

[0174] The service provider determines and dispatches authorization request(s) to those node(s) responsible for authorizing access to the interface in order to determine if the requesting node has access. In many situations, the authorizing node and the service providing node will be the same entity, and the dispatching and processing of the authorization request will be local operations invoked through a lightweight service interface binding such as a C function entry point.

[0175] Upon receiving the authorization response, if the requesting node is authorized, the service provider will fulfill the request. If not, an appropriate response message might be generated.

[0176] The response message is returned based on the service interface binding and requesting node's preferences. En route to the requesting node, the message may pass through other intermediate nodes that may provide necessary or "value added" services. For example an intermediate node might provide routing, trust negotiation, or delivery to a notification processing node that can deliver the message in a way acceptable to the requesting node. An example of a "value added" service is a coupon service that appends coupons to the message if it knows of the requesting node's interests.

2. SYSTEM ARCHITECTURE

[0177] Consider a sample embodiment of the NEMO system framework, as illustrated in **FIG. 1**, implementing a DRM application.

[0178] As noted above, NEMO nodes may interact by making service invocation requests and receiving responses. The NEMO framework supports the

construction of diverse and rich communication patterns ranging from a simple point to point interaction with a single service provider to a complex aggregation of a choreographed set of services from multiple service providers.

[0179] In the context of **FIG. 1**, the NEMO nodes interact with one another to provide a variety of services that, in the aggregate, implement a music licensing system. Music stored in Consumer Music Locker 110 can be extracted by Web Music Retailer 120 and provided to end users at their homes via their Entertainment Home Gateway 130. Music from Consumer Music Locker 110 may include rules that govern the conditions under which such music may be provided to Web Music Retailer 120, and subsequently to others for further use and distribution. Entertainment Home Gateway 130 is the vehicle by which such music (as well as video and other content) can be played, for example, on a user's home PC (e.g., via PC Software Video Player 140) or on a user's portable playback device (e.g., Portable Music Player 150). A user might travel, for example, with Portable Music Player 150 and obtain, via a wireless Internet connection (e.g., to Digital Rights License Service 160), a license to purchase additional songs or replay existing songs additional times, or even add new features to Portable Music Player 150 via Software Upgrade Service 170.

[0180] NEMO nodes can interact with one another, and with other devices, in a variety of different ways. A NEMO host, as illustrated in **FIG. 2a**, is some type of machine or device hosting at least one NEMO node. A host may reside within a personal area network 210 or at a remote location 220 accessible via the Internet. A host could, for example, be a server 230, a desktop PC 240, a laptop 250, or a personal digital assistant 260.

[0181] A NEMO node is a software agent that can provide services to other nodes (such as host 235 providing a 3rd party web service) as well as invoke other nodes' services within the NEMO-managed framework. Some nodes 270 are tethered to another host via a dedicated communication channel, such as Bluetooth. These hosts 240 and 250 are equipped with network connectivity and sufficient processing power to present a virtual node to other participating NEMO nodes.

[0182] As illustrated in **FIG. 2b**, a NEMO node can be a full peer within the local or personal area network 210. Nodes share the symmetric capability of exposing and invoking services; however, each node generally does not offer identical sets of services. Nodes may advertise and/or be specifically queried about the services they perform.

[0183] If an Internet connection is present, as shown in **FIG. 2c**, then local NEMO nodes (e.g., within personal area network 210) can also access the services of remote nodes 220. Depending on local network configuration and policy, it is also possible for local and remote nodes (e.g., Internet-capable NEMO hosts 280) to interoperate as NEMO peers.

[0184] As illustrated in **FIG. 2d**, not all NEMO nodes may be on hosts capable of communicating with other hosts, whether local or remote. A NEMO host 280 can provide a gateway service through which one node can invoke the services of another, such as tethered node 285 or nodes in personal area network 210.

[0185] As illustrated in **FIG. 2e**, a node 295 on a tethered device may access the services of other nodes via a gateway, as discussed above. It may also be accessed by other nodes via a proxy service on another host 290. The proxy service creates a virtual node running on the NEMO host. These proxy nodes can be full NEMO peers.

[0186] As illustrated in **FIG. 2f**, a NEMO host may provide dedicated support for tethered devices via NEMO node adapters. A private communication channel 296 is used between host/NEMO device adapter 297 and tethered node 298 using any suitable protocol. Tethered node 298 does not see, nor is it visible to, other NEMO peer nodes.

[0187] We next consider exemplary digital rights management (DRM) functionality that can be provided by NEMO-enabled devices in certain embodiments, or that can be used outside the NEMO context. As previously described, one of the primary goals of a preferred embodiment of the NEMO system framework is to support the development of secure, interoperable interconnections between media-related services spanning both commercial and consumer-oriented network tiers. In addition to service connectivity, interoperability between media-related services will often require coordinated management of usage rights as applied to the content available through those services. NEMO services and the exemplary DRM engine described herein can be used in combination to achieve interoperability that allows devices based on the NEMO framework to provide consumers with the perception of a seamless rendering and usage experience, even in the face of a heterogeneous DRM and media format infrastructure.

[0188] In the context of a DRM application, as illustrated in **FIG. 3**, a network of NEMO-enabled DRM devices may include content provider/server 310, which packages content for other DRM devices, as well as consumer PC player 330 and consumer PC packager/player 320, which can not only play protected content, but can also package content for delivery to portable device 340.

[0189] Within each DRM device, the DRM engine performs specific DRM functions (e.g., enforcing license terms, delivering keys to the host application, etc.), and relies on the host application for those services which can be most effectively provided by the host, such as encryption, decryption, and file management.

[0190] As will be discussed in greater detail below, in one embodiment the DRM engine includes a virtual machine (VM) designed to determine whether certain actions on protected content are permissible. This Control VM can be implemented as a simple stack-based machine with a minimal set of instructions. In one embodiment, it is capable of performing logical and arithmetic calculations, as well as querying state information from the host environment to check parameters such as system time, counter state, and so forth.

[0191] In one embodiment, the DRM engine utilizes a graph-based algorithm to verify relationships between entities in a DRM value chain. **FIG. 4** illustrates a conceptual embodiment of such a graph. The graph comprises a collection of nodes or vertices, connected by links. Each entity in the system can be represented by a vertex object. Only entities that need to be referenced by link objects, or be the recipient of cryptographically targeted information, need to have corresponding vertex objects. In one embodiment, a vertex typically represents a user, a device, or a group. Vertex objects also have associated attributes that represent certain properties of the entity associated with the vertex.

[0192] For example, **FIG. 4** shows two users (Xan and Knox), two devices (the Mac and a portable device), and several entities representing groups (members of the Carey family, members of the public library, subscribers to a particular music

service, RIAA-approved devices, and devices manufactured by a specific company). Each of these has a vertex object associated with it.

[0193] The semantics of the links may vary in an application-specific manner. For example, the directed edge from the Mac vertex to the Knox vertex may mean that Knox is the owner of the Mac. The edge from Knox to Public Library may indicate that Knox is a member of the Public Library. In one embodiment the DRM engine does not impose or interpret these semantics—it simply ascertains the existence or non-existence of paths within the graph. This graph of vertices can be considered an "authorization" graph in that the existence of a path or relationship (direct or indirect) between two vertices may be interpreted as an authorization for one vertex to access another vertex.

[0194] For example, because Knox is linked to the Carey family and the Carey family is linked to the Music Service, there is a path between Knox and the Music Service. The Music Service vertex is considered reachable from another vertex when there is a path from that vertex to the Music Service. This allows a control to be written that allows permission to access protected content based on the condition that the Music Service be reachable from the portable device in which the application that requests access (e.g., a DRM client host application) is executing.

[0195] For example, a content owner may create a control program to be interpreted by the Control VM that allows a particular piece of music to be played if the consuming device is owned by a member of the Public Library and is RIAA-approved. When the Control VM running on the device evaluates this control program, the DRM engine determines whether links exist between Portable Device and Public Library, and between Portable Device and RIAA Approved. The edges

and vertices of the graph may be static and built into devices, or may be dynamic and discovered through services communicating with the host application.

[0196] By not imposing semantics on the vertices and links, the DRM engine can enable great flexibility. The system can be adapted to many usage models, from traditional delegation-based policy systems to authorized domains and personal area networks.

[0197] In one embodiment, the DRM client can also reuse the authorization graph for content protection key derivation. System designers may chose to allow the existence of a link to also indicate the sharing of certain cryptographic information. In such cases, the authorization graph can be used to derive content keys without explicit cryptographic retargeting to consuming devices.

3. NODE ARCHITECTURE

3.1. Overview

[0198] Any type of device (physical or virtual), including consumer electronics equipment, networked services, or software clients, can potentially be NEMO-enabled, which means that the device's functionality may be extended in such a way as to enable participation in the NEMO system. In one embodiment, a NEMO-enabled device (node) is conceptually comprised of certain standard modules, as illustrated in **FIG. 5a**.

[0199] Native Services API 510 represents the logical set of one or more services that the device implements. There is no requirement that a NEMO node expose any service directly or indirectly. Native Service Implementation 520 represents the corresponding set of implementations for the native services API.

[0200] Service Access Point 530 provides support for invoking exposed service interfaces. It encapsulates the functionality necessary for allowing a NEMO node to make a service invocation request to a targeted set of service-providing NEMO nodes and to receive a set of responses. NEMO-enabled nodes may use diverse discovery, name resolution, and transport protocols, necessitating the creation of a flexible and extensible communication API. The Service Access Point can be realized in a variety of ways tailored to a particular execution environment and application framework style. One common generic model for its interface will be an interface capable of receiving XML messages in some form and returning XML messages. Other models with more native interfaces can also be supported.

[0201] NEMO Service Adaptation Layer 540 represents an optional layer through which an exposed subset of an entity's native services are accessed using one or more discoverable bindings. It provides a level of abstraction above the native services API, enabling a service provider to more easily support multiple types of service interface bindings. In situations where a service adaptation layer is not present, it may still be possible to interact with the service directly through the Service Access Point 530 if it supports the necessary communication protocols.

[0202] The Service Adaptation Layer 540 provides a common way for service providers to expose services, process requests and responses, and orchestrate services in the NEMO framework. It is the logical point at which services are published, and provides a foundation on which to implement other specific service interface bindings.

[0203] In addition to providing a common way of exposing a service provider's native services to other NEMO-enabled nodes, Service Adaptation Layer

540 also provides a natural place on which to layer components for supporting additional service interface bindings 560, as illustrated in **FIG. 5b**. By supporting additional service interface bindings, a service provider increases the likelihood that a compatible binding will be able to be negotiated and used either by a Service Access Point or through some other native API.

[0204] Referring back to **FIG. 5a**, Workflow Collator 550 provides supporting management of service messages and service orchestration. It provides a common interface allowing a node to manage and process collections of request and response messages. This interface in turn provides the basic building blocks to orchestrate services through management of the messages associated with those services. This interface typically is implemented by a node that supports message routing functionality as well as the intermediate queuing and collating of messages.

[0205] In some embodiments, the NEMO framework includes a collection of optional support services that facilitate an entity's participation in the network. Such services can be classified according to various types of functionality, as well as the types of entities requiring such services (e.g., services supporting client applications, as opposed to those needed by service providers). Typical supporting services include the following:

[0206] *WSDL Formatting and Manipulation Routines* – provide functionality for the creation and manipulation of WSDL-based service messages.

[0207] *Service Cache* – provides a common interface allowing a node to manage a collection of mappings between discovered nodes and the services they support.

[0208] *Notification Processor Interface* – provides a common service provider interface for extending a NEMO node that supports notification processing to some well-defined notification processing engine.

[0209] *Miscellaneous Support Functionality* – including routines for generating message IDs, timestamps, etc.

3.2. Basic Node Interaction

[0210] Before examining the individual architectural elements of NEMO nodes in greater detail, it is helpful to understand the manner by which such nodes interact and communicate with one another. Diverse communication styles are supported, ranging from synchronous and asynchronous RPC-style communication, to one-way interface invocations and client callbacks.

[0211] *Asynchronous RPC Delivery Style* – this model is particularly appropriate if there is an expectation that fulfilling the request will take an extended period of time and the client does not want to wait. The client submits a request with the expectation that it will be processed in an asynchronous manner by any service-providing nodes. In this case, the service-providing endpoint may respond indicating that it does not support this model, or, if the service-providing node does support this model, it will return a response that will carry a ticket that can be submitted to the given service-providing node in subsequent requests to determine if it has a response to the client's request.

[0212] In one embodiment, any service-providing endpoint that does support this model is obligated to cache responses to pending client requests based on an internal policy. If a client attempts to redeem a ticket associated with such a request and no response is available, or the response has been thrown away by the service-

providing node, then an appropriate error response is returned. In this embodiment, it is up to the client to determine when it will make such follow-on requests in attempting to redeem the ticket for responses.

[0213] *Synchronous RPC Delivery Style* – the client submits a request and then waits for one or more responses to be returned. A service-providing NEMO-enabled endpoint may respond indicating that it does not support this model.

[0214] *Message-Based Delivery Style* – the client submits a request indicating that it wants to receive any responses via a message notification associated with one or more of its notification handling service interfaces. A service-providing NEMO-enabled endpoint may respond indicating that it does not support this model.

[0215] From the client application's perspective, none of the interaction patterns above necessitates an architecture that must block and wait for responses, or must explicitly poll. It is possible to use threading or other platform-specific mechanisms to model both blocking and non-blocking semantics with the above delivery style mechanisms. Also, none of the above styles is intended to directly address issues associated with the latency of a given communication channel — only potential latency associated with the actual fulfillment of a request. Mechanisms to deal with the issues associated with communication channel latency should be addressed in the specific implementation of a component such as the Service Access Point, or within the client's implementation directly.

3.3. Service Access Point

[0216] As noted above, a Service Access Point (SAP) can be used as a common, reusable API for service invocation. It can encapsulate the negotiation and use of a transport channel. For example, some transport channels may require SSL

session setup over TCP/IP, while some channels may only support relatively unreliable communication over UDP/IP, and still others may not be IP-based at all.

[0217] A SAP can encapsulate the discovery of an initial set of NEMO nodes for message routing. For example, a cable set-top box may have a dedicated connection to the network and mandate that all messages flow through a specific route and intermediary. A portable media player in a home network may use UPnP discovery to find multiple nodes that are directly accessible. Clients may not be able, or may choose not, to converse directly with other NEMO nodes by exchanging XML messages. In this case, a version of the SAP may be used that exposes and uses whatever native interface is supported.

[0218] In a preferred embodiment, the SAP pattern supports the following two common communication models (although combinations of the two, as well as others, may be supported): (i) *Message Based* (as discussed above) – where the SAP forms XML request messages and directly exchanges NEMO messages with the service provider via some interface binding; or (ii) *Native* – where the SAP may interact with the service provider through some native communication protocol. The SAP may internally translate to/from XML messages defined elsewhere within the framework.

[0219] A sample interaction between two NEMO peer nodes is illustrated in **FIG. 6a**. Client node 610 interacts with service-providing node 660 using NEMO service access point (SAP) 620. In this example, web service protocols and standards are used both for exposing services and for transport. Service-providing node 660 uses its web services layer 670 (using, e.g., WSDL and SOAP-based messaging) to expose its services to clients such as node 610. Web services layer 630 of client node 610 creates and interprets SOAP messages, with help from mapping layer 640 (which

maps SOAP messages to and from SAP interface 620) and trust management processing layer 650 (which could, for example, leverage WS-Security using credentials conveyed within SOAP headers).

[0220] Another example interaction between NEMO nodes is illustrated in **FIG. 6b**. Service-providing node 682 interacts with client node 684 using SAP 686. In this example, service-providing node 682 includes a different but interoperable trust management layer than client 684. In particular, service-providing node 682 includes both a trust engine 688 and an authorization engine 690. In this example, trust engine 688 might be generally responsible for performing encryption and decryption of SOAP messages, for verifying digital certificates, and for performing other basic cryptographic operations, while authorization engine 690 might be responsible for making higher-level policy decisions. In the example shown in **FIG. 6b**, client node 684 includes a trust engine 692, but not an authorization engine. Thus, in this example, client node 684 might be capable of performing basic cryptographic operations and enforcing relatively simple policies (e.g., policies related to the level of message authenticity, confidentiality, or the like), but might rely on service providing node 682 to evaluate and enforce higher order policies governing the client's use of, and interaction with, the services and/or content provided by service providing node 682. It should be appreciated that **FIG. 6b** is provided for purposes of illustration and not limitation, and that in other embodiments client node 684 might also include an authorization engine, as might be the case if the client needed to adhere to a set of obligations related to a specified policy. Thus, it can be seen that different NEMO peers can contain different parts of the trust management framework depending on their requirements. **FIG. 6b** also illustrates that the

communication link between nodes can be transport agnostic. Even in the context of a SOAP processing model, any suitable encoding of data and/or processing rules can be used. For example, the XML security model could be replaced with another security model that supported a different encoding scheme.

[0221] A Service Access Point may be implemented in a variety of forms, such as within the boundaries of a client (in the form of a shared library) or outside the boundaries of the client (in the form of an agent running in a different process). The exact form of the Service Access Point implementation can be tailored to the needs of a specific type of platform or client. From a client's perspective, use of the Service Access Point may be optional, although in general it provides significant utility, as illustrated below.

[0222] The Service Access Point may be implemented as a static component supporting only a fixed set of service protocol bindings, or it may be able to support new bindings dynamically.

[0223] Interactions involving the Service Access Point can be characterized from at least two perspectives – a client-side which the requesting participant uses, and a service-side which interacts with other NEMO-enabled endpoints (nodes).

[0224] In one client-side embodiment, illustrated in **FIG. 7a**, Service Access Point 710 directly exchanges XML messages with client 720. Client 720 forms request messages 740 directly and submits them to Service Access Point 710, which generates and sends one or more response messages 750 to client 720, where they are collected, parsed and processed. Client 720 may also submit (when making requests) explicit set(s) of service bindings 730 to use in targeting the delivery of the request. These service bindings may have been obtained in a variety of ways. For example,

client 720 can perform service-discovery operations and then select which service bindings are applicable, or it can use information obtained from previous responses.

[0225] In another client-side embodiment, illustrated in **FIG. 7b**, Service Access Point 760 directly supports a native protocol 770 of client 780. Service Access Point 760 will translate messages internally between XML and that native protocol 770, thereby enabling client 780 to participate within the NEMO system. To effect such support, native protocol 770 (or a combination of native protocol 770 and the execution environment) must provide any needed information in some form to Service Access Point 760, which generates an appropriate request and, if necessary, determines a suitable target service binding.

[0226] On the service-side, multiple patterns of interaction between a client's Service Access Point and service-providing NEMO-enabled endpoints can be supported. As with the client-side, the interaction patterns can be tailored and may vary based on a variety of criteria, including the nature of the request, the underlying communication network, and the nature of the application and/or transport protocols associated with any targeted service bindings.

[0227] A relatively simple type of service-side interaction pattern is illustrated in **FIG. 7c**, in which Service Access Point 711 communicates directly with the desired service-providing node 712 in a point-to-point manner.

[0228] Turning to **FIG. 7d**, Service Access Point 721 may initiate communication directly with (and may receive responses directly from) multiple potential service providers 725. This type of interaction pattern may be implemented by relaying multiple service bindings from the client for use by Service Access Point 721; or a broadcast or multicast network could be utilized by Service Access Point

721 to relay messages. Based on preferences specified in the request, Service Access Point 721 may choose to collect and collate responses, or simply return the first acceptable response.

[0229] In **FIG. 7e**, Service Access Point 731 doesn't directly communicate with any targeted service-providing endpoints 735. Instead, requests are routed through an intermediate node 733 which relays the request, receives any responses, and relays them back to Service Access Point 731.

[0230] Such a pattern of interaction may be desirable if Service Access Point 731 is unable or unwilling to support directly any of the service bindings associated with service-providing endpoints 735, but can establish a relationship with intermediate node 733, which is willing to act as a gateway. Alternatively, the client may not be able to discover or otherwise determine the service bindings for any suitable service-providing nodes, but may be willing to allow intermediate node 733 to attempt to discover any suitable service providers. Finally, Service Access Point 731 may want to take advantage of intermediate node 733 because it supports more robust collection and collating functionality, which in turn permits more flexible communication patterns between Service Access Point 731 and service providers such as endpoint nodes 735.

[0231] In addition to the above basic service-side interaction patterns, combinations of such patterns or new patterns can be implemented within the Service Access Point. Although the Service Access Point is intended to provide a common interface, its implementation will typically be strongly tied to the characteristics of the communication models and associated protocols employed by given NEMO-enabled endpoints.

[0232] In practice, the Service Access Point can be used to encapsulate the logic for handling the marshalling and un-marshaling of I/O related data, such as serializing objects to appropriate representations, such as an XML representation (with a format expressed in WSDL), or one that envelopes XML-encoded objects in the proper format.

[0233] In a preferred embodiment, the SAP also encapsulates logic for communication via one or more supported application, session, and/or transport protocols, such as service invocation over HTTP using SOAP enveloping.

[0234] Finally, in some embodiments, the SAP may encapsulate logic for providing message integrity and confidentiality, such as support for establishing SSL/TLS sessions and/or signing/verifying data via standards such as XML-Signature and XML-Encryption. When the specific address of a service interface is unknown or unspecified (for example, when invoking a service across multiple nodes based on some search criteria), the SAP may encapsulate the logic for establishing an initial connection to a default/initial set of NEMO nodes where services can be discovered or resolved.

[0235] The following is an example, non-limiting embodiment of a high-level API description exported by one SAP embodiment:

[0236] *ServiceAccessPoint::Create(Environment[]) → ServiceAccessPoint* – this is a singleton interface that returns an initialized instance of a SAP. The SAP can be initialized based on an optional set of environmental parameters.

[0237] *ServiceAccessPoint::InvokeService(Service Request Message, Boolean) → Service Response Message* – a synchronous service invocation API is supported where the client (using WSDL) forms an XML service request message,

and receives an XML message in response. The API also accept a Boolean flag indicating whether or not the client should wait for a response. Normally, the flag will be true, except in the case of messages with no associated response, or messages to which responses will be delivered back asynchronously via another channel (such as via notification). The resulting message may also convey some resulting error condition.

[0238] *ServiceAccessPoint::ApplyIntegrityProtection(Boolean, Desc[]) → Boolean* – This API allows the caller to specify whether integrity protection should be applied, and to which elements in a message it should be applied.

[0239] *ServiceAccessPoint::ApplyConfidentiality(Boolean, Desc[]) → Boolean* – This API allows the caller to specify whether confidentiality should be applied and to which objects in a message it should be applied.

[0240] ServiceAccessPoint::SetKeyCallbacks(SigningKeyCallback, SignatureVerificationKeyCallback, EncryptionKeyCallback, DecryptionKeyCallback) → Boolean

As indicated in the previous APIs, when a message is sent or received it may contain objects which require integrity protection or confidentiality . This API allows the client to set up any necessary hooks between itself and the SAP to allow the SAP to obtain keys associated with a particular type of trust management operation. In one embodiment, the interface is based on callbacks supporting integrity protection through digital signing and verification, and confidentiality through encryption and decryption. In one embodiment, each of the callbacks is of the form:

KeyCallback(KeyDesc) → Key[]

where KeyDesc is an optional object describing the key(s) required and a list of appropriate keys is returned. Signatures are validated as part of receiving response

services messages when using the InvokeService(...) API. If a message element fails verification, an XML message can be returned from InvokeService(...) indicating this state and the elements that failed verification.

3.4. Service Adaptation Layer

[0241] As noted above, the Service Adaptation Layer provides a common way for service providers to expose their services, process requests and generate responses for services, and orchestrate services in the NEMO framework. It also provides a foundation on which other specific service interface bindings can be implemented. In one embodiment, WSDL is used to describe a service's interface within the system.

[0242] Such a service description might, in addition to defining how to bind to a service on a particular interface, also include a list of one or more authorization service providers that will be responsible for authorizing access to the service, a pointer to a semantic description of the purpose and usage of the service, and a description of the necessary orchestration for composite services resulting from the choreographed execution of one or more other services.

[0243] In addition to serving as the logical point at which services are exposed, the Service Adaptation Layer also preferably encapsulates the concrete representations of the NEMO data types and objects specified in NEMO service profiles for platforms that are supported by a given participant. It also contains a mechanism for mapping service-related messages to the appropriate native service implementation.

[0244] In one embodiment, the NEMO framework does not mandate how the Service Adaptation Layer for a given platform or participant is realized. In situations where a service-providing node does not require translation of its native service protocols – i.e., exposing its services only to client nodes that can communicate via

that native protocol – then that service-providing node need not contain a Service Adaptation Layer.

[0245] Otherwise, its Service Adaptation Layer will typically contain the following elements, as illustrated in **FIG. 8**:

[0246] *Entry Points* – a layer encapsulating the service interface entry points 810 and associated WSDL bindings. Through these access points, other nodes invoke services, pass parameter data, and collect results.

[0247] *Message Processing Logic* – a layer 820 that corresponds to the logic for message processing, typically containing a message pump 825 that drives the processing of messages, some type of XML data binding support 826, and low level XML parser and data representation support 827.

[0248] *Native Services* – a layer representing the native services available (onto which the corresponding service messages are mapped), including a native services API 830 and corresponding implementation 840.

3.5. Workflow Collator

[0249] In a preferred embodiment, a Workflow Collator (WFC) helps fulfill most nontrivial NEMO service requests by coordinating the flow of events of a request, managing any associated data including transient and intermediate results, and enforcing the rules associated with fulfillment. Examples of this type of functionality can be seen in the form of transaction coordinators ranging from simple transaction monitors in relational databases to more generalized monitors as seen in Microsoft MTS/COM+.

[0250] In one embodiment, the Workflow Collator is a programmable mechanism through which NEMO nodes orchestrate the processing and fulfillment of

service invocations. The WFC can be tailored toward a specific NEMO node's characteristics and requirements, and can be designed to support a variety of functionality ranging from traditional message queues to more sophisticated distributed transaction coordinators. A relatively simple WFC might provide an interface for storage and retrieval of arbitrary service-related messages. By building on this, it is possible to support a wide variety of functionality including (i) collection of service requests for more effective processing; (ii) simple aggregation of service responses into a composite response; (iii) manual orchestration of multiple service requests and service responses in order to create a composite service; and (iv) automated orchestration of multiple service requests and service responses in order to create a composite service.

[0251] A basic service interaction pattern begins with a service request arriving at some NEMO node via the node's Service Adaptation Layer. The message is handed off to the WSDL Message Pump that initially will drive and in turn be driven by the WFC to fulfill the request and return a response. In even more complex scenarios, the fulfillment of a service request might require multiple messages and responses and the participation of multiple nodes in a coordinated fashion. The rules for processing requests may be expressed in the system's service description language or using other service orchestration description standards such as BPEL.

[0252] When a message is given to the WFC, the WFC determines the correct rules for processing this request. Depending upon the implementation of the WFC, the service description logic may be represented in the form of a fixed state machine for a set of services that the node exposes or it may be represented in ways that support the processing of a more free form expression of the service processing logic.

[0253] In a preferred embodiment the WFC architecture is modular and extensible, supporting plug-ins. In addition to interpreting service composition and processing rules, the WFC may need to determine whether to use NEMO messages in the context of initiating a service fulfillment processing lifecycle, or as input in the chain of an ongoing transaction. In one embodiment, NEMO messages include IDs and metadata that are used to make these types of determinations. NEMO messages also can be extended to include additional information that may be service transaction specific, facilitating the processing of messages.

[0254] As discussed in greater detail below, notification services are directly supported by various embodiments of the NEMO system. A notification represents a message targeted at interested NEMO-enabled nodes received on a designated service interface for processing. Notifications may carry a diverse set of payload types for conveying information and the criteria used to determine if a node is interested in a notification is extensible, including identity-based as well as event-based criteria.

[0255] In one embodiment, illustrated in **FIG. 9a**, a service-providing NEMO node 910 provides a service that requires an orchestration process by its Workflow Collator 914 (e.g., the collection and processing of results from two other service providers) to fulfill a request for that service from client node 940.

[0256] When NEMO-enabled application 942 on client node 940 initiates a request to invoke the service provided by service provider 910, Workflow Collator 914 in turn generates messages to initiate its own requests (on behalf of application 942), respectively, to Service Provider "Y" 922 on node 920 and Service Provider "Z" 932 on node 930. Workflow Collator 914 then collates and processes the results from

these two other service-providing nodes in order to fulfill the original request from client node 940.

[0257] Alternatively, a requested service might not require the services of multiple service-providing nodes; but might instead require multiple rounds or phases of communication between the service-providing node and the requesting client node. As illustrated in **FIG. 9b**, when NEMO-enabled application 942 on client node 940 initiates a request to invoke the service provided by service provider 910, Workflow Collator 914 in turn engages in multiple phases of communication 950 with client node 940 in order to fulfill the original request. For example, Workflow Collator 914 may generate and send messages to client node 940 (via Access Point 944), receive and process the responses, and then generate additional messages (and receive additional responses) during subsequent phases of communication, ultimately fulfilling the original request from client node 940.

[0258] In this scenario, Workflow Collator 914 is used by service provider 910 to keep track (perhaps based on a service-specific session ID or transaction ID as part of the service request) of which phase of the operation it is in with the client for correct processing. As noted above, a state machine or similar mechanism or technique could be employed to process these multiple phases of communication 950.

[0259] **FIG. 9c** illustrates one embodiment of a relatively basic interaction, within service-providing node 960, between Workflow Collator 914 and Message Pump 965 (within the node's Service Adaptation Layer, not shown). As noted above, Workflow Collator 914 processes one or more service requests 962 and generates responses 964, employing a storage and retrieval mechanism 966 to maintain the state of this orchestration process. In this simple example, Workflow Collator 914 is able

to process multiple service requests and responses, which could be implemented with a fairly simple state machine.

[0260] For more complex processing, however, **FIG. 9d** illustrates a node architecture that can both drive or be driven in performing service orchestration. Such functionality includes the collection of multiple service requests, aggregation of responses into a composite response, and either manual or automated orchestration of multiple service requests and responses in order to create a composite service.

[0261] A variety of scenarios can be supported by the architecture surrounding Workflow Collator 914 in **FIG. 9d**. For example, by having a NEMO node combine its functionality with that of an external coordinator 970 that understands the semantics of process orchestration (such as a Business Process Language engine driven by a high level description of the business processes associated with services) or resource usage semantics (such as a Resource Description Framework engine which can be driven by the semantic meaning of resources in relationship to each other), it is possible to create more powerful services on top of simpler ones. Custom External BPL 972 and/or RDF 973 processors may leverage external message pump 975 to execute process descriptions via a manual orchestration process 966, i.e., one involving human intervention.

[0262] In addition to relying on a manually driven process that relies on an external coordinator working in conjunction with a NEMO node's message pump, it is also possible to create an architecture where modules may be integrated directly with Workflow Collator 914 to support an automated form of service coordination and orchestration 968. For example, for regular types of service orchestration patterns, such as those represented in BPEL and EBXML and communicated in the

web service bindings associated with a service interface, Workflow Collator 914 can be driven directly by a description and collection of request and response messages 967 that arrive over time. In this scenario, a composite response message is pushed to Message Pump 965 only when the state machine associated with the given orchestration processor plug-in (e.g., BPEL 982 or EBXML 983) has determined that it is appropriate.

[0263] Following is an embodiment of a relatively high-level API description exported by an embodiment of a NEMO Workflow Collator:

[0264] *WorkflowCollator::Create(Environment[])* → *WorkflowCollator* – this is a singleton interface that returns an initialized instance of a WFC. The WFC can be initialized based on an optional set of environmental parameters.

[0265] *WorkflowCollator::Store(Key[], XML Message)* → *Boolean* – this API allows the caller to store a service message within the WFC via a set of specified keys.

[0266] *WorkflowCollator::RetrieveByKey(Key[], XML Message)* → *XML Message[]* – this API allows the caller to retrieve a set of messages via a set of specified keys. The returned messages are no longer contained within the WFC.

[0267] *WorkflowCollator::PeekByKey(Key[], XML Message)* → *XML Message[]* – this API allows the caller to retrieve a set of messages via a set of specified keys. The returned messages are still contained within the WFC.

[0268] *WorkflowCollator::Clear()* → *Boolean* – this API allows the caller to clear any messages stored within the WFC.

[0269] As an alternative to the relatively rigid BPEL orchestration standard, another embodiment could permit a more ad hoc XML-based orchestration

description – e.g., for a more dynamic application, such as a distributed search. Consider the following description that could be interpreted by a NEMO Workflow Collator (and could possibly even replace an entire service given a sufficiently rich language):

[0270] <WSDL>

<NEMO Orchestration Descriptor>

<Control Flow>	e.g., EXECUTE Service A; If result = Yes then Service B; Else Service C
<Shared State/Context>	e.g., Device State
<Transactions>	e.g., State, Rollback, etc
<Trust/Authorization>	Note that Trust not necessarily transitive

3.6. Exemplary DRM Engine Architecture

[0271] In the context of the various embodiments of the NEMO node architecture described above, **FIG. 10** illustrates the integration of a modular embodiment of a DRM Engine 1000 into a NEMO content consumption device, thereby facilitating its integration into many different devices and software environments.

[0272] Host application 1002 typically receives a request to access a particular piece of content through its user interface 1004. Host application 1002 then sends the request, along with relevant DRM engine objects (preferably opaque to the host application), to DRM engine 1000. DRM engine 1000 may make requests for additional information and cryptographic services to host services module 1008 through well-defined interfaces. For example, DRM engine 1000 may ask host services 1008 whether a particular link is trusted, or may ask that certain objects be decrypted. Some of the requisite information may be remote, in which case host

services 1008 can request the information from networked services through a service access point 1014.

[0273] Once DRM engine 1000 has determined that a particular operation is permitted, it indicates this and returns any required cryptographic keys to host services 1008 which, under the direction of host application 1002, relies on content services 1016 to obtain the desired content and manage its use. Host services 1008 might then initiate the process of media rendering 1010 (e.g., playing the content through speakers, displaying the content on a screen, etc.), coordinated with cryptography services 1012 as needed.

[0274] The system architecture illustrated in **FIG. 10** is a relatively simple example of how the DRM engine can be used in applications, but it is only one of many possibilities. For example, in other embodiments, the DRM engine can be integrated into packaging applications under the governance of relatively sophisticated policy management systems. Both client (content consumption) and server (content packaging) applications of the DRM engine, including descriptions of the different types of DRM-related objects relied upon by such applications, will be discussed below, following a description of one embodiment of the internal architecture of the DRM engine itself.

[0275] DRM Engine 1100, illustrated in **FIG. 11**, relies on a virtual machine, control VM 1110, for internal DRM processing (e.g., executing control programs that govern access to content) within a broad range of host platforms, utilizing host environment 1120 (described above, and in greater detail below) to interact with the node's host application 1130 and, ultimately, other nodes within, e.g., the NEMO or other system.

[0276] In one embodiment, control VM 1110 is a virtual machine used by an embodiment of DRM Engine 1100 to execute control programs that govern access to content. Following is a description of the integration of control VM 1110 into the architecture of DRM engine 1100, as well as some of the basic elements of the control VM, including details about its instruction set, memory model, code modules, and interaction with host environment 1120 via system calls 1106.

[0277] In one embodiment, control VM 1110 is a relatively small-footprint virtual machine that is designed to be easy to implement using various programming languages. It is based on a stack-oriented instruction set that is designed to be minimalist in nature, without much concern for execution speed or code density. However, it will be appreciated that, if execution speed and/or code density were issues in a given application, conventional techniques (e.g., data compression) could be used to improve performance.

[0278] Control VM 1100 is suitable as a target for low or high level programming languages, and supports languages such as assembler, C, and FORTH. Compilers for other languages, such as Java or custom languages, could also be implemented with relative ease.

[0279] Control VM 1110 is designed to be hosted within DRM Engine 1100, including host environment 1120, as opposed to being run directly on a processor or in silicon. Control VM 1110 runs programs by executing instructions stored in Code Modules 1102. Some of these instructions can make calls to functions implemented outside of the program itself by making one or more System Calls 1106, which are either implemented by Control VM 1110 itself, or delegated to Host Environment 1120.

[0280] *Execution Model*

[0281] Control VM 1110 executes instructions stored in code modules 1102 as a stream of byte code loaded in memory 1104. Control VM 1110 maintains a virtual register called the program counter (PC) that is incremented as instructions are executed. The VM executes each instruction, in sequence, until the `OP_STOP` instruction is encountered, an `OP_RET` instruction is encountered with an empty call stack, or an exception occurs. Jumps are specified either as a relative jump (specified as a byte offset from the current value of PC), or as an absolute address.

[0282] *Memory Model*

[0283] In one embodiment, control VM 1110 has a relatively simple memory model. VM memory 1104 is separated into a data segment (DS) and a code segment (CS). The data segment is a single, flat, contiguous memory space, starting at address 0. The data segment is typically an array of bytes allocated within the heap memory of host application 1130 or host environment 1120. For a given VM implementation, the size of the memory space is preferably fixed to a maximum; and attempts to access memory outside of that space will cause faults and terminate program execution. The data segment is potentially shared between several code modules 1102 concurrently loaded by the VM. The memory in the data segment can be accessed by memory-access instructions, which can be either 32-bit or 8-bit accesses. 32-bit memory accesses are accomplished using the big-endian byte order. No assumptions are made with regard to alignment between the VM-visible memory and the host-managed memory (host CPU virtual or physical memory).

[0284] In one embodiment, the code segment is a flat, contiguous memory space, starting at address 0. The code segment is typically an array of bytes allocated within the heap memory of host application 1130 or host environment 1120.

[0285] Control VM 1110 may load several code modules, and all of the code modules may share the same data segment (each module's data is preferably loaded at a different address), but each has its own code segment (e.g., it is preferably not possible for a jump instruction from one code module 1102 to cause a jump directly to code in another code module 1102).

[0286] *Data Stack*

[0287] In a preferred embodiment, the VM has a notion of a data stack, which represents 32-bit data cells stored in the data segment. The VM maintains a virtual register called the stack pointer (SP). After reset, SP points to the end of the data segment, and the stack grows downward (when data is pushed onto the data stack, the SP registers are decremented). The 32-bit values on the stack are interpreted either as 32-bit addressed, or 32-bit signed, integers, depending on the instruction referencing the stack data.

[0288] *Call Stack*

[0289] In one embodiment, control VM 1110 manages a call stack for making nested subroutine calls. The values pushed onto this stack cannot be read or written directly by the memory-access instructions, but are used indirectly by the VM when executing `OP_JSR` and `OP_RET` instructions. For a given VM profile, the size of this return address stack is preferably fixed to a maximum, which will allow a certain number of nested calls that cannot be exceeded.

[0290] *Instruction Set*

[0291] In one embodiment, control VM 1110 uses a relatively simple instruction set. Even with a limited number of instructions; however, it is still possible to express simple programs. The instruction set is stack-based: except for the OP_PUSH instruction, none of the instructions have direct operands. Operands are read from the data stack, and results are pushed onto the data stack. The VM is a 32-bit VM: all the instructions in this illustrative embodiment operate on 32-bit stack operands, representing either memory addresses or signed integers. Signed integers are represented using a 2s complement binary encoding.

[0292] An illustrative instruction set used in one embodiment is shown below:

OP CODE	Name	Operands	Description
OP_PUSH	Push Constant	N (direct)	Push a constant on the stack
OP_DROP	Drop		Remove top of stack
OP_DUP	Duplicate		Duplicate top of stack
OP_SWAP	Swap		Swap top two stack elements
OP_ADD	Add	A, B	Push the sum of A and B (A+B)
OP_MUL	Multiply	A, B	Push the product of A and B (A*B)
OP_SUB	Subtract	A, B	Push the difference between A and B (A-B)
OP_DIV	Divide	A, B	Push the division of A by B (A/B)
OP_MOD	Modulo	A, B	Push A modulo B (A%B)
OP_NEG	Negate	A	Push the 2s complement negation of A (-A)
OP_CMP	Compare	A	Push -1 if A negative, 0 if A is 0, and 1 is a positive
OP_AND	And	A, B	Push bit-wise AND of A and B (A & B)
OP_OR	Or	A, B	Push the bit-wise OR of A and B (A | B)
OP_XOR	Exclusive Or	A, B	Push the bit-wise eXclusive OR of A and B (A ^ B)
OP_NOT	Logical Negate	A	Push the logical negation of A (1 if A is 0, and 0 if A is not 0)
OP_SHL	Shift Left	A, B	Push A logically shifted left by B bits (A << B)
OP_SHR	Shift Right	A, B	Push A logically shifted right by B bits (A >> B)

OP_JSR	Jump to Subroutine	A	Jump to subroutine at absolute address A
OP_JSRR	Jump to Subroutine (Relative)	A	Jump to subroutine at PC+A
OP_RET	Return from Subroutine		Return from subroutine
OP_BRA	Branch Always	A	Jump to PC + A
OP_BRP	Branch if Positive	A, B	Jump to PC+A if B > 0
OP_BRN	Branch if Negative	A, B	Jump to PC+A if B < 0
OP_BRZ	Branch if Zero	A, B	Jump to PC+A if B is 0
OP_JMP	Jump	A	Jump to A
OP_PEEK	Peek	A	Push the 32-bit value at address A
OP_POKE	Poke	A, B	Store the 32-bit value B at address A
OP_PEEKB	Peek Byte	A	Push the 8-bit value at address A
OP_POKEB	Poke Byte	A, B	Store the least significant bits of B at address A
OP_PUSHSP	Push Stack Pointer		Push the value of SP
OP_POPSP	Pop Stack Pointer	A	Set the value of SP to A
OP_CALL	System Call	A	Perform System Call with index A
OP_STOP	Stop		Terminate Execution

[0293] *Module Format*

[0294] In one embodiment, code modules 1102 are stored in an atom-based format that is essentially equivalent to the atom structure used in the MPEG-4 file format. An atom consists of 32 bits, stored as 4-octets in big-endian byte order, followed by a 4-octet type (usually octets that correspond to ASCII values of letters of the alphabet), followed by the payload of the atom (size-8 octets).

3.7. DRM Client-Server Architecture: Content Consumption and Packaging

[0295] As noted above, DRM client-side consuming applications (e.g., media players) consume DRM content (e.g., play a song, display a movie, etc.). DRM service-side packaging applications (typically residing on a server) package content (e.g., associate with the content relevant usage and distribution rights, cryptographic keys, etc.) targeted to DRM clients.

[0296] **FIG. 12a** illustrates one embodiment of the main architectural elements of a DRM client. Host application 1200 interfaces with a device user (e.g., the owner of a music player) through user interface 1210. The user might, for example, request access to protected content and receive metadata along with the content (e.g., text displaying the name of the artist and song title, along with the audio for the song itself).

[0297] Host application 1200, in addition to interacting with user interface 1210, also performs various functions necessary to implement the user's request, which may include managing interaction with the other DRM client modules to which it delegates certain functionality. For example, host application 1200 may manage interaction with the file system to extract the requested content. Host application also preferably recognizes the protected content object format and issues a request to the DRM engine 1220 to evaluate the DRM objects that make up the license (e.g., by running the relevant control program) to determine whether permission to access the protected content should be granted.

[0298] If permission is granted, Host Application 1200 might also need to verify required signatures and delegate to crypto services 1230 any other general purpose cryptographic functions required by DRM engine 1220. DRM Engine 1220

is responsible for evaluating the DRM objects, confirming or denying permission, and providing the keys to host application 1200 to decrypt the content.

[0299] Host services 1240 provides DRM Engine 1220 with access to data managed by (as well as certain library functions implemented by) host application 1200. Host application 1200 interacts with content services 1250 to access the protected content, passing to DRM engine 1220 only that portion of the content requiring processing. Content services 1250 acquires the content from external media servers and stores and manages the content, relying on the client's persistent storage mechanisms.

[0300] Once the content is cleared for access, host application 1200 interacts with media rendering engine 1260 (e.g., by delivering keys) to decrypt and render the content via the client's AV output facilities. Some of the information needed by DRM Engine 1220 may be available in-band with the content, and can be acquired and managed via content services 1250, while other information may need to be obtained through external NEMO DRM services or some other source.

[0301] In a preferred embodiment, all of the cryptographic operations (encryption, signature verification, etc.) are handled by crypto services 1230, which interacts indirectly with DRM engine 1220 via host services 1240, which forwards requests. Crypto services 1230 can also be used by media rendering engine 1260 to perform content decryption.

[0302] Turning to the service side, **FIG. 12b** illustrates an embodiment of the main architectural elements of an exemplary DRM service-side packaging node. Host application 1200 interfaces with a content packager (e.g., an owner or distributor of music content) through user interface 1210. The packager might, for example,

provide content and licensing information to host application 1200 so that the content can be protected (e.g., encrypted and associated with limited access rights) and distributed to various end user and intermediate content providing nodes.

[0303] Host application 1200, in addition to interacting with user interface 1210, can also perform various functions necessary to implement the packager's request, including, for example, managing interaction with the other DRM packaging modules to which it delegates certain functionality. For example, it may manage interaction with general crypto services 1235 to encrypt the content. It may also create a content object that contains or references the content and contains or references a license (e.g., after DRM packaging engine 1225 creates the DRM objects that make up the license). Metadata can be associated with the license that explains what the license is about in a human-readable way (e.g., for potential client users to view).

[0304] As noted above, host application 1200 interacts with the user via user interface 1210. It is responsible for getting information such as a content reference and the action(s) the packager wants to perform (e.g., who to bind the content to). It can also display information about the packaging process such as the text of the license issued and, if a failure occurs, the reason for this failure. Some information needed by host application 1200 may require the use of NEMO Services 1270 (e.g., to leverage services such as authentication or authorization as well as membership).

[0305] In one embodiment, host application 1200 delegates to media format services 1255 responsibility for managing all media format operations, such as transcoding and packaging. General crypto services 1235 is responsible for issuing and verifying signatures, as well as encrypting and decrypting certain data. The

request for such operations could be issued externally or from DRM packaging engine 1225 via host services 1240.

[0306] In one embodiment, content crypto services 1237 is logically separated from general crypto services 1235 because it is unaware of host application 1200. It is driven by media format services 1255 at content packaging time with a set of keys previously issued by DRM packaging engine 1225 (all of which is coordinated by host application 1200).

3.8. DRM Content Protection and Governance Objects

[0307] In an illustrative scenario, a content provider uses a host application that relies on a DRM packager engine to create a set of objects that protect the content and govern its use, including conveying the information necessary for obtaining the content encryption keys. The term, *license,* is used to encompass this set of objects.

[0308] In a preferred embodiment, the content and its license are logically separate, but are bound together by internal references using object IDs. The content and license are usually stored together, but could be stored separately if necessary or desirable. A license can apply to more than one item of content, and more than one license can apply to any single item of content.

[0309] **FIG. 13** illustrates an embodiment of such a license, including the relationships among the set of objects discussed below. Note that control object 1320 and controller object 1330 are both signed objects in this embodiment, so that the DRM client engine can verify that the control information comes from a trusted source prior to providing the host application with permission to access the protected content. In this embodiment, all of these objects, with the exception of content object 1300, are created by the DRM client engine.

[0310] *Content object* – Content object 1300 represents the encrypted content 1304, using a unique ID 1302 to facilitate the binding between the content and its associated key. Content object 1300 is an "external" object. The format and storage of encrypted content 1304 (e.g., MP4 movie file, MP3 music track, etc.) is determined by the host application (or delegated to a service), based in part upon the type of content. The format of the content also provides support for associating ID 1302 with encrypted content 1304. The packager's host application encrypts the content in a format-dependent manner, and manages content object 1300, using any available cryptosystem (e.g., using a symmetric cipher, such as AES).

[0311] *ContentKey object* – ContentKey object 1310 represents the encrypted key data 1314 (including a unique encryption key(s), optionally stored internally within the object), and also has a corresponding unique ID 1312. Preferably, this key data, if contained within ContentKey object 1310, is itself encrypted so that it can only be identified by those authorized to decrypt the content. ContentKey object 1310 also specifies which cryptosystem was used to encrypt this key data. This cryptosystem, an embodiment of which is discussed in greater detail below, is referred to as the "key distribution system."

[0312] *Control object* – Control object 1320 includes and protects the control program (e.g., control byte code 1324) that represents the rules that govern the use of the keys used to encrypt and decrypt the content. It also includes ID 1322 so that it can be bound to the corresponding ContentKey object. As noted above, control object 1320 is signed so that the DRM client engine can verify the validity of the binding between the ContentKey 1310 and control 1320, as well as the binding between the ContentKey ID 1312 and the encrypted key data 1314. The validity of control byte

code 1324 can optionally be derived by verifying a secure hash (e.g., control hash 1338, if available) contained in controller object 1330.

[0313] *Controller object* – Controller object 1330 represents the binding between the keys and the rules governing their control, using IDs 1312 and 1322, respectively, to bind the ContentKey 1310 and control 1320 objects. Controller object 1330 governs the use of protected content by controlling application of the rules to that content – i.e., by determining which control governs the use of which ContentKey object 1310. Controller object 1330 also contains a hash 1336 value for each of the ContentKey objects 1310 that it references, in order to prevent tampering with the binding between each ContentKey object 1310 and its corresponding encrypted key data 1314. As noted above, controller objects 1330 are preferably signed (e.g., by a packager application that has a certificate allowing it to sign controller objects, using public key or symmetric key signatures, as discussed below) to enable verification of the validity of the binding between the ContentKey 1310 and control 1320 objects, as well as the binding between the ContentKey ID 1312 and the encrypted key data 1314. As also noted above, controller object 1330 also optionally contains control hash 1338, which allows the validity of control object 1320 to be derived without having to separately verify its signature.

[0314] *Symmetric Key Signature* – In a preferred embodiment, a symmetric key signature is the most common type of signature for controller objects 1330. In one embodiment, this type of signature is implemented by computing a MAC (Message Authentication Code) of the controller object 1330, keyed with the same key as the key represented by the ContentKey object 1310.

[0315] *Public Key Signature* – In a preferred embodiment, this type of signature is used when the identity of the signer of the controller object 1330 needs to be asserted uniquely. This type of signature is implemented with a public key signature algorithm, signing with the private key of the principal who is asserting the validity of this object. When using this type of signature, the ContentKey binding information carried in the controller object 1330 preferably contains a hash 1336 of the key contained in the ContentKey object 1310, concatenated with a fingerprint of the signing private key (typically a hash of the private key). This binding ensures that the signer of the object has knowledge of the key used to protect the content.

[0316] *Protector object* – Protector object 1340 provides protected access to content by controlling the use of keys used to encrypt and decrypt that content. Protector object 1340 binds content object 1300 to ContentKey object 1310 in order to associate protected content with its corresponding key(s). To accomplish this binding, it includes references 1342 and 1344, respectively, to the IDs 1302 and 1312 of content 1300 and ContentKey 1310. In one embodiment, protector object 1340 contains information not only as to which key was used to encrypt one or more content items, but also as to which encryption algorithm was employed. In one embodiment, if content reference 1342 references more than one content object 1300, ContentKey reference 1344 may still reference only one ContentKey object 1310, indicating that all of those content items were encrypted using the same encryption algorithm and the same key.

3.9. DRM Node and Link Objects

[0317] While **FIG. 13** illustrates the content protection and governance objects created by DRM engines to control access to protected content, **FIG. 14**

illustrates the DRM objects that represent entities in the system (e.g., users, devices or groups), as well as the relationships among those entities.

[0318] While **FIG. 4**, discussed above, illustrates a conceptual embodiment of a node or authorization graph depicting these entities and their relationships, **FIG. 14** illustrates two types of objects that implement an embodiment of this conceptual graph: *vertex* (or *"node"*) objects (1400a and 1400b), which represent entities and their attributes, and *link* objects (1420), which represent the relationships among node objects. In one embodiment, the DRM engine, by executing control programs, instigates one or more usage patterns involving these objects – e.g., encrypting a song and associating it with a license that restricts its distribution to particular individuals. Yet, the DRM engine in this embodiment does not specify, implicitly or explicitly, the semantics attached to these objects (e.g., to which individuals the song may be distributed).

[0319] In one embodiment this semantic context, referred to as a DRM profile, is defined within the attributes of the node objects themselves. A DRM profile may include descriptions of these entities and the various roles and identities they represent, typically expressed using node attributes (1401a and 1401b). As discussed above, a link 1420 between two nodes 1400a and 1400b could represent various types of semantic relationships. For example, if one node was a "user" and the other was a "device," then link 1420 might represent "ownership." If the other node was a "user group" instead of a "device," then link 1420 might represent "membership." Link 1420 might be unidirectional in one scenario and bidirectional in another (e.g., representing two links between the same two nodes).

[0320] Node objects 1400a and 1400b also typically have object confidentiality protection asymmetric key pairs (e.g., private key 1405a and public key 1406a of node 1400a, and private key 1405b and public key 1406b of node 1400b) to limit confidential information to authorized portions of the node. Confidential information targeted at a node will be encrypted with that node's confidentiality protection public key. Optionally, a content protection asymmetric key pair (e.g., private key 1403a and public key 1403b of node 1400a, and private key 1403b and public key 1403b of node 1400b) can be used in conjunction with link objects when the system uses a ContentKey derivation system for ContentKey distribution, as discussed in greater detail below. Content items themselves may be protected with content protection symmetric keys, such as symmetric key 1402a of node 1400a and key 1402b of node 1400b.

[0321] As noted above, in one embodiment link objects (e.g., link 1420) represent relationships between nodes. The semantics of these relationships can be stored in node attributes (e.g., 1401a of node 1400a and 1401b of node 1400b), referenced from within the link objects (e.g., node reference 1422 to node 1400a and node reference 1424 to node 1400b). Link objects can also optionally contain cryptographic data (e.g., key derivation info 1426) that enables the link object to be used for ContentKey derivations, as discussed below.

[0322] In one embodiment the link object itself is a signed object, represented by a directed edge in a graph, such as in **FIG. 4** above. When there exists such a directed edge from one node (e.g., node X) to another (e.g., node Y), this "path" from node X to node Y indicates that node Y is "reachable" from node X. The existence of a path can be used by other DRM objects, e.g., as a condition of performing a

particular function. A control object might check to determine whether a target node is reachable before it allows a certain action to be performed on its associated content object.

[0323] For example, if node D represents a device that wants to perform the "play" action on a content object, a control that governs this content object might test whether a certain node U representing a certain user is reachable from node D (e.g., whether that user is the "owner" of that device), and only allow the "play" action to be performed if that condition is satisfied. To determine if node U is reachable, the DRM engine can run a control program to determine whether there exists a set of link objects that can establish a path (e.g., a direct or indirect relationship) between node D and node U. As noted above, in one embodiment the DRM engine is unaware of the semantics of the relationship; it simply determines the existence of a path, enabling the host application, for example, to interpret this path as a conditional authorization, permitting access to protected content.

[0324] In one embodiment the DRM engine verifies link objects before allowing them to be used to determine the existence of paths in the system node graph. The validity of a link object at any given time may depend upon the particular features of the certificate system (discussed below) used to sign link objects. For example, they may have limited "lifetimes" or be revoked or revalidated from time to time based on various conditions.

[0325] Also, in one embodiment the policies that govern which entities can sign link objects, which link objects can be created, and the lifetime of link objects are not directly handled by the DRM engine. Instead, they may leverage the node attributes information. To facilitate the task of enforcing certain policies, the system

may provide a way to extend standard certificate formats with additional constraint checking. These extensions make it possible to express validity constraints on certificates for keys that sign links, such that constraints (e.g., the type of nodes connected by the link, as well as other attributes), can be checked before a link is considered valid.

[0326] Finally, in one embodiment the link object may contain cryptographic data that provides the user with the nodes' content protection keys for key distribution. That cryptographic data may, for example, contain, in addition to metadata, the private and/or symmetric content protection keys of the "from" node, encrypted with the content protection public key and/or the content protection symmetric key of the "to" node. For example, an entity that has been granted the ability to create link objects that link device nodes and user nodes under a certain policy may check to ensure that it only creates links between node objects that have attributes indicating they are indeed representing a device, and nodes that have attributes indicating that they represent a user.

3.10. DRM Cryptographic Keys

[0327] An example embodiment of a DRM key distribution system is illustrated in **FIG. 15**. The basic principle behind the key distribution system shown in **FIG. 15** is to use link objects to distribute keys in addition to their primary purpose of establishing relationships between node objects.

[0328] As noted above, a control object may contain a control program that determines whether a requested operation should be permitted. That control program may check to determine whether a specific node is reachable via a collection of link objects. The key distribution system shown in **FIG. 15** leverages that search through

a collection of link objects to facilitate the distribution of a key such that it is available to the DRM engine that is executing the control program.

[0329] In one embodiment, each node object that uses the key distribution system has one or more keys. These keys are used to encrypt content keys and other nodes' key distribution keys. Link objects created for use in the same deployment contain some cryptographic data payload that allows key information do be derived when chains of links are processed by the DRM engine.

[0330] With nodes and links carrying keys this way, given a collection of links (e.g., from a node A to a node B . . . to a node Z), any entity that has access to the private keys of node A also has access to the private keys of node Z. Having access to node Z's private keys gives the entity access to any content key encrypted with those keys.

[0331] Node objects that participate in a key distribution system contain keys as part of their data. As illustrated in **FIG. 15**, in one embodiment each node (1500a, 1500b, and 1500c) has three keys:

[0332] *Public Key Kpub[N]* – This is the public part of a pair of public/private keys for the public key cipher. In one embodiment this key (1505a, 1505b and 1505c, respectively, in nodes 1500a, 1500b and 1500c) comes with a certificate (discussed below) so that its credentials can be verified by entities that want to bind confidential information to it cryptographically.

[0333] *Private Key Kpriv[N]* – This is the private part of the public/private key pair. The entity that manages the node is responsible for ensuring that this private key (keys 1515a, 1515b and 1515c, respectively, in nodes 1500a, 1500b and 1500c) is

kept secret. For that reason, in one embodiment this private key is stored and transported separately from the rest of the node information.

[0334] *Symmetric Key Ks[N]* – This key is used with a symmetric cipher (discussed below). Because this private key (keys 1525a, 1525b and 1525c, respectively, in nodes 1500a, 1500b and 1500c) is confidential, the entity that manages the node is responsible for keeping it secret.

[0335] The key distribution system illustrated in **FIG. 15** can be implemented using different cryptographic algorithms, though the participating entities will generally need to agree on a set of supported algorithms. In one embodiment, at least one public key cipher (such as RSA) and one symmetric key cipher (such as AES) are supported.

[0336] The following notation refers to cryptographic functions:

Ep(Kpub[N], M) means "the message M encrypted with the public key Kpub of node N, using a public key cipher"

Dp(Kpriv[N], M) means "the message M decrypted with the private key Kpriv of node N using a public key cipher"

Es(Ks[N], M) means "the message M encrypted with the symmetric key Ks of node N using a symmetric key cipher"

Ds(Ks[N], M) means "the message M decrypted with the symmetric key Ks of node N using a symmetric key cipher"

[0337] Targeting a "ContentKey" to a node means making that key available to the entities that have access to the private keys of that node. In one embodiment binding is done by encrypting the key using one or both of the following methods:

Public Binding: Create a ContentKey object that contains Ep(Kpub[N], CK)

Symmetric Binding: Create a ContentKey object that contains Es(Ks[N], CK)

[0338] In this embodiment, symmetric binding is preferably used whenever possible, as it uses a less computationally intensive algorithm that is less onerous on the receiving entity. However, the entity (e.g., a content packager) that creates the ContentKey object may not always have access to Ks[N]. In that case, public binding can be used, as Kpub[N] should be available, as it is not confidential information. Kpub[N] will usually be made available to entities that need to target ContentKeys, accompanied by a certificate that can be inspected by the entity to decide whether Kpub[N] is indeed the key of a node that can be trusted to handle the ContentKey in accordance with some agreed-upon policy.

[0339] To allow entities to have access to the distribution keys of all reachable nodes, in one embodiment link objects contain a "payload." That payload allows any entity that has access to the private keys of the link's "from node" to also have access to the private keys of the link's "to node." In this manner, an entity can decrypt any ContentKey targeted to a node that is reachable from its node.

[0340] Thus, returning to **FIG. 15**, link 1530a, which links node 1500a to node 1500b, contains a payload that is created by encrypting the private keys 1515b and 1525b of node 1500b with either the symmetric key 1515a of node 1500a or, if unavailable (e.g., due to its confidentiality), with the public key 1525a of node 1500a. Similarly, link 1530b, which links node 1500b to node 1500c, contains a payload that is created by encrypting the private keys 1515c and 1525c of node 1500c with either the symmetric key 1515b of node 1500b or, if unavailable, with the public key 1525b of node 1500b.

[0341] When a DRM engine processes link objects, it processes the payload of each link to update an internal chain 1550 of keys to which it has access. In one embodiment the payload of a link from node A to node B consists of either:

Public derivation information
Ep(Kpub[A], {Ks[B],Kpriv[B]})

or

Symmetric derivation information
Es(Ks[A], {Ks[B],Kpriv[B]})

Where {Ks[B],Kpriv[B} is a data structure containing Ks[B] and Kpriv[B].

[0342] The public derivation information is used to convey the private keys of node B, Ks[B] and Kpriv[B], to any entity that has access to the private key of node A, Kpriv[A]. The symmetric derivation information is used to convey the private keys of node B, Ks[B] and Kpriv[B], to any entity that has access to the symmetric key of node A, Kpriv[A].

[0343] Thus, with reference to key chain 1550, an entity that has access to the private keys of node 1500a (private key 1515a and symmetric key 1525a) enables the DRM engine to utilize these private keys 1560 as a "first link" in (and starting point in generating the rest of) key chain 1550. Scuba keys 1560 are used to decrypt 1555a the ContentKey object within link 1530a (using private key 1515a for public derivation if public binding via public key 1505a was used, or symmetric key 1525a for symmetric derivation if symmetric binding via symmetric key 1525a was used), resulting in the next link 1570 in key chain 1550 – i.e., the confidential keys of node 1500b (private key 1515b and symmetric key 1525b). The DRM engine uses these keys 1570 in turn to decrypt 1555b the ContentKey object within link 1530b (using private key 1515b for public derivation if public binding via public key 1505b was

used, or symmetric key 1525b for symmetric derivation if symmetric binding via symmetric key 1525b was used), resulting in the final link 1580 in key chain 1550 – i.e., the confidential keys of node 1500c (private key 1515c and symmetric key 1525c).

[0344] Since, in one embodiment, the DRM engine can process links in any order, it may not be able to perform a key derivation at the time a link is processed (e.g., because the keys of the "from" node of that link have not yet been derived). In that case, the link is remembered, and processed again when such information becomes available (e.g., when a link is processed in which that node is the "to" node).

3.11. DRM Certificates

[0345] As noted above, in one embodiment certificates are used to check the credentials associated with cryptographic keys before making decisions based on the digital signature created with those keys. In one embodiment, multiple certificate technologies can be supported, leveraging existing information typically available as standard elements of certificates, such as validity periods, names, etc. In addition to these standard elements, additional constraints can be encoded to limit potential usage of a certified key.

[0346] In one embodiment this is accomplished by using key-usage extensions as part of the certificate-encoding process. The information encoded in such extensions can be used to enable the DRM engine to determine whether the key that has signed a specific object was authorized to be used for that purpose. For example, a certain key may have a certificate that allows it to sign only those link objects in which the link is from a node with a specific attribute, and/or to a node with another specific attribute.

[0347] The base technology used to express the certificate typically is not capable of expressing such a constraint, as its semantics may be unaware of elements such as links and nodes. In one embodiment such specific constraints are therefore conveyed as key usage extensions of the basic certificate, including a "usage category" and a corresponding "constraint program."

[0348] The usage category specifies which type of objects a key is authorized to sign. The constraint program can express dynamic conditions based on context. In one embodiment a verifier that is being asked to verify the validity of such a certificate is required to understand the relevant semantics, though the evaluation of the key usage extension expression is delegated to the DRM engine. The certificate is considered valid only if the execution of that program generates a successful result.

[0349] In one embodiment, the role of a constraint program is to return a boolean value – e.g., "true" indicating that the constraint conditions are met, and "false" indicating that they are not met. The control program may also have access to some context information that can be used to reach a decision. The available context information may depend upon the type of decision being made by the DRM engine when it requests the verification of the certificate. For example, before using the information in a link object, a DRM engine may verify that the certificate of the key that signed the object allows that key to be used for that purpose. When executing the constraint program, the environment of the DRM engine is populated with information regarding the link's attributes, as well as the attributes of the nodes referenced by that link.

[0350] The constraint program embedded in the key usage extension is encoded, in one embodiment, as a code module (described above). This code module

preferably exports at least one entry point named, for example, "*EngineName.Certificate.<Category>.Check*", where Category is a name indicating which category of certificates need to be checked. Parameters to the verification program will be pushed onto the stack before calling the entry point. The number and types of parameters passed onto the stack depends on the category of certificate extension being evaluated.

4. SYSTEM OPERATION

4.1. Basic Node Interaction

[0351] Having examined various embodiments of the principal architectural elements of the NEMO system, including embodiments in the context of DRM applications, we now turn to the NEMO system in operation – i.e., the sequence of events within and among NEMO nodes that establish the foundation upon which application-specific functionality can be layered.

[0352] In one embodiment, before NEMO nodes invoke application-specific functionality, they go through a process of initialization and authorization. Nodes initially seek to discover desired services (via requests, registration, notification, etc.), and then obtain authorization to use those services (e.g., by establishing that they are trustworthy and that they satisfy any relevant service provider policies).

[0353] This process is illustrated in **FIG. 16**, which outlines a basic interaction between a Service Provider 1600 (in this embodiment, with functionality shared between a Service Providing Node 1610 and an Authorizing Node 1620) and a Service Requester 1630 (e.g., a client consumer of services). Note that this interaction need not be direct. Any number of Intermediary Nodes 1625 may lie in the path between the Service Requester 1630 and the Service Provider 1600. The

basic steps in this process, which will be described in greater detail below, are discussed from the perspectives of both the client Service Requester 1630 and Service Provider 1600.

[0354] From the perspective of the Service Requester 1630, the logical flow of events shown in **FIG. 16** is as follows:

[0355] *Service Discovery* – In one embodiment, Service Requester 1630 initiates a service discovery request to locate any NEMO-enabled nodes that provide the desired service, and obtain information regarding which service bindings are supported for accessing the relevant service interfaces. Service Requester 1630 may choose to cache information about discovered services. It should be noted that the interface/mechanism for Service Discovery between NEMO Nodes is just another service a NEMO Node chooses to implement and expose. The Service Discovery process is described in greater detail below, including other forms of communication, such as notification by Service Providers to registered Service Requesters.

[0356] *Service Binding Selection* – Once candidate service-providing Nodes are found, the requesting Node can choose to target (dispatch a request to) one or more of the service-providing Nodes based on a specific service binding.

[0357] *Negotiation of Acceptable Trusted Relationship with Service Provider* – In one embodiment, before two Nodes can communicate in a secure fashion, they must be able to establish a trusted relationship for this purpose. This may include an exchange of compatible trust credentials (e.g. X.500 certificates, tokens, etc.) in some integrity-protected envelope that may be used to determine identity; and/or it may include establishing a secure channel, such as an SSL channel, based on certificates both parties trust. In some cases, the exchange and negotiation of these credentials

may be an implicit property of the service interface binding (e.g. WS-Security if the WS-I XML Protocol is used when the interface is exposed as a web service, or an SSL request between two well-known nodes). In other cases, the exchange and negotiation of trust credentials may be an explicitly separate step. NEMO provides a standard and flexible framework allowing Nodes to establish trusted channels for communication. It is up to a given Node, based on the characteristics of the Node and on the characteristics of the service involved in the interaction, to determine which credentials are sufficient for interacting with another NEMO Node, and to make the decision whether it trusts a given Node. In one embodiment the NEMO framework leverages existing and emerging standards, especially in the area of security-related data types and protocols. For example, in one embodiment the framework will support using SAML to describe both credentials (evidence) given by service requestors to service providers when they want to invoke a service, as well as using SAML as a way of expressing authorization queries and authorization responses.

[0358] *Creation of Request Message* – The next step is for Requesting Node 1630 to create the appropriate request message(s) corresponding to the desired service. This operation may be hidden by the Service Access Point. As noted above, the Service Access Point provides an abstraction and interface for interacting with service providers in the NEMO framework, and may hide certain service invocation issues, such as native interfaces to service message mappings, object serialization/de-serialization, negotiation of compatible message formats, transport mechanisms or message routing issues, etc.

[0359] *Dispatching of Request* – Once the request message is created, it is dispatched to the targeted service-providing Node(s) – e.g., Node 1610. The

communication style of the request can be synchronous/asynchronous RPC style or message-oriented, based on the service binding and/or preferences of the requesting client. Interacting with a service can be done directly by the transmission and processing of service messages or done through more native interfaces through the NEMO Service Access Point.

[0360] *Receiving Response Message(s)* – After dispatching the request, Requesting Node 1610 receives one or more responses in reply. Depending on the specifics of the service interface binding and the preferences of Requesting Node 1610, the reply(s) can be returned in various ways, including an RPC-style response or notification message. As noted above, requests and replies can be routed to their targeted Node via other Intermediary Node(s) 1625, which may themselves provide a number of services, including: routing, trust negotiation, collation and correlation functions, etc. All services in this embodiment are "standard" NEMO services described, discovered, authorized, bound to, and interacted with within the same consistent framework. The Service Access Point may hide message-level abstractions from the Node. For example from the Node's perspective, invocation of a service may seem like a standard function invocation with a set of simple fixed parameters.

[0361] *Validation of Response re Negotiated Trust Semantics* – In one embodiment, Requesting Node 1630 validates the response message to ensure that it adheres to the negotiated trust semantics between it and the Service Providing Node 1610. This logic typically is completely encapsulated within the Service Access Point.

[0362] *Processing of Message Payload* – Finally, any appropriate processing is then applied based on the (application specific) message payload type and contents.

[0363] Following are the (somewhat similar) logical flow of events from the perspective of the Service Provider 1600:

[0364] *Service Support Determination* – A determination is first made as to whether the requested service is supported. In one embodiment, the NEMO framework doesn't mandate the style or granularity of how a service interface maps as an entry point to a service. In the simplest case, a service interface maps unambiguously to a given service, and the act of binding to and invoking that interface constitutes support for the service. However, it may be the case that a single service interface handles multiple types of requests, or that a given service type contains additional attributes which need to be sampled before a determination can be made as to whether the Node really supports the specifically desired functionality.

[0365] *Negotiation of Acceptable Trusted Relationship with Service Requester* – In some cases, it may be necessary for Service Provider 1600 to determine whether it trusts Requesting Node 1630, and establish a trusted communication channel. This process is explained in detail above.

[0366] *Dispatch Authorization Request to Nodes Authorizing Access to Service Interface* – Service Providing Node 1610 then determines whether Requesting Node 1630 is authorized or entitled to have access to the service, and, if so, under what conditions. This may be a decision based on local information, or on a natively supported authorization decision mechanism. If not supported locally, Service Providing Node 1610 may dispatch an authorization request(s) to a known NEMO authorization service provider (e.g., Authorizing Node 1620) that governs its services, in order to determine if the Requesting Node 1610 is authorized to have access to the requested services. In many situations, Authorizing Node 1620 and Service Providing

Node 1610 will be the same entity, in which case the dispatching and processing of the authorizing request will be local operations invoked through a lightweight service interface binding such as a C function entry point. Once again, however, since this mechanism is itself just a NEMO service, it is possible to have a fully distributed implementation. Authorization requests can reference identification information and/or attributes associated with the NEMO Node itself, or information associated with users and/or devices associated with the Node.

[0367] *Message Processing Upon Receipt of Authorization Response* – Upon receiving the authorization response, if Requesting Node 1630 is authorized, Service Provider 1600 performs the necessary processing to fulfill the request. Otherwide, if Requesting Node 1630 is not authorized, an appropriate "authorization denied" response message can be generated.

[0368] *Return Response Message* – The response is then returned based on the service interface binding and the preferences of Requesting Node 1630, using one of several communication methods, including an RPC-style response or notification message. Once again, as noted above, requests and replies can be routed to their targeted Node via other Intermediary Node(s) 1625, which may themselves provide a number of services, including routing, trust negotiation, collation and correlation functions, etc. An example of a necessary service provided by an Intermediary Node 1625 might be delivery to a notification processing Node that can deliver the message in a manner known to Requesting Node 1630. An example of a "value added" service might be, for example, a coupon service which associates coupons to the response if it knows of the interests of Requesting Node 1630.

4.2. Notification

[0369] As noted above, in addition to both asynchronous and synchronous RPC-like communication patterns, where the client specifically initiates a request and then either waits for responses or periodically checks for responses through redemption of a ticket, some NEMO embodiments also support a pure messaging type of communication pattern based on the notion of notification. The following elements constitute data and message types supporting this concept of notification in one embodiment:

[0370] *Notification* – a message containing a specified type of payload targeted at interested endpoint Nodes.

[0371] *Notification Interest* – criteria used to determine whether a given Node will accept a given notification. Notification interests may include interests based on specific types of identity (e.g., Node ID, user ID, etc.), events (e.g., Node discovery, service discovery, etc.), affinity groups (e.g., new jazz club content), or general categories (e.g., advertisements).

[0372] *Notification Payload* – the typed contents of a notification. Payload types may range from simple text messages to more complex objects.

[0373] *Notification Handler Service Interface* – the type of service provider interface on which notifications may be received. The service provider also describes the notification interests associated with the interface, as well as the acceptable payload types. A Node supporting this interface may be the final destination for the notification or an intermediary processing endpoint.

[0374] *Notification Processor Service* – a service that is capable of matching notifications to interested Nodes, delivering the notifications based on some policy.

[0375] *Notification Originator* – a Node that sends out a notification targeted to a set of interested Nodes and/or an intermediary set of notification processing Nodes.

[0376] The notification, notification interest, and notification payload are preferably extensible. Additionally, the notification handler service interface is preferably subject to the same authorization process as any other NEMO service interface. Thus, even though a given notification may match in terms of interest and acceptable payload, a Node may refuse to accept a notification based on some associated interface policy related to the intermediary sender or originating source of the notification.

[0377] **FIG. 17a** depicts a set of notification processing Nodes 1710 discovering 1715 a Node 1720 that supports the notification handler service. As part of its service description, node 1720 designates its notification interests, as well as which notification payload types are acceptable.

[0378] **FIG. 17b** depicts how notifications can be delivered. Any Node could be the originating source as well as processor of the notification, and could be responsible for delivering the notification to Node 1720, which supports the notification handler service. Thus, Node 1710a could be the originating notification processing Node; or such functionality might be split between Node 1710c (originating source of notification) and Node 1710b (processor of notification). Still another Node (not shown) might be responsible for delivery of the notification. Notification processors that choose to handle notifications from foreign notification-originating Nodes may integrate with a commercial notification-processing engine such as Microsoft Notification Services in order to improve efficiency.

4.3. Service Discovery

[0379] In order to use NEMO services, NEMO Nodes will need to first know about them. One embodiment of NEMO supports three dynamic discovery mechanisms, illustrated in **FIGS. 18a-c**:

[0380] *Client Driven* – a NEMO Node 1810a (in **FIG. 18a**) explicitly sends out a request to some set of targeted Nodes (e.g., 1820a) that support a "Service Query" service interface 1815a, the request asking whether the targeted Nodes support a specified set of services. If requesting Node1810a is authorized, Service Providing Node 1820a will send a response indicating whether it supports the requested interfaces and the associated service interface bindings. This is one of the more common interfaces that Nodes will support if they expose any services.

[0381] *Node Registration* – a NEMO Node 1810b (in **FIG. 18b**) can register its description, including its supported services, with other Nodes, such as Service Providing Node 1820b. If a Node supports this interface 1815b, it is willing to accept requests from other Nodes and then cache those descriptions based on some policy. These Node descriptions are then available directly for use by the receiving Node or by other Nodes that perform service queries targeted to Nodes that have cached descriptions. As an alternative to P2P registration, a Node could also utilize a public registry, such as a UDDI (Universal Discovery, Description and Integration) standard registry for locating services.

[0382] *Event-Based* – Nodes (such as Node 1810c in **FIG. 18c**) send out notifications 1815c to Interested Nodes 1820c (that are "notification aware" and previously indicated their interest), indicating a change in state (e.g., Node active/available), or a Node advertises that it supports some specific service. The notification 1815c can contain a full description of the node and its services, or just

the ID of the node associated with the event. Interested nodes may then choose to accept and process the notification.

4.4. Service Authorization and the Establishment of Trust

[0383] As noted above, in one embodiment, before a NEMO Node allows access to a requested service, it first determines whether, and under which conditions, the requesting Node is permitted access to that service. Access permission is based on a trust context for interactions between service requestor and service provider. As will be discussed below, even if a Node establishes that it can be trusted, a service providing Node may also require that it satisfy a specified policy before permitting access to a particular service or set of services.

[0384] In one embodiment NEMO does not mandate the specific requirements, criteria, or decision-making logic employed by an arbitrary set of Nodes in determining whether to trust each other. Trust semantics may vary radically from Node to Node. Instead, NEMO provides a standard set of facilities that allow Nodes to negotiate a mutually acceptable trusted relationship. In the determination and establishment of trust between Nodes, NEMO supports the exchange of credentials (and/or related information) between Nodes, which can be used for establishing a trusted context. Such trust-related credentials may be exchanged using a variety of different models, including the following:

[0385] *Service-Binding Properties* – a model where trust credentials are exchanged implicitly as part of the service interface binding. For example, if a Node 1920a (in **FIG. 19a**) exposes a service in the form of an HTTP Post over SSL, or as a Web Service that requires a WS-Security XML Signature, then the actual properties

of this service binding may communicate all necessary trust-related credentials 1915a with a Requesting Node 1910a.

[0386] *Request/Response Attributes* – a model where trust credentials are exchanged through WSDL request and response messages (see **FIG. 19b**) between a Requesting Node 1910b and a Service Providing Node 1920b, optionally including the credentials as attributes of the messages 1915b. For example, digital certificates could be attached to, and flow along with, request and response messages, and could be used for forming a trusted relationship.

[0387] *Explicit Exchange* – a model where trust credentials are exchanged explicitly through a service-provider interface (1915c in **FIG. 19c**) that allows querying of information related to the trust credentials that a given node contains. This is generally the most involved model, typically requiring a separate roundtrip session in order to exchange credentials between a Requesting Node 1910c and a Service Providing Node 1920c. The service interface binding itself provides a mutually acceptable trusted channel for explicit exchange of credentials.

[0388] In addition to these basic models, NEMO can also support combinations of these different approaches. For example, the communication channel associated with a semi-trusted service binding may be used to bootstrap the exchange of other security-related credentials more directly, or exchanging security-related credentials (which may have some type of inherent integrity) directly and using them to establish a secure communication channel associated with some service interface binding.

[0389] As noted above, trust model semantics and the processes of establishing trust may vary from entity to entity. In some situations, mutual trust

between nodes may not be required. This type of dynamic heterogeneous environment calls for a flexible model that provides a common set of facilities that allow different entities to negotiate context-sensitive trust semantics.

4.5. Policy-Managed Access

[0390] In one embodiment (as noted above), a service providing Node, in addition to requiring the establishment of a trusted context before it allows a requesting Node to access a resource, may also require that the requesting Node satisfy a policy associated with that resource. The policy decision mechanism used for this purpose may be local and/or private. In one embodiment, NEMO provides a consistent, flexible mechanism for supporting this functionality.

[0391] As part of the service description, one can designate specific NEMO Nodes as "authorization" service providers. In one embodiment an authorization service providing Node implements a standard service for handling and responding to authorization query requests. Before access is allowed to a service interface, the targeted service provider dispatches an "Authorization" query request to any authorizing Nodes for its service, and access will be allowed only if one or more such Nodes (or a pre-specified combination thereof) respond indicating that access is permitted.

[0392] As illustrated in **FIG. 20**, a Requesting Node 2010 exchanges messages 2015 with a Service Providing Node 2020, including an initial request for a particular service. Service Providing Node 2020 then determines whether Requesting Node 2010 is authorized to invoke that service, and thus exchanges authorization messages 2025 with the authorizing Nodes 2025 that manage access to the requested service, including an initial authorization request to these Nodes 2030. Based on the

responses it receives, Service Providing Node 2020 then either processes and returns the applicable service response, or returns a response indicating that access was denied.

[0393] Thus, the Authorization service allows a NEMO Node to participate in the role of policy decision point (PDP). In a preferred embodiment, NEMO is policy management system neutral; it does not mandate how an authorizing Node reaches decisions about authorizations based on an authorization query. Yet, for interoperability, it is preferable that authorization requests and responses adhere to some standard, and be sufficiently extensible to carry a flexible payload so that they can accommodate different types of authorization query requests in the context of different policy management systems. In one embodiment, support is provided for at least two authorization formats: (1) a simple format providing a very simple envelope using some least common denominator criteria, such as input, a simple requestor ID, resource ID, and/or action ID, and (2) the standard "Security Assertion Markup Language" (SAML) format to envelope an authorization query.

[0394] In one embodiment, an authorizing Node must recognize and support at least a predefined "simple" format and be able to map it to whatever native policy expression format exists on the authorizing Node. For other formats, the authorizing Node returns an appropriate error response if it does not handle or understand the payload of an "Authorization" query request. Extensions may include the ability for Nodes to negotiate over acceptable formats of an authorization query, and for Nodes to query to determine which formats are supported by a given authorizing service provider Node.

4.6. Basic DRM Node Interaction

[0395] Returning to the specific NEMO instance of a DRM application, **FIG. 21** is a DRM Node (or Vertex) Graph that can serve to illustrate the interaction among DRM Nodes, as well as their relationships. Consider the following scenario in which portable device 2110 is a content playback device (e.g., an iPod1). Nip1 is the Node that represents this device. Kip1 is the content encryption key associated with Nip1. "User" is the owner of the portable device, and Ng is the Node that represents the user. Kg is the content encryption key associated with Ng.

[0396] PubLib is a Public Library. Npl represents the members of this library, and Kpl is the content encryption key associated with Npl. ACME represents all the ACME-manufactured Music Players. Namp represents that class of devices, and Kamp is the content encryption key associated with this group.

[0397] L1 is a link from Nip1 to Ng, which means that the portable device belongs to the user (and has access to the user's keys). L2 is a link from Ng to Npl, which means that the user is a member of the Public Library (and has access to its keys). L3 is a link from Nip1 to Namp, which means that the portable device is an ACME device (mere membership, as the company has no keys). L4 is a link from Npl to Napl, which is the Node representing all public libraries (and has access to the groupwide keys).

[0398] C1 is a movie file that the Public Library makes available to its members. Kc1 is a key used to encrypt C1. GB[C1] (not shown) is the governance information for C1 (e.g., rules and associated information used for governing access to the content). E(a,b) means 'b' encrypted with key 'a'.

[0399] For purposes of illustration, assume that it is desired to set a rule that a device can play the content C1 as long as (a) the device belongs to someone who is a member of the library and (b) the device is manufactured by ACME.

[0400] The content C1 is encrypted with Kc1. The rules program is created, as well as the encrypted content key RK[C1] = E(Kamp, E(Kpl, Kc1)). Both the rules program and RK[C1] can be included in the governance block for the content, GB[C1].

[0401] The portable device receives C1 and GB[C1]. For example, both might be packaged in the same file, or received separately. The portable device received L1 when the user first installed his device after buying it. The portable device received L2 when the user paid his subscription fee to the Public Library. The portable device received L3 when it was manufactured (e.g., L3 was built in).

[0402] From L1, L2 and L3, the portable device is able to check that Nip1 has a graph path to Ng (L1), Npl (L1+L2), and Namp (L3). The portable device wants to play C1. The portable device runs the rule found in GB[C1]. The rule can check that Nip1 is indeed an ACME device (path to Namp) and belongs to a member of the public library (path to Npl). Thus, the rule returns "yes", and the ordered list (Namp, Npl).

[0403] The portable device uses L1 to compute Kg, and then L2 to compute Kpl from Kg. The portable device also uses L3 to compute Kamp. The portable device applies Kpl and Kamp to RK[C1], found in GB[C1], and computes Kc1. It then uses Kc1 to decrypt and play C1.

[0404] When Node keys are symmetric keys, as in the previous examples, the content packager needs to have access to the keys of the Nodes to which it wishes to

"bind" the content. This can be achieved by creating a Node that represents the packager, and a link between that Node and the Nodes to which it wishes to bind rules. This can also be achieved "out of band" through a service, for instance. But in some situations, it may not be possible, or practical to use symmetric keys. In that case, it is possible to assign a key pair to the Nodes to which a binding is needed without shared knowledge. In that case, the packager would bind a content key to a Node by encrypting the content key with the target Node's public key. To obtain the key for decryption, the client would have access to the Node's private key via a link to that Node.

[0405] In the most general case, the Nodes used for the rules and the Nodes used for computing content encryption keys need not be the same. It is natural to use the same Nodes, since there is a strong relationship between a rule that governs content and the key used to encrypt it, but it is not necessary. In some systems, some Nodes may be used for content protection keys that are not used for expressing membership conditions, and vice versa, and in some situations, two different graphs of Nodes can be used, one for the rules and one for content protection. For example, a rule could say that all members of group Npl can have access to content C1, but the content key Kc1 may not be protected by Kpl, but may instead by protected by the node key Kapl of node Napl, which represents all public libraries, not just Npl. Or a rule could say that you need to be a member of Namp, but the content encryption key could be bound only to Npl.

4.7. Operation of DRM Virtual Machine (VM)

[0406] The discussion with respect to **FIG. 21** above described the operation of a DRM system at a high (Node and Link) level, including the formation and

enforcement of content governance policies. **FIG. 22** depicts an exemplary code module 2200 of a DRM engine's VM that implements the formation and enforcement of such content governance policies.

[0407] Four main elements of illustrative Code Module 2200, shown in **FIG. 22**, include:

[0408] *pkCM Atom:* The pkCM Atom 2210 is the top-level Code Module Atom. It contains a sequence of sub-atoms.

[0409] *pkDS Atom:* The pkDS Atom 2220 contains a memory image that can be loaded into the Data Segment. The payload of the Atom is a raw sequence of octet values.

[0410] *pkCS Atom:* The pkCS Atom 2230 contains a memory image that can be loaded into the Code Segment. The payload of the Atom is a raw sequence of octet values.

[0411] *pkEX Atom:* The pkEX Atom 2240 contains a list of export entries. Each export entry consists of a name, encoded as an 8-bit name size, followed by the characters of the name, including a terminating 0, followed by a 32-bit integer representing the byte offset of the named entry point (this is an offset from the start of the data stored in the pkCS Atom).

4.7.1. *Module Loader*

[0412] In one embodiment, the Control VM is responsible for loading Code Modules. When a Code Module is loaded, the memory image encoded in pkDS Atom 2220 is loaded at a memory address in the Data Segment. That address is chosen by the VM Loader, and is stored in the DS pseudo-register. The memory image encoded

in the pkCS Atom 2230 is loaded at a memory address in the Code Segment. That address is chosen by the VM Loader, and is stored in the CS pseudo-register.

4.7.2. *System Calls*

[0413] In one embodiment, Control VM Programs can call functions implemented outside of their Code Module's Code Segment. This is done through the use of the `OP_CALL` instruction, that takes an integer stack operand specifying the System Call Number to call. Depending on the System Call, the implementation can be a Control VM Byte Code routine in a different Code Module (for instance, a library of utility functions), directly by the VM in the VM's native implementation format, or delegated to an external software module, such as the VM's Host Environment.

[0414] In one embodiment, several System Call Numbers are specified:

[0415] *SYS_NOP = 0*: This call is a no-operation call. It just returns (does nothing else). It is used primarily for testing the VM.

[0416] *SYS_DEBUG_PRINT = 1*: Prints a string of text to a debug output. This call expects a single stack argument, specifying the address of the memory location containing the null-terminated string to print.

[0417] *SYS_FIND_SYSCALL_BY_NAME = 2:* Determines whether the VM implements a named System Call. If it does, the System Call Number is returned on the stack; otherwise the value -1 is returned. This call expects a single stack argument, specifying the address of the memory location containing the null-terminated System Call name that is being requested.

4.7.3. *System Call Numbers Allocation*

[0418] In one embodiment, the Control VM reserves System Call Numbers 0 to 1023 for mandatory System Calls (System Calls that have to be implemented by all profiles of the VM).

[0419] System Call Numbers 16384 to 32767 are available for the VM to assign dynamically (for example, the System Call Numbers returned by SYS_FIND_SYSCALL_BY_NAME can be allocated dynamically by the VM, and do not have to be the same numbers on all VM implementations).

4.7.4. *Standard System Calls*

[0420] In one embodiment, several standard System Calls are provided to facilitate writing Control Programs. Such standard system calls may include a call to obtain a time stamp from the host, a call to determine if a node is Reachable, and/or the like. System calls preferably have dynamically determined numbers (e.g., their System Call Number can be retrieved by calling the SYS_FIND_SYSCALL_BY_NAME System Call with their name passed as the argument).

4.8. Interfaces Between DRM Engine Interface and Host Application

[0421] Following are some exemplary high level descriptions of the types of interfaces provided by an illustrative DRM (client consumption) engine to a Host Application:

[0422] *SystemName::CreateSession(hostContextObject)* → *Session*

Creates a session given a Host Application Context. The context object is used by the DRM engine to make callbacks into the application.

[0423] *Session::ProcessObject(drmObject)*

This function should be called by the Host Application when it encounters certain types of objects in the media files that can be identified as belonging to the DRM subsystem. Such objects include content control programs, membership tokens, etc. The syntax and semantics of those objects is opaque to the Host Application.

[0424] *Session::OpenContent(contentReference) → Content*

The host application calls this function when it needs to interact with a multimedia content file. The DRM engine returns a Content object that can be used subsequently for retrieving DRM information about the content, and interacting with such information.

[0425] *Content::GetDrmInfo()*

Returns DRM metadata about the content that is otherwise not available in the regular metadata for the file.

[0426] *Content::CreateAction(actionInfo) → Action*

The Host Application calls this function when it wants to interact with a Content object. The actionInfo parameter specifies what type of action the application needs to perform (e.g., Play), as well as any associated parameters, if necessary. The function returns an Action object that can then be used to perform the action and retrieve the content key.

[0427] *Action::GetKeyInfo()*

Returns information that is necessary for the decryption subsystem to decrypt the content.

[0428] *Action::Check()*

Checks whether the DRM subsystem will authorize the performance of this action (i.e., whether Action::Perform() would succeed).

[0429] *Action::Perform()*

Performs the action, and carries out any consequences (with their side effects) as specified by the rule(s) that governs this action.

[0430] Following are some exemplary high level descriptions of the type of interface provided by an illustrative Host Application to a DRM (client consumption) engine:

[0431] *HostContext::GetFileSystem(type) → FileSystem*

Returns a virtual FileSystem object to which the DRM subsystem has exclusive access. This virtual FileSystem will be used to store DRM state information. The data within this FileSystem is readable and writeable only by the DRM subsystem.

[0432] *HostContext::GetCurrentTime()*

Returns the current date/time as maintainted by the host system.

[0433] *HostContext::GetIdentity()*

Returns the unique ID of this host.

[0434] *HostContext::ProcessObject(dataObject)*

Gives back to the host services a data object that may have been embedded inside a DRM object, but that the DRM subsystem has identified as being managed by the host (e.g., certificates).

[0435] *HostContext::VerifySignature(signatureInfo)*

Checks the validity of a digital signature to a data object. Preferably, the signatureInfo object contains information equivalent to the information found

in an XMLSig element. The Host Services are responsible for managing the keys and key certificates necessary to validate the signature.

[0436] *HostContext::CreateCipher(cipherType, keyInfo) → Cipher*

Creates a Cipher object that the DRM subsystem can use to encrypt and decrypt data. A minimal set of cipher types will preferably be defined, and for each a format for describing the key info required by the cipher implementation.

[0437] *Cipher::Encrypt(data)*

The Cipher object referred to above, used to encrypt data.

[0438] *Cipher::Decrypt(data)*

The Cipher object referred to above, used to decrypt data.

[0439] *HostContext::CreateDigester(digesterType) → Digester*

Creates a Digester object that the DRM subsystem can use to compute a secure hash over some data. A minimal set of digest types will be defined.

[0440] Digester::Update(data)

The Digester object referred to above, used to compute the secure hash.

[0441] *Digester::GetDigest()*

The Digester object referred to above, used to obtain the secure hash computed by the DRM subsystem.

[0442] Following are some exemplary high level descriptions of the type of interface provided by an illustrative DRM (service-side packaging) engine to a Host Application:

[0443] *SystemName::CreateSession(hostContextObject) → Session*

Creates a session given a Host Application Context. The context object is used by the DRM Packaging engine to make callbacks into the application.

[0444] *Session::CreateContent(contentReferences[]) → Content*

The Host Application will call this function in order to create a Content object that will be associated with license objects in subsequent steps. Having more than one content reference in the contentReferences array implies that these are bound together in a bundle (one audio and one video track for example), and that the license issued should be targeted to these as one indivisible group.

[0445] *Content::SetDrmInfo(drmInfo)*

The drmInfo parameter specifies the metadata of the license that will be issued. The structure will be read and will act as a guideline to compute the license into bytecode for the VM.

[0446] *Content::GetDRMObjects(format) → drmObjects*

This function is called when the Host Application is ready to get the drmObjects that the DRM Packaging engine created. The format parameter will indicate the format expected for these objects (e.g., XML or binary atoms).

[0447] *Content::GetKeys() → keys[]*

This function is called by the Host Application when it needs the keys in order to encrypt the content. In one embodiment there will be one key per content reference.

[0448] Following are some exemplary high level descriptions of the type of interface provided by an illustrative Host Application to a DRM (service-side packaging) engine:

[0449] *HostContext::GetFileSystem(type) → FileSystem*

Returns a virtual FileSystem object to which the DRM subsystem has exclusive access. This virtual FileSystem would be used to store DRM state information.

The data within this FileSystem should only be readable and writeable by the DRM subsystem.

[0450] *HostContext::GetCurrentTime() → Time*

Returns the current date/time as maintainted by the host system.

[0451] *HostContext::GetIdentity() → ID*

Returns the unique ID of this host.

[0452] *HostContext::PerformSignature(signatureInfo, data)*

Some DRM objects created by the DRM Packaging engine will have to be trusted. This service, provided by the host, will be used to sign the specified object.

[0453] *HostContext::CreateCipher(cipherType, keyInfo) → Cipher*

Creates a Cipher object that the DRM Packaging engine can use to encrypt and decrypt data. This is used to encrypt the content key data in the ContentKey object.

[0454] *Cipher::Encrypt(data)*

The Cipher object referred to above, used to encrypt data.

[0455] *Cipher::Decrypt(data)*

The Cipher object referred to above, used to decrypt data.

[0456] *HostContext::CreateDigester(digesterType) → Digester*

Creates a Digester object that the DRM Packaging engine can use to compute a secure hash over some data.

[0457] *Digester::Update(data)*

The Digester object referred to above, used to comput the secure hash.

[0458] *Digester::GetDigest()*

The Digester object referred to above, used to obtain the secure hash computed by the DRM subsystem.

[0459] *HostContext::GenerateRandomNumber()*

Generates a random number that can be used for generating a key.

5. SERVICES

5.1. Overview

[0460] Having described the NEMO/DRM system from both an architectural and operational perspective, we now turn our attention to an illustrative collection of services, data types, and related objects ("profiles") that can be used to implement the functionality of the system.

[0461] As noted above, a preferred embodiment of the NEMO architecture employs a flexible and portable way of describing the syntax of requests and responses associated with service invocation, data types used within the framework, message enveloping, and data values exposed by and used within the NEMO framework. WSDL 1.1 and above provides sufficient flexibility to describe and represent a variety of types of service interface and invocation patterns, and has sufficient abstraction to accommodate bindings to a variety of different endpoint Nodes via diverse communication protocols.

[0462] In one embodiment, we define a profile to be a set of thematically related data types and interfaces defined in WSDL. NEMO distinguishes a "Core" profile (which includes the foundational set of data types and service messages necessary to support fundamental NEMO Node interaction patterns and infrastructural functionality) from an application-specific profile, such as a DRM

Profile (which describes the Digital Rights Management services that can be realized with NEMO), both of which are discussed below.

[0463] It will be appreciated that many of the data types and services defined in these profiles are abstract, and should be specialized before they are used. Other profiles are built on top of the Core profile.

5.2. NEMO Profile Hierarchy

[0464] In one embodiment, the definition of service interfaces and related data types is structured as a set of mandatory and optional profiles that build on one another and may be extended. The difference between a profile and a profile extension is a relatively subtle one. In general, profile extensions don't add new data types or service type definitions. They just extend existing abstract and concrete types.

[0465] **FIG. 23** illustrates an exemplary profile hierarchy for NEMO and DRM functionality. The main elements of this profile hierarchy include:

[0466] *Core Profile* – At the base of this profile hierarchy lies Core Profile 2300, which preferably shares both NEMO and DRM functionality. This is the profile on which all other profiles are based. It includes a basic set of generic types (discussed below) that serve as the basis for creating more complex types in the framework. Many of the types in the Core Profile are abstract and will need to be specialized before use.

[0467] *Core Profile Extensions* – Immediately above Core Profile 2300 are the Core Profile Extensions 2320, which are the primary specializations of the types in Core Profile 2300, resulting in concrete representations.

[0468] *Core Services Profile* – Also immediately above Core Profile 2300, the Core Services Profile 2310 defines a set of general infrastructure services, also discussed below. In this profile, the service definitions are abstract and will need to be specialized before use.

[0469] *Core Services Profile Extensions* – Building upon both Core Profile Extensions 2320 and Core Services Profile 2310 are the Core Services Profile Extensions 2330, which are the primary specializations of the services defined in Core Services Profile 2310, resulting in concrete representations.

[0470] *DRM Profile* – Immediately above Core Profile 2300 lies DRM Profile 2340, upon which other DRM-related profiles are based. DRM Profile 2340 includes a basic set of generic types (discussed below) that serve as the basis for creating more complex DRM-specific types. Many of the types in DRM Profile 2340 are abstract and will need to be specialized before use.

[0471] *DRM Profile Extensions* – Building upon DRM Profile 2340 are the DRM Profile Extensions 2350, which are the primary specializations of the types in DRM Profile 2340, resulting in concrete representations.

[0472] *DRM Services Profile* – Also building upon DRM Profile 2340 is DRM Services Profile 2360, which defines a set of general DRM services (discussed below). In this profile, the service definitions are abstract and need to be specialized before use.

[0473] *Specific DRM Profile* – Building upon both DRM Profile Extensions 2350 and DRM Services Profile 2360 is the Specific DRM Profile 2370, which is a further specialization of the DRM services defined in DRM Services Profile 2360.

This profile also introduces some new types and further extends certain types specified in Core Profile Extensions 2320.

5.3. NEMO Services and Service Specifications

[0474] Within this profile hierarchy lies, in one embodiment, the following main service constructs (as described in more detail above):

[0475] *Peer Discovery* - the ability to have peers in the system discover one another.

[0476] *Service Discovery* - the ability to discover and obtain information about services offered by different peers.

[0477] *Authorization* - the ability to determine if a given peer (e.g., a Node) is authorized to access a given resource (such as a service).

[0478] *Notification* - services related to the delivery of targeted messages, based on specified criteria, to a given set of peers (e.g., Nodes).

[0479] Following are definitions (also discussed above) of some of the main DRM constructs within this example profile hierarchy:

[0480] *Personalization* - services to obtain the credentials, policy, and other objects needed for a DRM-related endpoint (such as a CE device, music player, DRM license server, etc.) to establish a valid identity in the context of a specific DRM system.

[0481] *Licensing Acquisition* - services to obtain new DRM licenses.

[0482] *Licensing Translation* - services to exchange one new DRM license format for another.

[0483] *Membership* - services to obtain various types of objects that establish membership within some designated domain.

[0484] The NEMO/DRM profile hierarchy can be described, in one embodiment, as a set of Generic Interface Specifications (describing an abstract set of services, communication patterns, and operations), Type Specifications (containing the data types defined in the NEMO profiles), and Concrete Specifications (mapping abstract service interfaces to concrete ones including bindings to specific protocols). One embodiment of these specifications, in the form of Service Definitions and Profile Schemas, is set forth in Appendix C hereto.

6. ADDITIONAL APPLICATION SCENARIOS

[0485] **FIG. 24** illustrates a relatively simple example of an embodiment of NEMO in operation in the context of a consumer using a new music player to play a DRM-protected song. As shown below, however, even this simple example illustrates many different potential related application scenarios. This example demonstrates the bridging of discovery services – using universal plug and play (UPnP) based service discovery as a mechanism to find and link to a UDDI based directory service. It also details service interactions between Personal Area Network (PAN) and Wide Area Network (WAN) services, negotiation of a trusted context for service use, and provisioning of a new device and DRM service.

[0486] Referring to **FIG. 24**, a consumer, having bought a new music player 2400, desires to play a DRM-protected song. Player 2400 can support this DRM system, but needs to be personalized. In other words, Player 2400, although it includes certain elements (not shown) that render it both NEMO-enabled and DRM-capable, must first perform a personalization process to become part of this system.

[0487] Typically, a NEMO client would include certain basic elements illustrated in **FIGS. 5a** and **6** above, such as a Service Access Point to invoke other

Node's services, Trust Management Processing to demonstrate that it is a trusted resource for playing certain protected content, as well as a Web Services layer to support service invocations and the creation and receipt of messages. As discussed below, however, not all of these elements are necessary to enable a Node to participate in a NEMO system.

[0488] In some embodiments, client nodes may also include certain basic DRM-related elements, as illustrated in **FIG 12a** and **13-15** above, such as a DRM client engine and cryptographic services (and related objects and cryptographic keys) to enable processing of protected content, including decrypting protected songs, as well as a media rendering engine to play those songs. Here, too, some such elements need not be present. For example, had Player 2400 been a music player that was only capable of playing unprotected content, it might not require the core cryptographic elements present in other music players.

[0489] More specifically, in the example shown in **FIG. 24**, Player 2400 is wireless, supports the UPnP and Bluetooth protocols, and has a set of X.509 certificates it can use to validate signatures and sign messages. Player 2400 is NEMO-enabled in that it can form and process a limited number of NEMO service messages, but it does not contain a NEMO Service Access Point due to resource constraints.

[0490] Player 2400, however, is able to participate in a Personal Area Network (PAN) 2410 in the user's home, which includes a NEMO-enabled, Internet-connected, Home Gateway Device 2420 with Bluetooth and a NEMO SAP 2430. The UPnP stacks of both Player 2400 and Gateway 2420 have been extended to support a new service profile type for a "NEMO-enabled Gateway" service, discussed below.

[0491] When the user downloads a song and tries to play it, Player 2400 determines that it needs to be personalized, and initiates the process. For example, Player 2400 may initiate a UPnP service request for a NEMO gateway on PAN 2410. It locates a NEMO gateway service, and Gateway 2420 returns the necessary information to allow Player 2400 to connect to that service.

[0492] Player 2400 then forms a NEMO Personalization request message and sends it to the gateway service. The request includes an X.509 certificate associated with Player 2400's device identity. Gateway 2420, upon receiving the request, determines that it cannot fulfill the request locally, but has the ability to discover other potential service providers. However, Gateway 2420 has a policy that all messages it receives must be digitally signed, and thus it rejects the request and returns an authorization failure stating the policy associated with processing this type of request.

[0493] Player 2400, upon receiving this rejection, notes the reason for the denial of service and then digitally signs (e.g., as discussed above in connection with **FIG. 15**) and re-submits the request to Gateway 2420, which then accepts the message. As previously mentioned, Gateway 2420 cannot fulfill this request locally, but can perform service discovery. Gateway 2420 is unaware of the specific discovery protocols its SAP 2430 implementation supports, and thus composes a general attribute-based service discovery request based on the type of service desired (personalization), and dispatches the request via SAP 2430.

[0494] SAP 2430, configured with the necessary information to talk to UDDI registries, such as Internet-Based UDDI Registry 2440, converts the request into a native UDDI query of the appropriate form and sends the query. UDDI Registry 2440 knows of a service provider that supports DRM personalization and returns the query

results. SAP 2430 receives these results and returns an appropriate response, with the necessary service provider information, in the proper format, to Gateway 2420.

[0495] Gateway 2420 extracts the service provider information from the service discovery response and composes a new request for Personalization based on the initial request on behalf of Player 2400. This request is submitted to SAP 2430. The service provider information (in particular, the service interface description of Personalization Service 2450) reveals how SAP 2430 must communicate with a personalization service that exposes its service through a web service described in WSDL. SAP 2430, adhering to these requirements, invokes Personalization Service 2450 and receives the response.

[0496] Gateway 2420 then returns the response to Player 2400, which can use the payload of the response to personalize its DRM engine. At this point, Player 2400 is provisioned, and can fully participate in a variety of local and global consumer oriented services. These can provide full visibility into and access to a variety of local and remote content services, lookup, matching and licensing services, and additional automated provisioning services, all cooperating in the service of the consumer. As explained above, various decryption keys may be necessary to access certain protected content, assuming the consumer and Player 2400 satisfy whatever policies are imposed by the content provider.

[0497] Thus, a consumer using a personal media player at home can enjoy the simplicity of a CE device, but leverage the services provided by both gateway and peer devices. When the consumer travels to another venue, the device can rediscover and use most or all of the services available at home, and, through new gateway services, be logically connected to the home network, while enjoying the services

available at the new venue that are permitted according to the various policies associated with those services. Conversely, the consumer's device can provide services to peers found at the new venue.

[0498] Clearly, utilizing some or all of these same constructs (NEMO Nodes, SAPs, Service Adaptation Layers, various standards such as XML, WSDL, SOAP, UDDI, etc.), many other scenarios are possible, even within the realm of this DRM music player example. For example, Player 2400 might have contained its own SAP, perhaps eliminating the need for Gateway 2420. UDDI Registry 2440 might have been used for other purposes, such as locating and/or licensing music content. Moreover, many other DRM applications could be constructed, e.g., involving a licensing scheme imposing complex usage and distribution policies for many different types of audio and video, for a variety of different categories of users. Also, outside of the DRM context, virtually any other service-based applications could be constructed using the NEMO framework.

[0499] As another example, consider the application of NEMO in a business peer-to-peer environment. Techniques for business application development and integration are quickly evolving beyond the limits of traditional tools and software development lifecycles as practiced in most IT departments. This includes the development of word processing documents, graphic presentations, and spreadsheets. While some would debate whether these documents in their simplest form represent true applications, consider that many forms of these documents have well defined and complex object models that are formally described. Such documents or other objects might include, for example, state information that can be inspected and updated during the lifecycle of the object, the ability for multiple users to work on the objects

concurrently, and/or additional arbitrary functionality. In more complicated scenarios, document-based information objects can be programmatically assembled to behave like full-fledged applications.

[0500] Just as with traditional software development, these types of objects can also benefit from source control and accountability. There are many systems today that support document management, and many applications directly support some form of document control. However most of these systems in the context of distributed processing environments exhibit limitations, including a centralized approach to version management with explicit check-in and checkout models, and inflexible (very weak or very rigid) coherence policies that are tied to client rendering applications or formats particularly within the context of a particular application (e.g., a document).

[0501] Preferred embodiments of NEMO can be used to address these limitation by means of a P2P policy architecture that stresses capability discovery and format negotiation. It is possible to structure the creation of an application (e.g., a document) in richer ways, providing multiple advantages. Rich policy can be applied to the objects and to the structure of the application. For example, a policy might specify some or all of the following:

- Only certain modules can be modified.
- Only object interfaces can be extended or implementations changed.
- Deletions only allowed but not extensions.
- How updates are to be applied, including functionality such as automatic merging of non-conflicting updates, and application of updates before a given peer can send any of its updates to other peers.

- Policy-based notification such that all peers can be notified of updates if they choose, in order to participate in direct synchronization via the most appropriate mechanisms.
- Support updates from different types of clients based on their capabilities.

[0502] In order to achieve this functionality, the authoring application used by each participant can be a NEMO-enabled peer. For the document that is created, a template can be used that describes the policy, including who is authorized and what can be done to each part of the document (in addition to the document's normal formatting rules). As long as the policy engine used by the NEMO peer can interpret and enforce policy rules consistent with their semantics, and as long as the operations supported by the peer interfaces allowed in the creation of the document can be mapped to a given peer's environment via the Service Adaptation Layer, then any peer can participate, but may internally represent the document differently.

[0503] Consider the case of a system consisting of different NEMO peers using services built on the NEMO framework for collaboration involving a presentation document. In this example, a wireless PDA application is running an application written in Java, which it uses for processing and rendering the document as text. A different implementation running under Microsoft Windows® on a desktop workstation processes the same document using the Microsoft Word® format. Both the PDA and the workstation are able to communicate, for example, by connection over a local area network, thus enabling the user of the PDA and the user of the workstation to collaborate on the same document application. In this example:

- NEMO peers involved in the collaboration can discover each other, their current status, and their capabilities.

- Each NEMO peer submits for each committable change, its identity, the change, and the operation (e.g., deletion, extension, etc.).
- All changes are propagated to each NEMO peer. This is possible because each NEMO peer can discover the profile and capabilities of another peer if advertised. At this point the notifying peer can have the content change encoding in a form acceptable by the notified peer if it is incapable of doing so. Alternatively the accepting peer may represent the change in any format it sees fit upon receipt at its interface.
- Before accepting a change the peer verifies that it is from an authorized NEMO participant.
- The change is applied based on the document policy.

[0504] As another example, consider the case of a portable wireless consumer electronics (CE) device that is a NEMO-enabled node (X), and that supports DRM format A, but wants to play content in DRM format B. X announces its desire to render the content as well as a description of its characteristics (e.g., what its identity is, what OS it supports, its renewability profile, payment methods it supports, and/or the like) and waits for responses back from other NEMO peers providing potential solutions. In response to its query, X receives three responses:

(1) Peer 1 can provide a low quality downloadable version of content in clear MP3 form for a fee of $2.00.

(2) Peer 2 can provide high quality pay-per-play streams of content over a secure channel for $0.50 per play.

(3) Peer 3 can provide a software update to X that will permit rendering of content in DRM format B for a fee of $10.00.

[0505] After reviewing the offers, X determines that option one is the best choice. It submits a request for content via offer number one. The request includes an assertion for a delegation that allows Peer 1 to deduct $2.00 from X's payment account via another NEMO service. Once X has been charged, then X is given back in a response from Peer 1 a token that allows it to download the MP3 file.

[0506] If instead, X were to decide that option three was the best solution, a somewhat more complicated business transaction might ensue. For example, option three may need to be represented as a transactional business process described using a NEMO Orchestration Descriptor (NOD) implemented by the NEMO Workflow Collator (WFC) elements contained in the participating NEMO enabled peers. In order to accomplish the necessary software update to X, the following actions could be executed using the NEMO framework:

- X obtains permission from its wireless service provider (B) that it is allowed to receive the update.
- Wireless service provider B directly validates peer three's credentials in order to establish its identity.
- X downloads from B a mandatory update that allows it to install 3rd party updates, their is no policy restriction on this, but this scenario is the first triggering event to cause this action.
- X is charged for the update that peer three provides.
- X downloads the update from peer three.

[0507] In this business process some actions may be able to be carried out concurrently by the WFC elements, while other activities may need to authorized and executed in a specific sequence.

[0508] Yet another example of a potential application of the NEMO framework is in the context of online gaming. Many popular multiplayer gaming environment networks are structured as centralized, closed portals that allow online gamers to create and participate in gaming sessions.

[0509] One of the limitations of these environments is that the users generally must have a tight relationship with the gaming network and must have an account (usually associated with a particular game title) in order to use the service. The typical gamer must usually manage several game accounts across multiple titles across multiple gaming networks and interact with game-provider-specific client applications in order to organize multiple player games and participate within the networks. This is often inconvenient, and discourages online use.

[0510] Embodiments of the NEMO framework can be used to enhance the online gaming experience by creating an environment that supports a more federated distributed gaming experience, making transparent to the user and the service provider the details of specific online game networks. This not only provides a better user experience, thereby encouraging adoption and use of these services, but can also reduce the administrative burden on game network providers.

[0511] In order to achieve these benefits, gaming clients can be personalized with NEMO modules so that they can participate as NEMO peers. Moreover, gaming networks can be personalized with NEMO modules so that they can expose their administrative interfaces in standardized ways. Finally, NEMO trust management can be used to ensure that only authorized peers interact in intended ways.

[0512] For example, assume there are three gaming network providers A, B, and C, and two users X and Y. User X has an account with A, and User Y has an

account with B. X and Y both acquire a new title that works with C and want to play each other. Using the NEMO framework, X's gaming peer can automatically discover online gaming provider C. X's account information can be transmitted to C from A, after A confirms that C is a legitimate gaming network. X is now registered with C, and can be provisioned with correct tokens to interact with C. User Y goes through the same process to gain access to C using its credentials from B. Once both X and Y are registered they can now discover each other and create an online gaming session.

[0513] This simple registration example can be further expanded to deal with other services that online gaming environments might provide, including, e.g., game token storage (e.g., in lockers), account payment, and shared state information such as historical score boards.

[0514] While several examples were presented in the context of enterprise document management, online gaming, and media content consumption, it will be appreciated that the NEMO framework and the DRM system described herein can be used in any suitable context, and are not limited to these specific examples.

APPENDIX A

SYSTEMS AND METHODS FOR PEER-TO-PEER SERVICE ORCHESTRATION

[0515] This Appendix corresponds to U.S. Provisional Patent Application No. 60/476,357, entitled Systems and Methods for Peer-to-Peer Service Orchestration, by William Bradley and David Maher, filed June 5, 2003 ("the Bradley et al. application"). References to "the Bradley et al. application" throughout the overall PCT application submitted herewith pertain to the information here, in Appendix A.

COPYRIGHT AUTHORIZATION

[0516] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the WIPO or U.S. Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

FIELD OF THE INVENTION

[0517] The present invention relates generally to the distribution and use of digital information and services. More specifically, systems and methods are disclosed for providing and/or supporting peer-to-peer service orchestration.

BACKGROUND AND SUMMARY OF THE INVENTION

[0518] Today, significant barriers exist to the goal of an interoperable and secure world of media-related services. For example, multiple, overlapping de facto and formal standards inhibit straightforward interoperability; it is often difficult to locate and connect to needed services; implementation technologies are often not

interoperable; and there are often impedance mismatches between different trust and protection models

[0519] While emerging web service standards are beginning to address some of these issues on the web, such approaches are incomplete. In addition, such approaches fail to address these issues across multiple tiers of network nodes spanning personal and local area networks, home, enterprise, and department gateways, and wide area networks. Nor do they enable dynamic orchestration of both simple and complex services using a variety of service interface bindings both local and remote (e.g. WS-I, Java RMI, DCOM, C function invocation, .Net, etc.) allowing the integration of many legacy applications.

[0520] Embodiments of the present invention can be used to solve or address some or all of these problems. For example, in some embodiments service providers can use the service discovery protocols most appropriate for the device, service, or application that participates in a given heterogeneous service network. Thus, BlueTooth, UpNp, rendezvous, JINI, UDDI, etc. can be integrated in the same service, and each node can use the discovery service(s) most appropriate for the device that hosts the node.

[0521] In the media world, the systems and interfaces required or favored by the major sets of stakeholders (e.g., content publishers, distributors, retail services, consumer device providers, and consumers) often differ widely. Thus, it would be desirable to unite the capabilities provided by these entities into integrated services that can rapidly evolve into optimal configurations meeting the needs of all participating entities. Embodiments of the present invention can be used to meet this goal by using peer-to-peer ("P2P") service orchestration.

[0522] While the advantages of P2P frameworks have already been seen for such things as music and now video distribution, P2P technology can be used much more extensively. For example, there are opportunities in various enterprise services, and especially in the interactions between two or more enterprises. While enterprises are most often organized hierarchically, and their information systems often reflect that organization, when people in two enterprises interact, they often will do so more effectively through peer interfaces. The receiving person/service in company A can solve problems or get useful information more directly by talking to the shipping person in company B. Traversing hierarchies or unnecessary interfaces is often not useful. Shipping companies (such as FedEx and UPS) realize this and allow direct visibility into their processes, allowing events to be directly monitored by customers. Companies and municipalities are organizing their services through enterprise portals, allowing crude forms of self-service. However, so far, peer-to-peer frameworks do not allow one enterprise to expose its various service interfaces to its customers and suppliers in such a way as to allow those entities to interact at natural peering levels, allowing those entities to orchestrate the enterprise's services in ways that best suit them. Embodiments of the present invention will allow this.

[0523] Systems and methods are disclosed for providing service orchestration in a networked computing environment. It should be appreciated that the disclosed embodiments can be implemented in numerous ways, including as processes, apparatuses, systems, devices, methods, or computer readable media. Several inventive embodiments are described below.

[0524] In one set of embodiments, a services framework is provided that enables various stakeholders in the consumer or enterprise media space (e.g.,

consumers, content providers, device manufacturers, service providers, enterprise departments) to find each other, establish a trusted relationship, and exchange value in rich and dynamic ways through trusted service interfaces. This services framework can be described as a platform for enabling interoperable, secure, media-related ecommerce in a world of heterogeneous consumer devices, media formats, communication protocols and security mechanisms.

[0525] These and other features, advantages, and embodiments of the present invention will be illustrated in more detail by the following detailed description and the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0526] Embodiments of the present invention will be readily understood by referring to the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

[0527] Fig. 1 shows an example embodiment of the MediaDrive framework.

[0528] Fig. 2 is a conceptual view of a MediaDrive node.

[0529] Fig. 3 illustrates a generic interaction pattern between MediaDrive nodes.

[0530] Fig. 4 illustrates an embodiment of the Service Adaptation Layer.

[0531] Fig. 5a shows a service proxy involved in a client-side MSDL interaction.

[0532] Fig. 5b shows a service proxy involved in a client-side native interaction.

[0533] Fig. 6 shows a service proxy involved in a service-side point-to-intermediary interaction pattern.

[0534] Fig. 7 illustrates an interaction pattern of a MediaDrive workflow collator.

[0535] Fig. 8 depicts a set of notification processing nodes discovering a node that supports the notification handler service.

[0536] Fig. 9 illustrates the delivery of notifications in accordance with a preferred embodiment.

[0537] Fig. 10 illustrates a client-driven scenario where a requesting node makes a service query request to a targeted service providing node.

[0538] Fig. 11 illustrates a node registration scenario where a requesting node wants to register its description with another node.

[0539] Fig. 12 illustrates a MediaDrive-enabled node being notified of the existence of a service.

[0540] Fig. 13 illustrates the process of establishing trust based on an explicit exchange.

[0541] Fig. 14 illustrates policy-managed access to a service.

[0542] Fig. 15 illustrates how MediaDrive can be used in the context of delivering CRM-related services to different stakeholders.

DETAILED DESCRIPTION

[0543] A detailed description of the inventive body of work is provided below. While this description is provided in conjunction with several embodiments, it should be understood that the invention is not limited to any one embodiment, but instead encompasses numerous alternatives, modifications, and equivalents. For example, while some embodiments are described in the context of consumer-oriented content and applications, those skilled in the art will recognize that the disclosed systems and methods are readily adaptable for broader application. For example, without limitation, embodiments of the present invention could be readily applied in the context of enterprise content and applications. In addition, while numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention, the present invention may be practiced without some or all of these details. Moreover, for the purpose of clarity, certain technical material that is known in the art related to the invention has not been described in detail in order to avoid unnecessarily obscuring the present invention

[0544] Embodiments of a novel, policy managed, peer-to-peer, service orchestration framework (the MediaDrive framework) are described herein. The framework was originally designed to support the formation of grass-roots, self-organizing service networks that can enable rich media experiences. As one of ordinary skill in the art will appreciate, the framework itself is novel, as are many of its features, aspects, and applications.

[0545] In a preferred embodiment, the framework is designed to enable the dynamic configuration and evolution of many different types of services using a services description language (the MediaDrive Services Description Language

("MSDL")). Services can be distributed across peer-to-peer communicating nodes, each providing message routing and orchestration using a message pump and workflow collator designed specifically for this framework. Distributed policy management of MediaDrive service interfaces can be used to help provide trust and security, thereby facilitating commercial exchange of value.

[0546] Peer-to-peer computing is often defined as the sharing of resources (such as hard drives and processing cycles) among computers and other intelligent devices. *See* http://www.intel.com/cure/peer.htm. Here, P2P may be viewed as a communication model allowing network nodes to symmetrically consume and provide services of all sorts. P2P messaging and workflow collation allow rich services to be dynamically created from a heterogeneous set of more primitive services. This enables examination of the possibilities of P2P computing when the shared resources are services of many different types, even using different service bindings. Embodiments of MediaDrive can be used to provide a media services framework enabling stakeholders (e.g., consumers, content providers, device manufacturers, and service providers) to find one other, to interact, exchange value, and to cooperate in rich and dynamic ways. The different types of services that embodiments of MediaDrive can be used to coordinate range from the basic (discovery, notification, search, and file sharing) to more complex higher level services (such as lockers, licensing, matching, authorization, payment transaction, and update), and combinations of any or all of these.

[0547] Enterprise applications of MediaDrive are particularly intriguing as enterprises move to service-oriented architectures. Through use of peer to peer service orchestration, and distributed policy management, enterprises can allow their

departments to publish service interfaces that can be configured into rich, personalized services by customers and suppliers, breaking down artificial hierarchies, and allowing both internal and external entities to optimize their interaction with the enterprise.

[0548] For ease of explanation, the following nomenclature will be used:

[0549] *Service* – Any well-defined functionality offered by some provider. This could be, for example, a low-level functionality offered within a device such as a cell phone (e.g. a voice recognition service), or a multi-faceted functionality offered over the world-wide web (e.g. a shopping service).

[0550] *Service Interface* – Any well-defined way of interacting with one or more services.

[0551] *Service Binding* – The conventions and protocols used to invoke a service interface. These may be represented in a variety of well-defined ways, such as the WS-I standard XML protocol, RPC based on the WSDL definition, or a function invocation from a DLL.

[0552] *Service Orchestration* – The assembly and coordination of services into manageable, coarser-grained services, reusable components, or full applications that adhere to rules specified by the service provider. Examples include rules based on provider identity, type of service, method in which services are accessed, order in which services are composed, etc.

[0553] *Governance* – The process of exercising authoritative or dominating influence over some item such as a music file, a document, or a service operation such as the ability to download and install a software upgrade. Governance typically interacts with trust management, policy management, and content protection.

[0554] *Peer to Peer (P2P) Philosophy* – A communication model supporting symmetric access to services for participants.

[0555] *MediaDrive Node* – A participant within the MediaDrive Framework. In a preferred embodiment, a node may act in multiple roles including that of a service consumer and/or a service provider. Nodes may be implemented in a variety of forms including consumer electronic devices, software agents such as media players, or virtual service providers such as content search engines, DRM license providers, or content lockers. MediaDrive services may be orchestrated to provide more robust composite services.

[0556] *MediaDrive Service Description Language (MSDL)* – In one embodiment, an XML Schema based language that defines and describes the extensible set of data types and messages that are used for communication between nodes in an embodiment of the MediaDrive framework.

7. Logical Model

[0557] Fig. 1 illustrates a relatively simple instance of the MediaDrive framework. In a preferred embodiment, the MediaDrive framework consists of a logically connected set of nodes that interact in a P2P fashion.

[0558] In a preferred embodiment, some characteristics of this interaction are:

- MediaDrive nodes interact by making service invocation requests and receiving responses. The format and payload of the request and response messages are defined in MSDL. The MediaDrive framework supports the construction of diverse communication patterns, ranging from direct interaction with a single service provider to a complex aggregation of a choreographed set of services from multiple service providers. In a preferred embodiment, the framework supports the basic mechanisms for using existing service choreography standards and also allows service providers to use their own conventions.

- A service interface may have one or more service bindings. In a preferred embodiment, the description of services bindings is expressed in MSDL. In this embodiment, a MediaDrive node may invoke the interface of another node as long as that node's interface binding can be expressed in MSDL, and as long as the requesting node can support the conventions and protocols associated with the binding. For example, if a node supports a web service interface, a requesting node may be required to support SOAP, HTTP, WS-Security, etc.
- Any service interface may be controlled (e.g., rights managed) in a standardized fashion directly providing aspects of rights management. Interactions between MediaDrive nodes can be viewed as governed operations.

[0559] Virtually any type of device (physical or virtual) can be viewed as potentially MediaDrive-enabled, and able to implement key aspects of the MediaDrive framework. Device types include, for example, consumer electronics equipment, networked services, or software clients. In a preferred embodiment, a MediaDrive-enabled device (node) typically includes some or all of the following logical modules:

- Native Services API – the set of one or more services that the device implements. There is no requirement that a MediaDrive node expose any service directly or indirectly in the MediaDrive framework.
- Native Service Implementation – the corresponding set of implementations for the native services API.
- MediaDrive Service Adaptation Layer – the logical layer through which an exposed subset of an entity's native services is accessed using one or more discoverable bindings described in MSDL.
- MediaDrive Framework Support Library – components that provide support functionality for working with the MediaDrive Framework including support for invoking service interfaces, message processing, service orchestration, etc.

[0560] Fig. 2 provides a conceptual view of an exemplary embodiment of a MediaDrive node. The actual design and implementation corresponding to each of the above modules will typically vary from device to device.

8. Basic Interaction Pattern

[0561] Fig. 3 depicts a typical logical interaction pattern between two MediaDrive nodes, a service requester and a service provider.

[0562] From the requesting node's perspective, the flow of events will typically be:

- Make a service discovery request to locate any MediaDrive-enabled nodes that can provide the necessary service using specified service bindings. A node may choose to cache information about discovered services. The interface/mechanism for service discovery between MediaDrive nodes can be just another service that a MediaDrive node chooses to implement.
- Once candidate service providing nodes are found, the requesting node may choose to dispatch a request to one or more of the service providing nodes based on a specific service binding.
- In a preferred embodiment, two nodes that wish to communicate securely with each other will establish a trusted relationship for the purpose of exchanging MSDL messages. For example, they may negotiate a set of compatible trust credentials (e.g., X.500 certificates, device keys, etc) that may be used in determining identity, authorization, establishing a secure channel, etc. In some cases, the negotiation of these credentials may be an implicit property of the service interface binding (e.g., WS-Security if WS-I XML Protocol is used, or an SSL request between two well-known nodes). In some cases, the negotiation of trust credentials may be an explicitly separate step. In a preferred embodiment, it is up to a given node to determine what credentials are sufficient for interacting with another MediaDrive node, and to make the decision that it can trust a given node.

- The requesting node creates the appropriate MSDL request message(s) that correspond to the requested service.
- Once the message is created, it is dispatched to the targeted service providing node(s). The communication style of the request may, for example, be synchronous or asynchronous RPC style, or message-oriented based on the service binding. Dispatching of service requests and receiving of responses may be done directly by the device or through the MediaDrive Service Proxy. The service proxy (described below) provides an abstraction and interface for sending messages to other participants in the MediaDrive framework, and may hide certain service binding issues such as compatible message formats, transport mechanisms, message routing issues, or the like.
- After dispatching a request, the requesting node will typically receive one or more responses. Depending on the specifics of the service interface binding and the requesting node's preferences, the response(s) may be returned in a variety of ways, including, for example, an RPC-style response or a notification message. The response, en-route to the targeted node(s), may pass through other intermediate nodes that may provide a number of MediaDrive services, including, e.g., routing, trust translation, collation and correlation functions, etc.
- The requesting node validates the response(s) to ensure it adheres to the negotiated trust semantics between it and the service providing node.
- Appropriate processing is then applied based on the message payload type and contents.

[0563] Referring once again to Fig. 3, from the service providing node's perspective, the flow of events is:

- Determine if the requested service is supported. In a preferred embodiment, the MediaDrive framework does not mandate the style or granularity of how a service interface maps as an entry point to a service. In the simplest case, a service interface may map unambiguously to a given service and the act of binding to and invoking it may constitute support for the service. However, in some embodiments a single service interface may handle multiple types of requests and

a given service type may contain additional attributes that need to be examined before a determination can be made that the node supports the specifically desired functionality.

- In some cases it may be necessary for the service provider to determine if it trusts the requesting node and to negotiate a set of compatible trust credentials. In a preferred embodiment, regardless of whether the service provider determines trust, any policy associated with the service interface will still apply.
- The service provider determines and dispatches authorization request(s) to those MediaDrive node(s) responsible for authorizing access to the interface in order to determine if the requesting node has access. In many situations, the authorizing node and the service providing node will be the same entity, and the dispatching and processing of the authorization request will be local operations invoked through a lightweight MediaDrive service interface binding such as a C function entry point.
- Upon receiving the authorization response, if the requesting node is authorized, the service provider will fulfill the request. If not, an appropriate response message can be generated.
- The response message is returned based on the service interface binding and requesting node's preferences. En route to the requesting node, the message may pass through other intermediate nodes that may provide necessary or "value added" services. For example an intermediate node might provide routing, trust translation, or delivery to a notification processing node that can deliver the message in a way acceptable to the requesting node. An example of a "value added" service is a coupon service that appends coupons to the message if it knows of the requesting node's interests.

9. MSDL

[0564] The syntax of messages associated with service invocation are preferably described in a relatively flexible and portable manner, as are the core data types used within the MediaDrive framework. In a preferred embodiment, this is

accomplished using a service description language that will be referred to herein as the MediaDrive Service Description Language, or MSDL. In addition, MSDL provides simple ways for referencing semantic descriptions associated with described services.

[0565] In a preferred embodiment, MSDL is an XML schema-based description language that embodies an extensible set of data types enabling the description and composition of services and their associated interface bindings. Many of the object types in MSDL are polymorphic and can be extended to support new functionality. In a preferred embodiment, a basic MSDL profile defines a minimum set of data types and messages for supporting MediaDrive interaction patterns and infrastructural functionality. This basic profile will be referred to as the "MSDL Core." MSDL users may either directly in an ad-hoc manner, or through some form of standardization process, define other MSDL profiles built on top of the MSDL Core adding new data and service types and extending existing ones. In one embodiment, the MSDL Core includes definitions for some or all of the following major basic data types:

- Node – representation of a participant in the MediaDrive framework.
- Device – encapsulates the representation of a virtual or physical device.
- User – encapsulates the representation of a client user.
- Content Reference – encapsulates the representation of a reference or pointer to a content item. Such a reference will typically leverage other standardized ways of describing content format, location, etc.
- DRM Reference – encapsulates the representation of a reference or pointer to a description of a digital rights management format.

- Service – encapsulates the representation of a set of well-defined functionality exported from a node.
- Service Binding – encapsulates a specific way to communicate with a service.
- Request – encapsulates a request for a service to a set of targeted nodes.
- Request Input – encapsulates the input for a request.
- Response – encapsulates a response associated with a request.
- Request Result – encapsulates the results within a response associated with some request.

[0566] In one embodiment, the MSDL Core includes definitions for some or all of the following basic services:

- Authorization – a request or response to authorize some participant to access a service.
- Message Routing – a request or response to provide message routing functionality, including the ability to have the service providing node forward the message or collect and collate messages.
- Node Registration – a request or response to perform registration operations for a node, thereby allowing the node to be discovered through an intermediate node.
- Node Discovery (Query) – a request or response related to the discovery of MediaDrive nodes.
- Notification – a request or response to send or deliver targeted notification messages to a given set of nodes.
- Service Discovery (Query) – a request or response related to the discovery of services provided by some set of one or more nodes.
- Security Credential Exchange – a request or response related to allowing nodes to exchange security related information, such as key pairs, certificates, or the like.

- Upgrade – represents a request or response related to receiving a functionality upgrade. In one embodiment, this service is purely abstract, with other profiles providing concrete representations.

10. Architectural overview

10.1. Service Adaptation Layer

[0567] Referring once again to Fig. 2, the Service Adaptation Layer can be used to expose the native service(s) of a participant in the MediaDrive framework. In a preferred embodiment, MSDL is used to describe how to specifically bind to the service on one or more interfaces. A service description may also include other information, including some or all of the following:

- A list of one or more service providers that will be responsible for authorizing access to the service.
- A pointer to a semantic description of the service's purpose and usage.
- If the service is a composite service resulting from the choreographed execution of a set of other services, a description of the necessary orchestration.

[0568] In addition to serving as the logical point at which services are published in MSDL, in a preferred embodiment the Service Adaptation Layer also encapsulates the concrete representations of the MSDL data types and objects contained within any MSDL profiles for those platforms that are supported by a given participant. It also contains the logic for mapping MSDL messages corresponding to service requests to the appropriate native service implementation.

[0569] Typically, an implementation of the Service Adaptation Layer includes the elements shown in Fig. 4. Specifically, a typical Service Adaptation Layer implementation will include the following elements in a layered design:

- A layer that encapsulates the service interface entry points as described in the MSDL bindings of the service interface. Through these access points, service invocation is made, input parameter data is passed in, and results flow out.
- A layer that corresponds to the logic for MSDL message processing. This typically contains some sort of message pump that drives the processing, some type of XML data binding support, and low level XML parser and data representation support.
- A layer that represents the available native services onto which the corresponding MSDL described service messages are mapped.

10.2. Framework Support Library

[0570] Referring once again to Fig. 2, the Framework Support Library provides a collection of optional support functions that make it easier to enable an entity to participate in the MediaDrive framework. How the support library is factored (collections of objects, collections of functions, etc.) and how it is implemented will vary depending on the targeted platform. There may be multiple versions of the support library supporting different functionality. In the general case, the functionality available within the support library includes some or all of the following:

- Service Proxy – encapsulates the functionality that enables a MediaDrive node to make a service invocation request to a targeted set of service-providing nodes and receive a set of responses.
- MSDL Message Manipulation Routines – functionality for manipulating MSDL message parts.
- Service Cache Interface – a common service provider interface allowing a node to manage mappings between discovered nodes and the services they support.
- Workflow Collator Interface – service provider interface allowing a node to manage and process collections of MSDL messages. It provides the basic building

blocks to orchestrate services. Currently, this interface is most often implemented by a node that supports message routing and supports the intermediate queuing and collating of messages.

- Notification Processor Interface – a service provider interface for hooking a MediaDrive node that supports notification processing to some well-defined notification processing engine.
- Miscellaneous Support Functionality – routines for generating message ids, timestamps, etc.

10.3. Communication Styles

[0571] It is helpful to give a basic overview of the general communication styles supported in preferred embodiments of the MediaDrive Framework and utilized in different Support Library components such as the Service Proxy. There are a few basic types of communication patterns, including:

[0572] *Asynchronous RPC Delivery Style* – This model is appropriate if there is an expectation that fulfilling a request will take an extended period of time and the client does not want to wait. Here, the client submits a request with the expectation that it will be processed in an asynchronous manner by some service-providing node(s). A service-providing endpoint may respond with an indication that it does not support this model, or, if the service-providing node does support this model, it may return a response that carries a ticket that can be submitted to the service-providing node in subsequent requests to determine if it has a response to the client's request. In a preferred embodiment, service-providing endpoints that support this model will cache responses to pending requests based on some internal policy. If a client attempts to redeem a ticket associated with such a request, and no response is available, or the response has been thrown away by the service-providing node, then an appropriate error response can be returned.

[0573] *Synchronous RPC Delivery Style* – Here, the client submits a request then waits for one or more responses to be returned. A service-providing MediaDrive-enabled endpoint may respond with an indication that it does not support this model.

[0574] *Message Based Delivery Style* – Here, the client submits a request indicating that it wants to receive any responses via a message notification associated with one or more of its notification handling service interfaces. A node may respond indicating that it does not support this model. Exemplary embodiments of the notification framework are described below.

[0575] None of the interaction patterns above necessitates blocking and waiting for responses or explicit polling. It is possible to use threading or other platform-specific mechanisms to model both blocking and non-blocking semantics with the above delivery style mechanisms.

10.4. Service Proxy

[0576] MediaDrive-enabled endpoints may use diverse discovery, name resolution, and transport protocols. As a result, it will typically be desirable to use a flexible and extensible communication API. The Service Proxy API provides a common interface to the functionality for allowing a MediaDrive node to make service requests and receive responses for a set of targeted MediaDrive nodes. A service proxy may be implemented in a variety of forms within the boundaries of a client (in the form of a shared library) or outside the boundaries of the client (in the form of an agent running in a different process). The exact form of the service proxy implementation is preferably tailored towards the needs of the platform or client. Use of the service proxy is optional, although in general it provides significant utility. Some of the common uses of the service proxy include:

- A common, reusable API for service invocation.
- Encapsulation of issues surrounding the negotiation and use of a transport channel. Some transport channels require SSL session setup over TCP/IP, some channels may only support unreliable communication over UDP/IP, and other channels may not be IP based at all. In a preferred embodiment, clients are generally shielded from these details.
- Encapsulation of issues surrounding the discovery of an initial set of MediaDrive nodes for message routing. For example, a cable set-top box may have a dedicated connection to the network and mandate that all messages flow through a specific route and intermediary, whereas a portable media player in a home network may use UPnP discovery to find multiple endpoints that are directly accessible. Again, clients are preferably shielded from these details.
- "Legacy" clients may be unable or may choose not to converse directly with other MediaDrive nodes by the creation of MSDL messages. Here, a version of the service proxy may be created that understands whatever native protocol the client supports.

[0577] The service proxy may be implemented as a static component, only supporting a fixed set of protocol bindings, or it may be implemented such that it is able to dynamically support new bindings. Two examples of common service proxy models are: (1) the service proxy directly receives and returns MSDL messages to the client, and (2) the service proxy supports some native protocol understood by the client. The service proxy can be viewed as having two sides: a client-side that the requesting participant uses, and a service-side that interacts with other MediaDrive-enabled endpoints.

[0578] Two primary patterns of interaction between the client and the service proxy are:

[0579] Pattern 1 (Figure 5a). The client forms MSDL request messages directly and submits them to the service proxy. It receives one or more responses in

MSDL format that it will typically need to collect, parse and process. The client may also submit explicit set(s) of service bindings to use in targeting the delivery of the request. These service bindings may have been obtained by the client performing service discovery operations, or by using information obtained from previous request responses.

[0580] Pattern 2 (Figure 5b). The service proxy accommodates a client-side model where the client may interact through some native communication protocol that the service proxy will accept, translating to/from MSDL to allow the client to participate within the MediaDrive framework. Here, the native protocol or a combination of the native protocol and the execution environment provide any needed information for the service proxy to generate appropriate MSDL requests and determine suitable target service bindings if necessary.

[0581] On the service side, multiple types of patterns of interaction between the service proxy and the service-providing MediaDrive-enabled endpoints are supported, and, as with the client-side, the interaction patterns can be tailored and may vary based on a variety of criteria including nature of request, underlying communication network, nature of application, and/or transport protocols associated with any targeted service bindings.

[0582] In one pattern, the service proxy initiates communication with multiple potential service providers directly, and receives responses back. This type of interaction pattern may conform to multiple service bindings being relayed from the client for use by the service proxy, or it may be indicative of a broadcast or multicast network being present that a service proxy uses in relaying messages. The service

proxy may choose to collect and collate responses or just return the first acceptable response.

[0583] In the model shown in Fig. 6, the service proxy does not directly communicate with a targeted service-providing endpoint, but instead routes requests through an intermediate node that relays the requests, receives any responses, and sends them back to the service proxy. This interaction pattern may be due to:

- Service proxy is unable/unwilling to directly support any of the service bindings associated with service-providing endpoints but can establish a relationship with an intermediate node that is willing to act as a gateway.
- Client may be unable to discover or otherwise determine the service bindings for any suitable service providing nodes, and is willing to allow an intermediate node to attempt to discover any suitable service providers.
- Service proxy wants to take advantage of an intermediate node that supports more robust collection and collating functionality allowing for more flexible communication patterns between the service proxy and service providers.

[0584] Alternatively, or in addition to the above basic service-side interaction patterns, it is possible to have implementations of the service proxy that implement combinations of the above patterns or new patterns.

10.5. Workflow Collator

[0585] In a preferred embodiment, the fulfillment of most non-trivial MediaDrive service requests will often require some type of coordinating mechanism that can manage the flow of events of a request, manage any associated data including transient and intermediate results, and ensure that the rules (if any) associated with fulfillment are followed. This functionality can be provided in the form of transaction coordinators ranging from simple transaction monitors in relational databases to more generalized agents as seen in Microsoft MTS/COM+. In a preferred embodiment, the

Workflow Collator is a programmable mechanism through which MediaDrive nodes orchestrate the processing and fulfillment of service invocations.

[0586] In a preferred embodiment, the Workflow Collator can be tailored towards a specific MediaDrive endpoint's characteristics and requirements, and is designed to support a variety of functionality, ranging from traditional message queuing to more sophisticated distributed transaction coordination. A simple Workflow Collator provides an interface for storage and retrieval of arbitrary MSDL messages. By building on this, it is possible to support a wide variety of functionality, including:

- Collection of service requests for more effective processing.
- Simple aggregation of service responses into a composite response.
- Manual orchestration of multiple service requests and service responses in order to create a composite service.
- Automated orchestration of multiple service requests and service responses in order to create a composite service.

[0587] Fig. 7 illustrates another potential interaction pattern. The interaction pattern begins with a service request arriving at a MediaDrive node and being accepted through the node's Service Adaptation Layer. The message is handed off to the MSDL Message Pump that will initially drive and in turn be driven by the Workflow Collator to fulfill the request and return a response. This pattern is the basis for a number of service processing patterns. In a common simple case, the fulfillment of a service request can be represented in a single request/response interaction pattern where the processing is idempotent and the rules for processing the request are completely expressed as part of the request in MSDL. But, the fulfillment of a service request may depend on multiple MSDL messages arriving, processed asynchronously

over some period of time. In this case, the processing is part of a potentially long-lived transaction monitored by the Workflow Collator.

[0588] In even more complex scenarios, the fulfillment of a given service request may require the participation of multiple nodes in a coordinated fashion. The rules for processing the request may be expressed in MSDL that can optionally wrap messages from other service orchestration description standards such as BPEL. In the fulfillment of a request, a node's Workflow Collator may depend on functionality through some set of local interfaces or it may interact with other MediaDrive nodes.

[0589] When a MSDL request message is given to the Workflow Collator, it determines what the correct rules are for processing the request. Depending upon the implementation of the Workflow Collator, the service description logic may be represented in the form of a fixed state machine for a set of services that the node exposes, or it may be represented in ways that support the processing of a more free form expression of the service processing logic such as may be described using BPEL4WS. The Workflow Collator architecture is preferably modular and extensible, supporting plug-ins. In a preferred embodiment, MSDL itself supports a simple way of expressing the rules for orchestrating multiple MediaDrive services into composite services.

[0590] In addition to interpreting service composition and processing rules, in a preferred embodiment the workflow collator will need to determine whether the MSDL message should be used in the context of initiating a service fulfillment processing lifecycle, or whether the MSDL message is just further input in the chain of an ongoing transaction. MSDL messages include IDs and metadata that can be used for purposes of making these types of determinations; it is also possible to extend

MSDL messages to include additional information that may be service transaction specific, facilitating the processing of messages.

10.6. Notification Services

[0591] In addition to both asynchronous and synchronous RPC-like communication patterns, where the client specifically initiates a request and then either waits for responses or periodically checks for responses through redemption of a ticket, preferred embodiments of the MediaDrive framework also support a pure messaging type of communication pattern based on the notion of notification. The following elements constitute the core MSDL data and message types used in the notification framework of a preferred embodiment:

- Notification – a message containing a specified type of payload targeted at interested MediaDrive-enabled endpoints (nodes).
- Notification Interest – criteria used to determine whether a given node will accept a given notification. Notification interests may include interests based on specific types of identity (e.g., node id, user id, etc.), events (e.g., node discovery, service discovery, etc.), affinity groups (e.g., new jazz club content), or general categories (e.g., advertisements).
- Notification Payload – the typed contents of a notification. Payload types may range from simple text messages to more complex objects.
- Notification Handler Service Interface – the type of service provider interface on which notifications may be received. The service provider also describes the notification interests associated with the interface, as well as the acceptable payload types. A node supporting this interface may be the final destination for the notification or an intermediary processing endpoint.
- Notification Processor Service – a node that is capable of matching notifications to interested nodes, delivering the notifications based on some policy.

- Notification Originator – a node that sends out a notification targeted to a set of interested nodes and/or an intermediary set of notification processing nodes.

[0592] As with all MSDL data types and messages, the notification, notification interest, and notification payload are preferably extensible. Additionally, the notification handler service interface is preferably subject to the same authorization process that any other MediaDrive service interface is subject to. Thus, even though a given notification may match in terms of interest and acceptable payload, a node may refuse to accept a notification based on some associated interface policy related to the intermediary sender or originating source of the notification.

[0593] Fig. 8 depicts a set of notification processing nodes discovering a node that supports the notification handler service. As part of its service description, the node that supports receiving notifications designates what its notification interests are and what notification payload types are acceptable.

[0594] Fig. 9 depicts how notifications can be delivered. Any node could be the originating source as well as processor of the notification, and could be responsible for delivering it. Notification processors that choose to handle notifications from foreign notification-originating nodes may integrate with a commercial notification-processing engine such as Microsoft Notification Services in order to improve efficiency.

10.7. Service Discovery

[0595] Preferred embodiments of the MediaDrive Framework support a set of flexible models for discovering services. These models can be mapped onto existing discovery models. In one embodiment, the framework supports three basic models of service discovery:

- Client Driven – a MediaDrive node explicitly sends out a request to some set of targeted nodes that support a "Service Query" service interface asking whether they support a specified set of services. If the requesting node is authorized, the service providing node will report that it supports the requested interfaces and the associated service interface bindings. This is one of the more common interfaces that nodes will support if they expose any services.
- Node Registration – a node can register its description, including its supported services, with other nodes. If a node supports this interface, it is willing to accept requests from other nodes and then cache those descriptions based on some policy. These node descriptions are then available directly for use by the receiving node or to other nodes that perform service queries targeted to nodes that have cached descriptions.
- Event-Based – Nodes send out notifications indicating a change in state (e.g., node active/available), or a node advertises that it supports some specific service. The notification can contain a full description of the node and its services, or just the ID of the node associated with the event. Interested nodes may then choose to accept and process the notification.

[0596] Fig. 10 depicts the client-driven scenario where some requesting node makes a service query request to a targeted service providing node and receives a response indicating whether the service provider supports the requested service(s).

[0597] Fig. 11 depicts the node registration scenario where some requesting node wants to register its description with some other targeted node that is willing to receive its description and cache it based on some internal policy. The requesting node makes a registration request to a targeted service providing node and receives a response indicating that the service provider cached its description.

[0598] Fig. 12 depicts a MediaDrive-enabled node being alerted to the existence of a service via a notification. Such a notification may come in many

forms, including notifications related to a general status change in a node's state, or a node explicitly sending out a notification related to a service's availability.

10.8. Service Authorization

[0599] Before a MediaDrive node allows access to a specified service, it preferably first determines if the requesting node is permitted access. In a preferred embodiment, allowing access is based on two different but related criteria. The first is whether the service-providing node trusts the requesting node for this communication, and the second is whether there is a policy permitting the requesting node to access the requested service's interface.

10.8.1. Establishing Trust

[0600] In a preferred embodiment, the MediaDrive framework does not mandate the specific requirements, criteria, or decision-making logic for how an arbitrary set of nodes come to trust each other. Trust semantics may vary radically from node to node. The MediaDrive framework strives to provide a standard set of facilities that allow nodes to negotiate a mutually acceptable trusted relationship. In the determination and establishment of trust between nodes, preferred embodiments of the MediaDrive framework support the exchange of credentials between nodes that can be used for establishing a trusted relationship. Within the MediaDrive framework, trust-related credentials may be exchanged using a variety of different models. For example:

- Service-Binding Properties – a model where trust credentials are exchanged implicitly as part of the service interface binding. For example, if a node exposes its service in the form on an HTTP Post over SSL, or as a Web Service that requires WS-Security XML Signature then the actual properties of this service binding may communicate all necessary trust-related credentials.

- Request/Response Attributes – a model where trust credentials are exchanged through MSDL request and response messages, optionally including the credentials as attributes on the messages.
- Explicit Exchange – a model where trust credentials are exchanged explicitly through a service-provider interface that allows querying of information related to the trust credentials that a given node contains. This is generally the most involved model, typically requiring a separate roundtrip session in order to exchange credentials. The service interface binding itself provides a mutually acceptable trusted channel.

[0601] Trust-related credentials may be exchanged directly as attributes on request and response messages. For example, digital certificates could be attached to, and flow along with, request and response messages, and could be used for forming a trusted relationship.

[0602] In the scenario depicted in Fig. 13, nodes may exchange credentials directly via a security credential service if a given node supports this interface.

[0603] In addition to these basic models, the MediaDrive framework preferably supports combinations of these different approaches. For example, the communication channel associated with a semi-trusted service binding may be used to bootstrap the exchange of other security-related credentials more directly or exchanging security-related credentials (which may have some type of inherent integrity) directly and using them for the establishment of a secure communication channel associated with some service interface binding.

[0604] In summary, trust model semantics and the processes that must be followed to establish trust will vary from entity to entity. In some situations mutual trust between nodes may not be required. This type of dynamic heterogeneous

environment calls for a flexible model that provides a common set of facilities that allow different entities to negotiate context-sensitive trust semantics.

10.8.2. Policy-Managed Access

[0605] In addition to establishing a trusted relationship between a set of interacting nodes, before a service providing node allows a requesting node to access a service it preferably must also determine whether this access is permitted based on any policy associated with the service interface. The MediaDrive framework provides a consistent, flexible way of making this determination. In a preferred embodiment, as part of the service description, one or more MediaDrive nodes will be designated as authorization service providers. Each of these authorization service providing nodes will be required to implement a standard service for handling and responding to "Authorization" query requests. For example, before access is allowed to a service interface, the targeted service provider may dispatch an "Authorization" query request to all MediaDrive nodes that authorize access to its service, and access may be allowed only if all authorizing entities respond indicating that access is permitted.

[0606] As Fig. 14 depicts, the service providing node makes authorization requests to all nodes that manage access to the requested service, and then, based on the decision(s), either processes and returns the applicable service response, or returns a response indicating that access was denied. In addition to providing a consistent mechanism, it is desirable for the mechanism by which authorizing nodes reach a decision be as open as practically possible so that the most appropriate policy decision making mechanism can be implemented for a specific platform and client. MediaDrive is preferably policy management system neutral; it does not mandate how

an authorizing node reaches decisions about authorizations to services based on an authorization query.

[0607] While MediaDrive preferably does not mandate the specific policy engine or policy language used by an authorizing node for a given platform, in a preferred embodiment it is required for interoperability that the authorization request have some standard form that can be exchanged and handled. The "Authorization" query request, which is extensible, is capable of carrying a flexible payload so that it can accommodate different types of authorization query requests in the context of different policy management systems. In one embodiment, support is provided for at least two authorization formats: (1) a simple format providing a very simple envelope using some least common denominator criteria as input, such as a simple requestor ID, resource ID, and action ID, and (2) a SAML Format that provides a way of passing an enveloped Security Assertion Markup Language form of authorization query.

[0608] In a preferred embodiment, it is a requirement that any authorizing node recognize and support at least the "simple" format and be able to map it to whatever native policy expression format exists on the authorizing node. For other formats, the authorizing node should return an appropriate error response if it does not handle or understand the payload of an "Authorization" query request. Extensions to the basic MediaDrive framework may include the ability for nodes to negotiate over acceptable formats of authorization query, and for nodes to query to determine what formats are supported by a given authorizing service provider node.

[0609] In one embodiment, the "Authorization" query response, though extensible, consists of a simple form to indicate if access is granted or not. Other

forms of this response for this service are possible, including the specification of obligations that the requesting node might have to fulfill if access to the given interface is to be granted.

11. Consumer Media Applications

[0610] Embodiments of MediaDrive can be used to link various consumer devices to a number of different services in a multi-tiered environment that includes personal and local area networks, home gateways, and IP and non-IP based wide area networks. For example, interoperability has been successfully demonstrated in one interconnected system using cell phones, game platforms, PDAs, PCs, web-based content services, discovery services, notification services, and update services. A present embodiment supports multiple media formats (e.g., MP4, WMF, et al), multiple discovery protocols (for discovery of services over Bluetooth and through registries such as UDDI, LDAP, Active Directory), and IP-based notification and Wireless SMS notification on the same device. The orchestration feature is used to help the consumer overcome interoperability barriers. When there is a query for content, the various services are coordinated in order to fulfill the request, including discovery, search, matching, update, rights exchange, and notification services. In preferred embodiments, it is desirable for a consumer to be able to use most any device, make a wish for content, and be nearly instantly fulfilled with the content and the rights and rendering capabilities to both use and share the content. The orchestration capability allows the consumer to view home and Internet-based content caches from virtually any device at any point in a dynamic multi-tiered network. This capability can be expanded to more advanced services that promote sharing of streams

and playlists, making impromptu broadcasts and narrowcasts easy to discover and connect to, and using many different devices while ensuring rights are respected.

[0611] Beyond the consumer-centric side, it is generally desirable to find ways to provide an end-to-end interoperable media distribution system that does not rely on a single set of standards for media format, rights management, and fulfillment protocols. In the value chain that includes content originators, distributors, retailers, service providers, device manufacturers, and consumers, there are a number of localized needs in each segment. This is especially true in the case of rights management, where content originators often need to express rights of use that may apply differently in various contexts to different downstream value chain elements. A consumer gateway typically has a much narrower set of concerns, and an end-user device typically has an even simpler set of concerns, namely just playing the content.

[0612] With a sufficiently automated system of dynamically self-configuring distribution services, content originators can produce and package content, express rights, and confidently rely on value added by other service providers to nearly instantly provide the content to all interested consumers, no matter where they are or what kind of device they are using. Embodiments of MediaDrive can be used to fulfill (or approach fulfillment of) this goal, or aspects thereof, and can provide means for multiple service providers to innovate and introduce new services that benefit both consumers and service providers without having to wait for, or depend on, a monolithic set of end-to-end standards. One possible approach allows digital rights management to be decomposed into components with a more natural separation of concerns that focus on policy management of service interfaces rather than on copy protection. This approach has the potential to change the tension between consumers

and content providers in the digital content domain, as content is provided through a rich infrastructure of services that provide consumers with useful information and instant gratification. Policy management can limit the extent to which pirates can leverage those legitimate services. Indeed, MediaDrive is designed to allow the network effect to encourage the evolution of a very rich set of legitimate services resulting in value that will exceed that which can be provided by pirates.

12. General Enterprise Applications

[0613] Policy managed P2P service orchestration appears to have some interesting applications in the business world to areas such as supply chain management (SCM) and customer relationship management (CRM). CRM is about finding, getting, and retaining customers. It is a core element in almost any customer-centric eBusiness strategy where it is advantageous to focus on the customer using integrated information systems and contact center implementations that allow the customer to communicate via any desired communication channel. This is also true for suppliers. Most enterprises have some sort of hierarchy and departmentalization, and the departmental and process view that works well internally does not necessarily work well for customers and suppliers. It is desirable to allow an enterprise's departments to publish a comprehensive set of service interfaces for their department, so that those services can be adapted by different customers and suppliers (according to policy) in ways that optimize their own processes and services. This could be a significant step beyond the enterprise information portal.

[0614] Today, most enterprises implement CRM using relatively expensive pre-packaged application suites that appear as monolithic software applications using proprietary client/server infrastructure. These suites interoperate within an enterprise

after a significant amount of configuration. Web services allow the use of open protocols and a distributed web-service infrastructure where communication not only occurs over closed LANs but also over MANs, and WANs. This model allows remote Application Service Providers (ASPs) to manage and deliver CRM application capabilities as collections of services to multiple entities within a business from diverse data centers. The business unit or department becomes a portal and in some cases an aggregator for these services. This can be seen in the context of the evolution that the IT industry is undergoing in the area of Service Oriented Architectures/Applications (SOA). Both the business's and the customer's interface to the application related services can range from simple browser-based applications to larger footprint PC-centric applications.

[0615] This new model still faces some significant hurdles that impose barriers to its acceptance in the mainstream: For example:

[0616] *Information Confidentiality and Access Control.* When an ASP provides a CRM solution, the business may encounter issues surrounding the placement of sensitive customer data in the hands of another company. The company must ensure that only legitimate entities can access this information. From the customer's perspective they now potentially have a new relationship (whether implicit or explicit) with another company. What is need are infrastructural mechanisms that allow both the customer and the business to exercise access control over their information.

[0617] *Application Extensibility.* An ASP often aligns with a particular small set of technology vendors and their associated products or services. Businesses often prefer the flexibility to leverage a diverse set of technologies from different vendors.

The application framework shouldn't necessarily be constrained by a particulars ASP's current business relationships or technology choices, instead it would be desirable to have infrastructural mechanisms that allow an ASP to quickly and effectively integrate with new functionality provided from different technology vendors or other ASPs. The application framework should also allow for integration of CRM services provided directly by business units themselves based, for example, on in-house developed extensions.

[0618] *Service Level Agreement (SLA).* A business's service requirements will usually not be static. Instead, they will typically vary based on changing business conditions, including new functionality driven by competition, operational load characteristics based on internal business activities and the number of customers being supported for a given activity, and the ability to leverage lower cost service alternatives. In some cases these changes may require human intervention over a significant period of time, but in other cases the changes can be automated or quickly reconfigured based on a set of specified rules. For both of these cases one can view the CRM application services as being orchestrated into a new configuration based on the existing environment and the authorized agreement between the ASP and the business unit. These types of applications can be characterized as being self-organizing.

[0619] Embodiments of MediaDrive can play a role in addressing some or all of these needs. For example, embodiments of MediaDrive can be used to provide the necessary infrastructural building blocks required to address the above requirements in building dynamic, service-oriented, self-organizing applications such as CRM and

SCM where only trusted, authorized entities have access to the pertinent services and information.

[0620] Fig. 15 illustrates a scenario showing how MediaDrive can be used in the context of delivering CRM-related services to different stakeholders. In the scenario shown in Fig. 15, a business unit is able to use a variety of different ASPs to provide the services necessary for constructing different types of CRM applications or providing different views on customer data to different departments within its organization or to its customers.

[0621] The scenario illustrates several aspects that use the functionality provided by preferred embodiments of the MediaDrive framework, including some or all of the following:

- Implicit or explicit interoperability between different CRM technology vendors. By describing the necessary interfaces and providing the appropriate bindings (using Service Adaptation Layer) either at the level of the ASP or directly on top of existing CRM software suites, it becomes possible for a business to use a single ASP or multiple ASPs with diverse CRM technologies from different providers.
- By accessing customer data through the MediaDrive framework, it is possible to leverage its rich policy management framework for access control.
- A department can access the necessary CRM related services for its application through the service proxy, thereby insulating the business unit from issues surrounding compatibility of transport protocol, data formats, etc.
- Interactions between departments, data warehouses, and applications service providers in the MediaDrive framework can be governed. Preferably, only legitimately authorized entities can access the services, and service invocation will occur only if the different MediaDrive framework participants can establish an acceptable level of trust in their interactions.

- Service composition and orchestration. CRM applications are self-organizing in that the application consists of services from several different ASPs, orchestrated together, or the application is delivered from a single ASP but still consists of multiple services whose composition and use need to be dynamically configured for different customer views. The functionality exposed by the service providers can change over time driven by the service level agreement (SLA) and the execution environment. The MediaDrive Workflow Collator can be used to orchestrate services. The rules used for controlling this interaction may be specified directly in MSDL or codified in some standard way in the SLA, as long as there is a plug-in written for the Workflow Collator that can understand and interpret the description. Recall that the Workflow Collator works in collaboration with the MediaDrive Message Pump to collect and process the necessary messages in order to determine how to provide the specified services based on the requirements.

[0622] In a current embodiment, MediaDrive is implemented using a set of distributed policy management services based on InterTrust's PolicyWorks and the PolicySpeak policy expression language. However, MediaDrive is designed to use other policy management frameworks and languages such as IBM's PolicyDirector and XACML, and, in a preferred embodiment, MediaDrive offers a policy management services interface, so that any suitable policy management frameworks and/or policy expression languages could be used.

[0623] By way of background, PolicyWorks is a set of software components that can be configured to originate, deploy, manage, and enforce rules, rights and policies regarding the use of network-based content, services, and resources within and among enterprises as well as between enterprises and consumers. Noting that enterprise policies or controls often need to interact with means for identifying users, credentials, directories, content metadata, and other policies and controls, PolicyWorks is designed to be efficiently integrated into an enterprise and to leverage

existing applications, directories, and identification systems. A goal of PolicyWorks is to simplify and unify the secure management of policies and rules across a variety of major applications ranging from network and file system access management, enterprise information portals and asset management, to supply chain management, web services, digital locker services, and consumer media subscription applications.

[0624] In a preferred embodiment, the PolicyWorks architecture is based on a logical model that exposes a number of basic interfaces:

- Governed Operation – This interface is used by applications such as those mentioned above to govern the actions of these applications regarding a specific resource, content item, or service. Virtually any item can be governed using PolicyWorks, providing persistent security protection, access control, usage monitoring and auditing, or transaction management. Through this interface, an application can make an item, service, or resource available, under policy control, with assurance that rules and requirements for use are followed. A governance component encapsulates the methods and operations necessary to govern the item. PolicyWorks allows the use of a number of customizable and off-the-shelf solutions for such governance, such as native document protection mechanisms in Microsoft Office and Adobe Acrobat, and the code resource and service protection mechanisms afforded by J2EE EJBs and the Microsoft .Net Framework.
- Policy Query – This interface is used to determine precisely what actions are allowed with a governed item, resource, or service for a given user, set of conditions, environment data, and any other relevant factors. A response to a query can include a permission and/or additional conditions and obligations that are concomitant with use of the item. The governance component, upon receiving a response to the query, takes responsibility for enforcing the policy, and executing (or causing to execute) any obligations and monitoring the consequences of the requested use.
- Credential Query – This interface is used by the PolicyManager to determine attributes of the user that may play a role in the policy associated with a governed

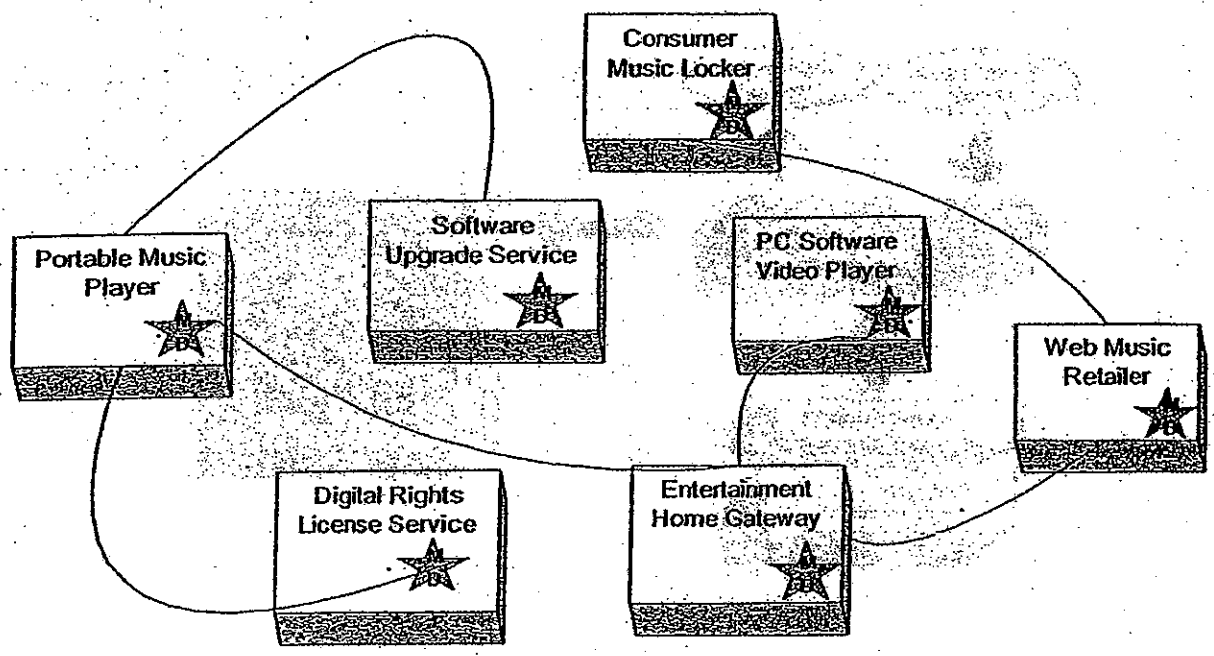
operation. This logical model is enriched in an actual implementation, so that, for example, the PolicyManager component has access to a number of interfaces and resources, both internal and external, to determine the response to a Policy Query, as there may be a number of directories, and sources of rules and metadata concerning the actors (users, policy creators, resource managers, etc) and the Governed items themselves.

- Security Service – This interface is used to obtain identity management services for the requester of a governed action, and to obtain persistent protection services for governed items.
- Policy Admin – This interface is used to administer the distributed policy management databases of rules and controls that pertain to governed items. Various tools for authoring and maintaining policies use this interface. These tools can provide very straightforward ways to apply and administer policy. For example, using this interface one can create a tool that associates policy with a shared folder in a file system, and allow any user to apply policy to a document (or other object) simply by dragging and dropping the document in the folder. Other tools can use more sophisticated authoring techniques that use any number of policy expression languages to specify policy in very precise and targeted terms.

[0625] SCM is another application where members of virtually any enterprise can view their external vendor networks as one big virtual stockroom, with shipping orders policy managed by Media Drive services, and billing coordinated through MediaDrive compatible notification services.

[0626] Although the foregoing invention has been described in some detail for purposes of clarity, it will be apparent that certain changes and modifications may be made without departing from the principles of the present invention. It should be noted that there are many alternative ways of implementing both the processes and apparatuses of the present invention. Accordingly, the present embodiments are to be

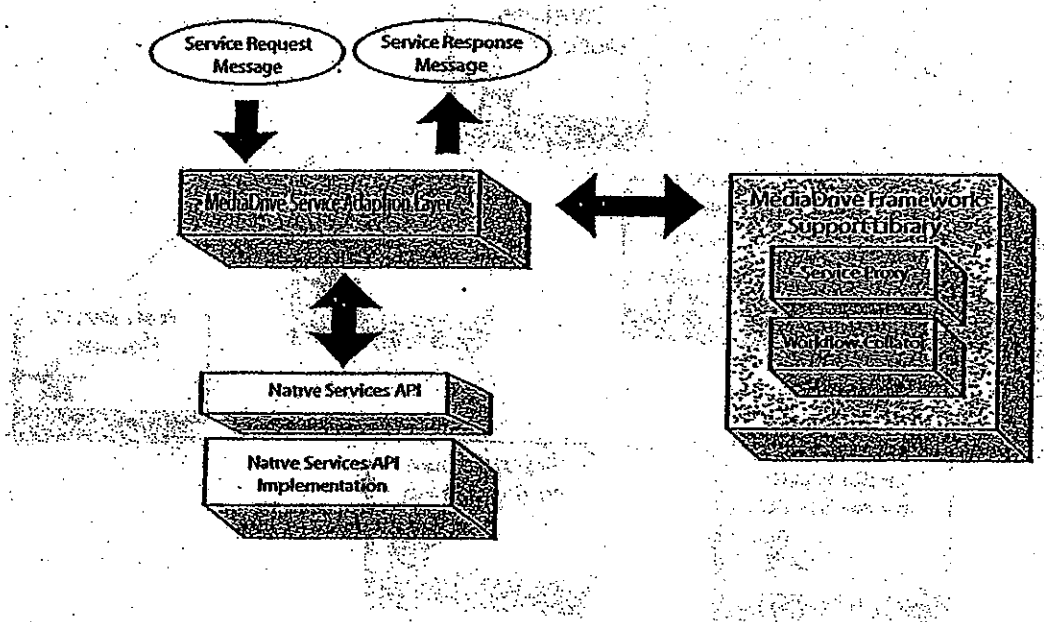
considered as illustrative and not restrictive, and the invention is not to be limited to the specific details given herein.



Example Embodiment of MediaDrive Framework

Fig. 1

(part of Appendix A)



Conceptual View of a MediaDrive Node

Fig. 2

(part of Appendix A)

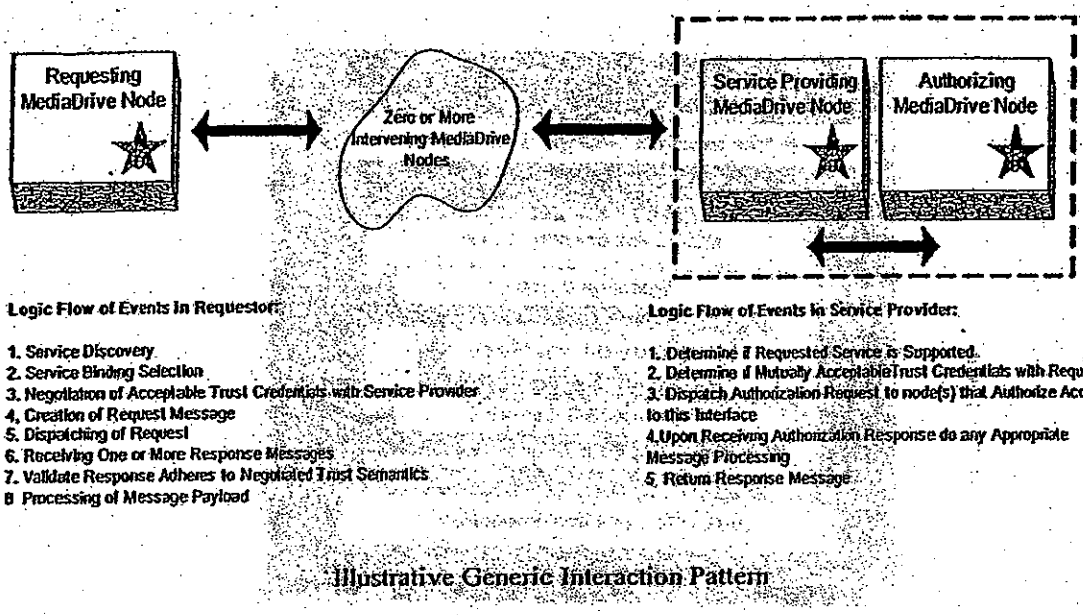
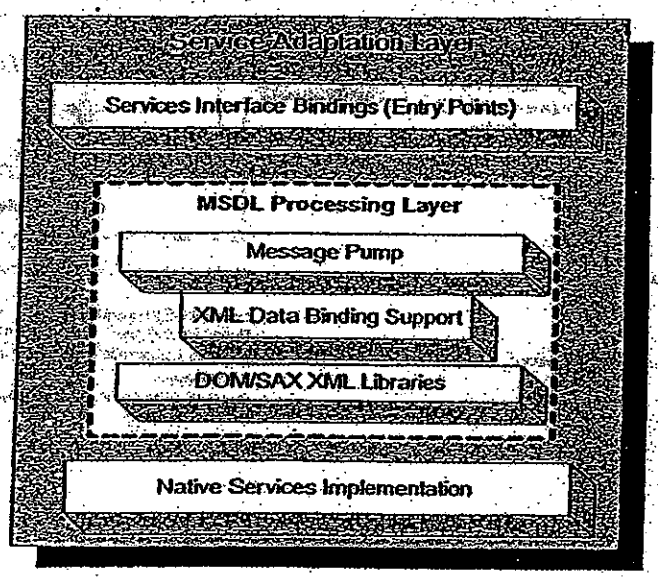


Fig. 3

(part of Appendix A)



Service Adaptation Layer

Fig. 4

(part of Appendix A)



Service Proxy – Client Side MSDL Interaction Pattern

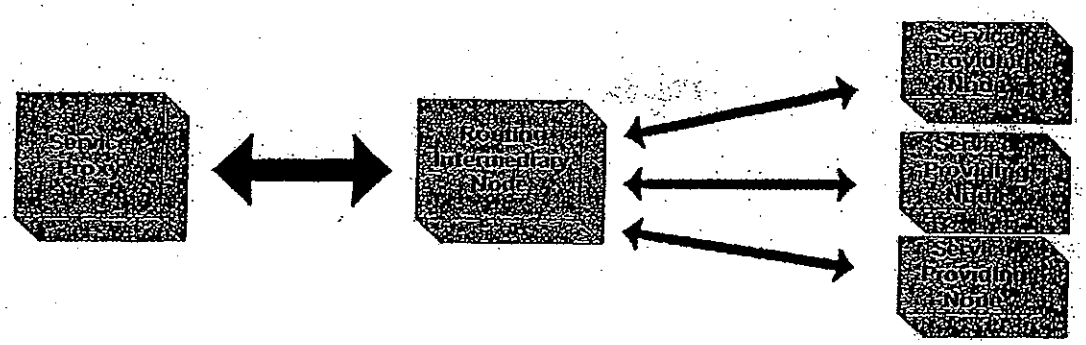
Fig. 5a



'Service Proxy – Client-Side Native Interaction Pattern

Fig. 5b

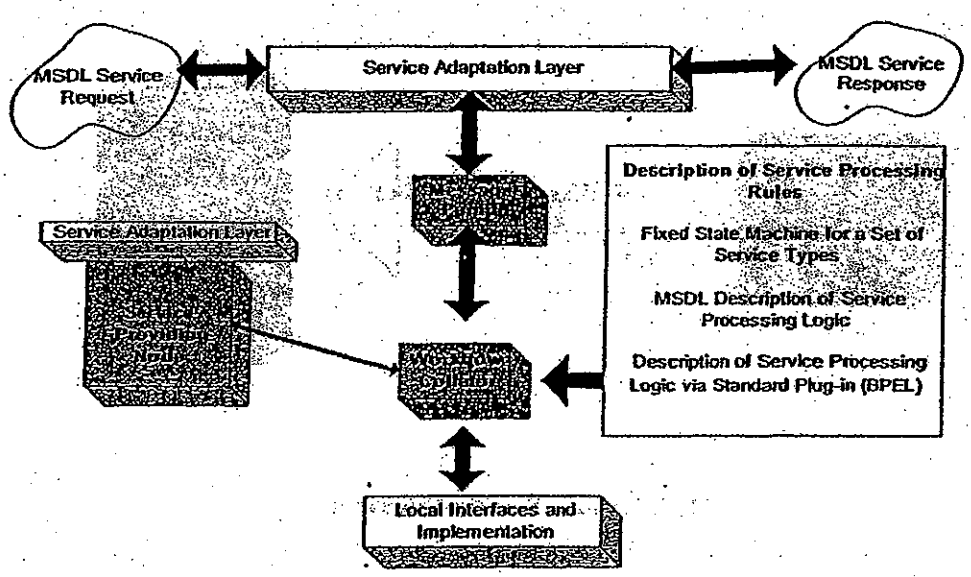
(part of Appendix A)



Service Proxy – Service-Side Point-to-Intermediary Interaction Pattern

Fig. 6

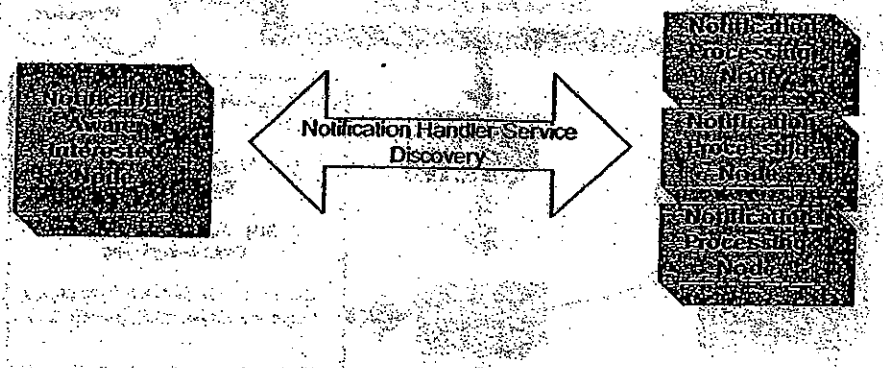
(part of Appendix A)



Interaction Pattern of MediaDrive Workflow Collator

Fig. 7

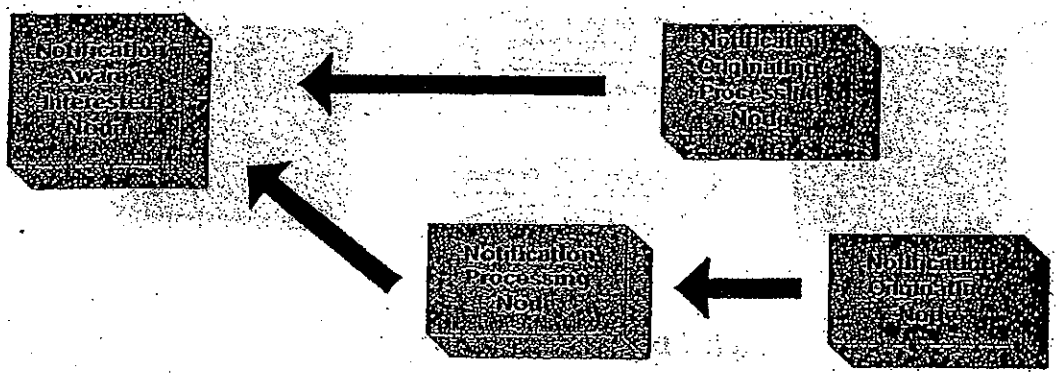
(part of Appendix A)



Notification – Service Discovery

Fig. 8

(part of Appendix A)



Notification – Delivery

Fig. 9

(part of Appendix A)

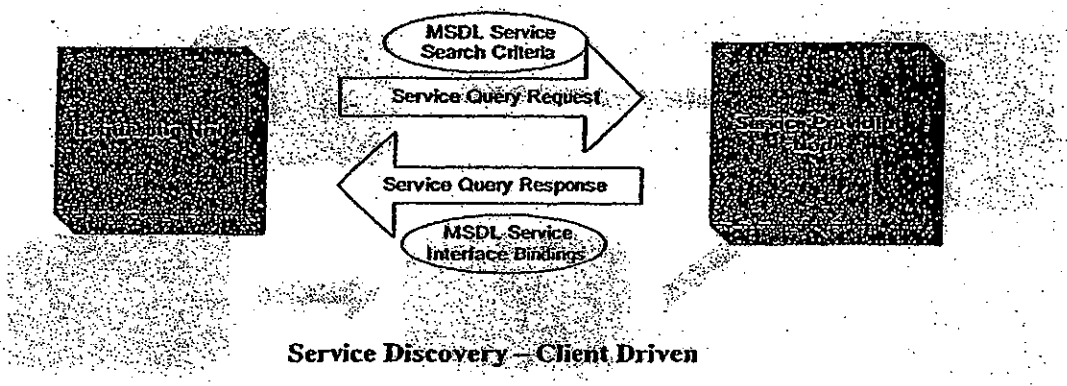
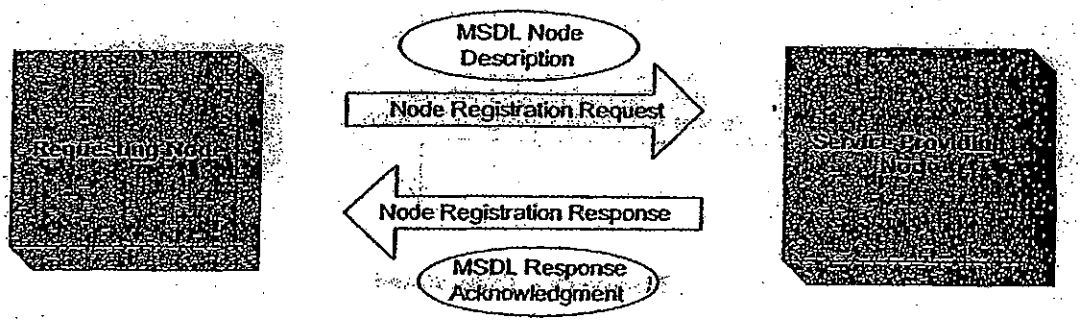


Fig. 10

(part of Appendix A)



Service Discovery – Node Registration

Fig. 11

(part of Appendix A)

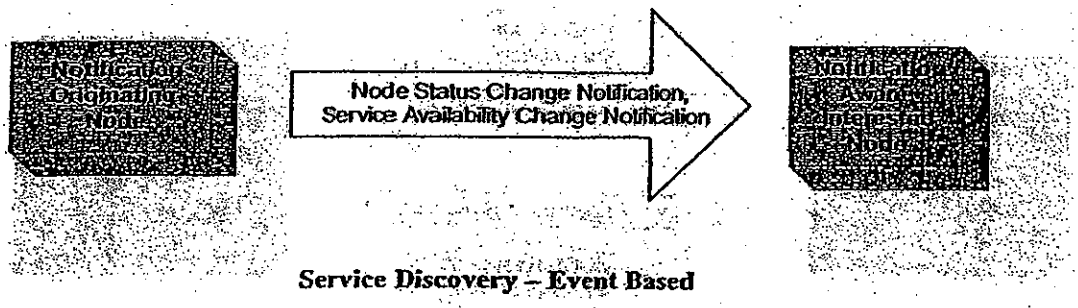
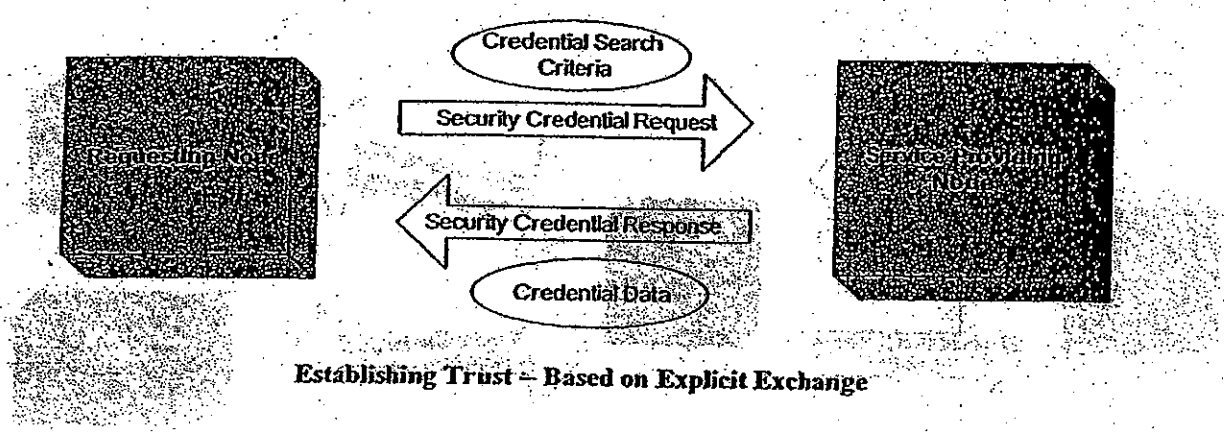


Fig. 12

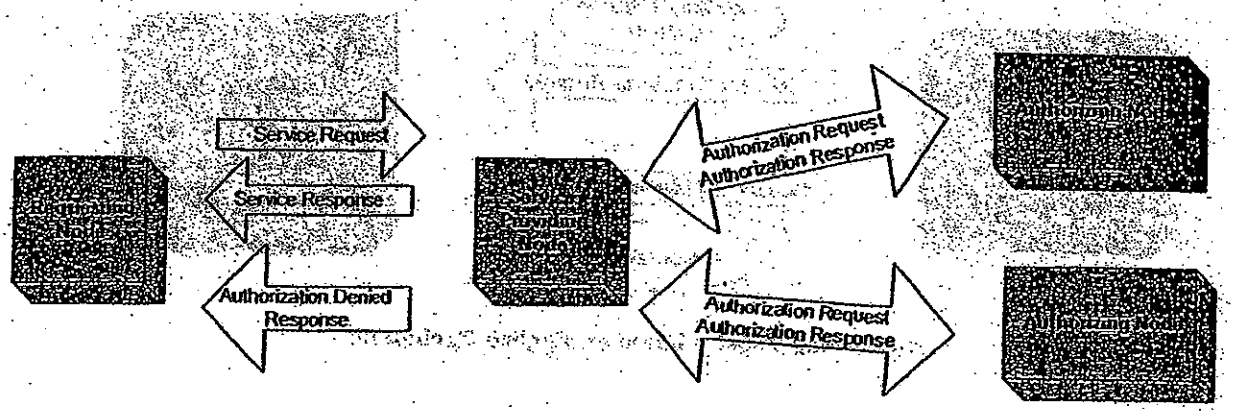
(part of Appendix A)



Establishing Trust – Based on Explicit Exchange

Fig. 13

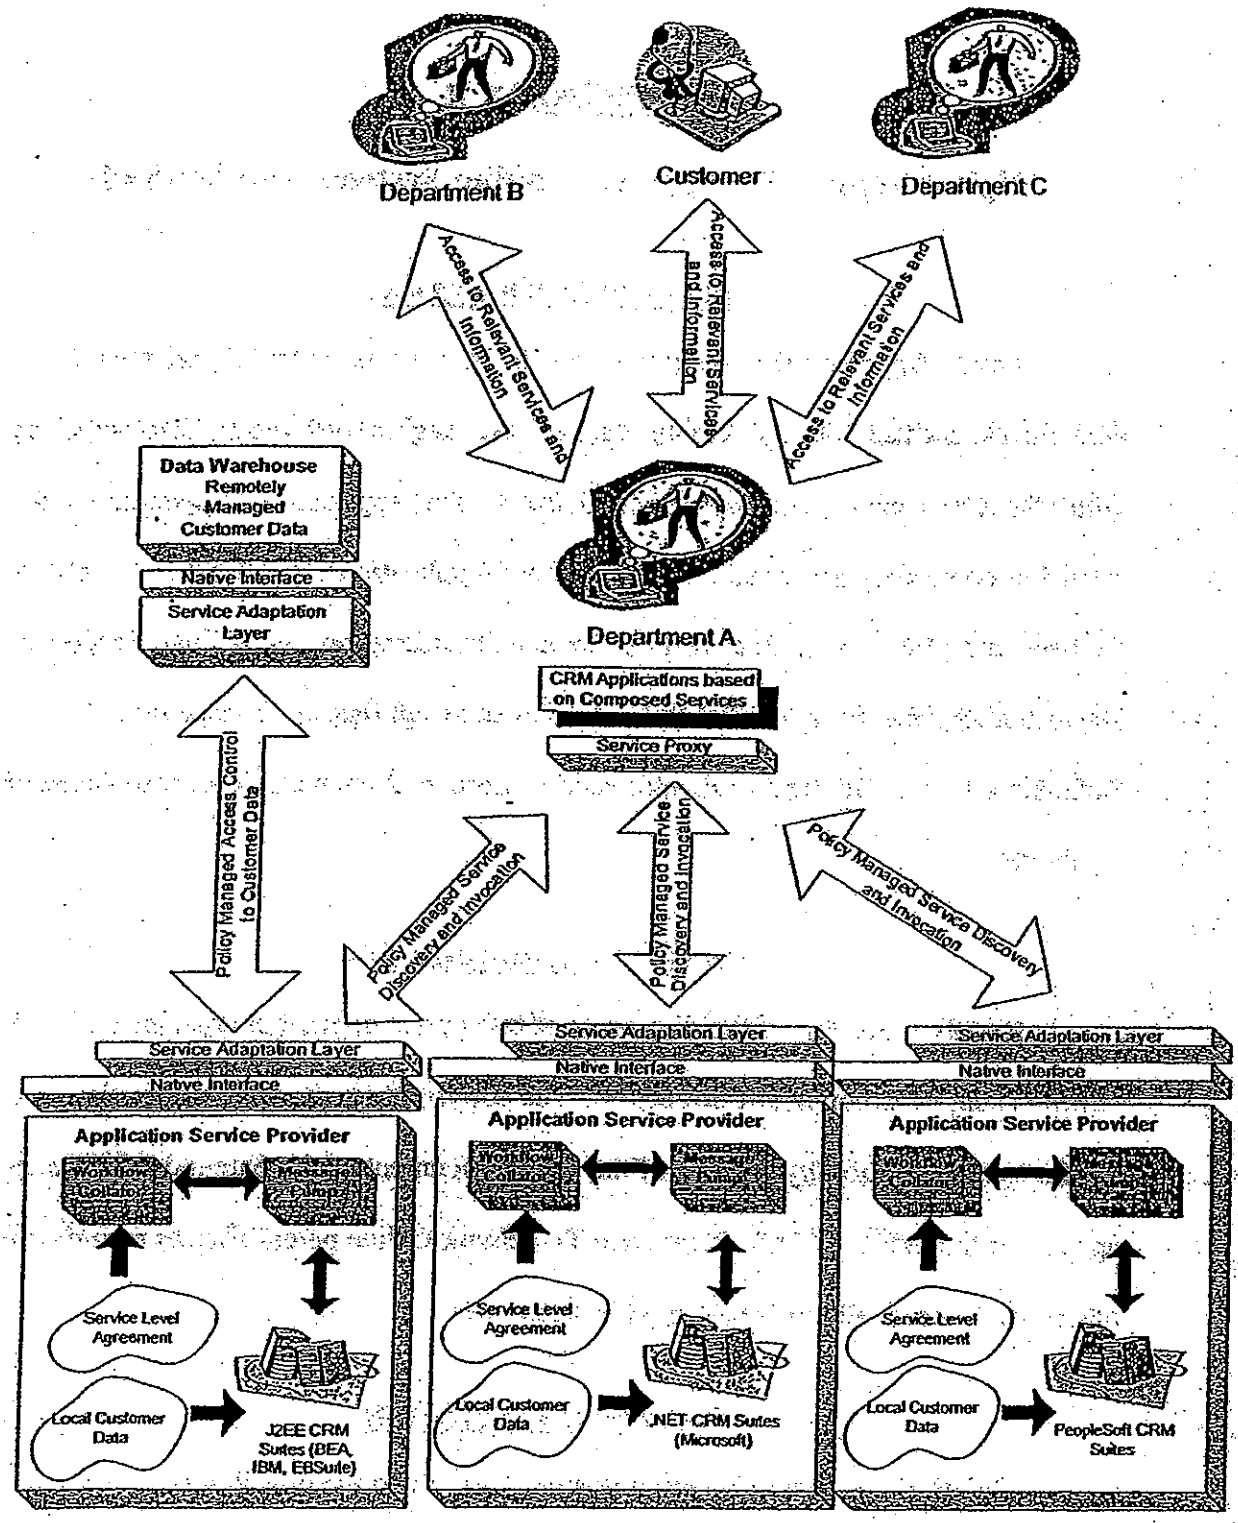
(part of Appendix A)



Policy Managed Access

Fig. 14

(part of Appendix A)



Exposing CRM Services with MediaDrive

Fig. 15

(part of Appendix A)

APPENDIX B

Digital Rights Management Engine Systems and Methods

RELATED APPLICATIONS

[0627] Appendix B corresponds to U.S. Provisional Patent Application No. 60/504,524, entitled "Digital Rights Management Engine Systems and Methods," by Gilles Boccon-Gibod, filed September 15, 2003. This application (Appendix B) is related to commonly-assigned U.S. Provisional Application No. 60/476,357, entitled Systems and Methods for Peer-to-Peer Service Orchestration, by William Bradley and David Maher, filed June 5, 2003 ("the Bradley et al. application"; Appendix A). References to "the Bradley et al. application," here, in Appendix B, refer to Appendix A, above.

COPYRIGHT AUTHORIZATION

[0628] A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure, as it appears in the WIPO or U.S. Patent and Trademark Office patent files or records, but otherwise reserves all copyright rights whatsoever.

FIELD OF THE INVENTION

[0629] The present invention relates generally to the distribution and use of digital information and services. More specifically, systems and methods are disclosed for providing, supporting, and/or using a digital rights management engine.

BACKGROUND

[0630] Networks such as the Internet have become the predominant medium for digital content delivery. The emergence of standards for web services promises to accelerate this trend, enabling companies to provide sophisticated services that interoperate with multiple software platforms and cooperate with other web services via standardized mechanisms.

SUMMARY OF THE INVENTION

[0631] Embodiments of the present invention provide a digital rights management engine that leverages the web services model. In contrast to many conventional digital rights management (DRM) systems, which require relatively sophisticated and heavyweight client-side engines to handle protected content, the present invention enables client-side DRM engines to be relatively simple, enforcing the governance policies set by richer policy management systems operating at the service level. Embodiments of the present invention also provide increased flexibility in the choice of media formats and cryptographic protocols, and can facilitate interoperability between DRM systems.

[0632] Preferred embodiments of the present invention provide a simple, open, and flexible client-side DRM engine that can be used to build powerful DRM-enabled applications. In a preferred embodiment, the DRM engine is designed to integrate easily into a web services environment, such as that described in the Bradley et al. application, and into virtually any host environment or software architecture. In a preferred embodiment, the DRM engine is independent of particular media formats and cryptographic protocols, allowing designers the flexibility to use standardized or proprietary technologies as required. In addition, the governance model used by a

preferred embodiment is relatively simple, yet can be used to express sophisticated relationships and business models.

[0633] Preferred embodiments of the present invention can be used to achieve some or all of the following goals:

[0634] *Simplicity.* In a preferred embodiment, the DRM engine uses a minimalist stack-based Virtual Machine (VM) to execute control programs (e.g., programs that enforce governance policies). For example, in one embodiment the VM may consist of only a few pages of code.

[0635] *Modularity.* In a preferred embodiment, the DRM engine is designed to function as a single module integrated into a larger DRM-enabled application. Many of the functions that were once performed by monolithic DRM kernels (such as cryptography services) can be requested from the host environment, which may provide these services to other code modules. This allows designers to incorporate standard or proprietary technologies with relative ease.

[0636] *Flexibility.* Because of its modular design, preferred embodiments of the DRM engine can be used in a wide variety of software environments, from embedded devices to general-purpose PCs.

[0637] *Open.* Preferred embodiments of the DRM engine are suitable for use as reference software, so that code modules and APIs can be implemented by users in virtually any programming language and in systems that they control completely. In a preferred embodiment, the system does not force users to adopt particular content formats or restrict content encoding.

[0638] *Semantically Agnostic.* In a preferred embodiment, the DRM engine is based on a simple graph-based model that turns authorization requests into queries

about the structure of the graph. The vertices in the graph represent entities in the system, and directed edges represent relationships between these entities. However, the DRM engine does not need to be aware of what these vertices and edges represent in any particular application.

[0639] *Seamless Integration with Web Services.* The DRM client engine can use web services in several ways. For example, vertices and edges in the authorization graph can be dynamically discovered through services. Content and content licenses may also be discovered and delivered to the DRM engine through sophisticated web services. Although the DRM engine of the preferred embodiment can be configured to leverage web services in many places, its architecture is independent of web services and it can be used as a stand-alone client-side DRM kernel.

[0640] *Simplified Key Management.* In a preferred embodiment, the authorization graph topology can be reused to simplify the derivation of content protection keys without requiring cryptographic retargeting. The key derivation method is an optional but powerful feature of the DRM engine—the system can also, or alternatively, be capable of integrating with other key management systems.

[0641] *Separation of Governance, Encryption, and Content.* In a preferred embodiment, the controls that govern content are logically distinct from the cryptographic information used to enforce the governance. Similarly, the controls and cryptographic information are logically distinct from content and content formats. Each of these elements can be delivered separately or in a unified package, thus allowing a high degree of flexibility in designing a content delivery system.

[0642] Preferred embodiments of the present invention can be used to achieve some or all of the foregoing goals; however, it should be appreciated that the present invention can also be practiced by systems that do not achieve these goals.

[0643] In a preferred embodiment, the DRM engine includes a virtual machine (VM) designed to determine whether certain actions on protected content are permissible. This Control VM can be implemented as a simple stack-based machine with a minimal set of instructions. In one embodiment, it is capable of performing logical and arithmetic calculations, as well as querying state information from the host environment to check parameters such as system time, counter state, and so forth.

[0644] In a preferred embodiment, the DRM engine utilizes a graph-based algorithm to verify relationships between entities in a DRM value chain. FIG. 1 shows an example of such an authorization graph. As shown in FIG. 1, the graph comprises a collection of nodes, connected by links. In a preferred embodiment, the semantics of the links may vary in an application-specific manner. For example, the directed edge from the *Mac* vertex to the *Knox* vertex may mean that Knox is the owner of the Mac. The edge from *Knox* to *Public Library* may indicate that Knox is a member of the Public Library. In a preferred embodiment, the DRM engine does not impose or interpret these semantics—it simply ascertains the existence or non-existence of paths within the graph.

[0645] For example, a content owner may create a control program to be interpreted by the Control VM that allows a particular piece of music to be played if the consuming device is owned by a member of the Public Library and is RIAA-approved. When the Control VM running on the device evaluates this control program, the DRM engine determines whether links exist between *Portable Device*

and *Public Library*, and between *Portable Device* and *RIAA Approved*. The edges and vertices of the graph may be static and built into devices, or may be dynamic and discovered through services communicating with the host application.

[0646] By not imposing semantics on the nodes and links, the DRM engine of the preferred embodiment enables great flexibility. The system can be adapted to many usage models, from traditional delegation-based policy systems to authorized domains and personal area networks.

[0647] In a preferred embodiment, the DRM client can also reuse the authorization graph for content protection key derivation. System designers may chose to allow the existence of a link to also indicate the sharing of certain cryptographic information. In such cases, the authorization graph can be used to derive content keys without explicit cryptographic retargeting to consuming devices.

[0648] Preferred embodiments of the DRM engine are modular, thereby enabling integration into many different devices and software environments. FIG. 2 illustrates a typical integration of the DRM client engine into a content consumption device. In a system such as that shown in FIG. 2, a typical content consumption event may proceed as follows:

[0649] The Host Application 202 receives a request to access a particular piece of content, typically through its User Interface 204. The Host Application 202 sends the request, along with all relevant DRM engine objects (preferably opaque to the Host Application) to the DRM engine 206. The DRM engine 206 makes requests for additional information and cryptographic services to the Host Services module 208 through well-defined interfaces. For example, the DRM engine 206 may ask the Host Services 208 whether a particular link is trusted, or may ask that certain objects

be decrypted. Some of the requisite information may be non-local, in which case the Host Services 208 can request the information from networked services through a service access point 214.

[0650] Once the DRM engine 206 has determined that a particular operation is permitted, it indicates this and returns any required cryptographic keys to Host Services 208, which initiates Media Rendering 210 (e.g., playing the content through speakers, displaying the content on a screen, etc.), coordinated with Cryptography Services 212 as needed.

[0651] The system architecture shown in FIG. 2 is a simple example of how the DRM engine can be used in applications, but it is only one of many possibilities. For example, in other embodiments the DRM engine can be integrated into packaging applications under the governance of relatively sophisticated policy management systems.

[0652] It should be appreciated that the disclosed embodiments can be implemented in numerous ways, including as processes, apparatuses, systems, devices, methods, or computer readable media. These and other features, advantages, and embodiments of the present invention will be illustrated in more detail by the following detailed description and the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

[0653] Embodiments of the present invention will be readily understood by referring to the following detailed description in conjunction with the accompanying drawings, wherein like reference numerals designate like structural elements, and in which:

[0654] FIG. 1 shows an authorization graph in accordance with an embodiment of the present invention.

[0655] FIG. 2 shows a digital rights management system in accordance with embodiments of the present invention.

[0656] FIG. 3 shows various objects that a DRM engine can use for content protection and governance in preferred embodiments of the present invention.

[0657] FIG. 4 illustrates node and link objects in accordance with embodiments of the present invention.

[0658] FIG. 5 illustrates the integration of a control virtual machine and a DRM engine in accordance with embodiments of the present invention.

[0659] FIG. 6 illustrates a code module in accordance with an embodiment of the present invention.

[0660] FIG. 7 illustrates the used of a node graph to facilitate content encryption and decryption.

DETAILED DESCRIPTION

[0661] Embodiments of a novel, digital rights management engine are described herein. As one of ordinary skill in the art will appreciate, the engine itself is novel, as are many of its features, aspects, components, and applications. A detailed description of the inventive body of work is provided below. While this description is provided in conjunction with several embodiments, it should be understood that the invention is not limited to any one embodiment, but instead encompasses numerous alternatives, modifications, and equivalents. For example, while some embodiments are described in the context of consumer-oriented content and applications, those skilled in the art will recognize that the disclosed systems and methods are readily

adaptable for broader application. For example, without limitation, embodiments of the present invention could be readily applied in the context of enterprise content and applications. In addition, while numerous specific details are set forth in the following description in order to provide a thorough understanding of the present invention, the present invention may be practiced without some or all of these details. Moreover, for the purpose of clarity, certain technical material that is known in the art related to the invention has not been described in detail in order to avoid unnecessarily obscuring the invention.

[0662] In a preferred embodiment, the DRM client engine operates on a group of basic objects or building blocks. FIG. 3 provides a high level illustration of these objects. As shown in FIG. 3, in a preferred embodiment the data represented by a Content object 302 is encrypted using a key. That key is represented by a ContentKey object 304, and the binding between the content and the key is represented by a Protector object 306. The rules that govern the use of the key to decrypt the content are represented by a Control object 308, and the binding between the ContentKey 304 and the Control 308 is represented by a Controller object 310. Compliant systems make use of the content decryption key under governance of the rules expressed by byte code in the Control object 308.

[0663] *Content Object 302:* In a preferred embodiment, Content object 302 is an "external" object. Its format and storage are not under the control of the DRM engine, but rather are under the control of the content management subsystem of the Host Application (for instance, it could be an MP4 movie file, an MP3 music track, etc.). However, in a preferred embodiment the format of the content is such that it allows the association of an ID with the content payload. The content payload is

encrypted in a format-dependent manner (e.g., with a symmetric cipher, such as AES).

[0664] *ContentKey Object 304:* In a preferred embodiment, ContentKey object 304 includes a unique encryption key and associated ID. The purpose of the ID is to enable Protector objects 306 and Controller Objects 310 to make references to ContentKey objects 304. The actual key data encapsulated in the ContentKey object 304 is itself encrypted so that it can be read only by the recipients that are authorized to decrypt the content. The ContentKey Object 304 also preferably specifies the cryptosystem that was used to encrypt the key data. The cryptosystem used to protect the content key data is called the Key Distribution System. Different Key Distributions Systems can be used.

[0665] *Protector Object 306:* In a preferred embodiment, Protector objects 306 contain the information that allows the DRM engine to find out which key was used to encrypt the data contained in Content objects 302. Protector objects 306 also contain information about which encryption algorithm was used to encrypt that data. A Protector object 306 contains one or more IDs that reference Content objects 302, and an ID that references the ContentKey object 304 that holds the key that was used to encrypt the data. In a preferred embodiment, if the Protector 304 points to more than one Content object 302, all of those Content objects represent data that has been encrypted using the same encryption algorithm and the same key.

[0666] *Control Object 308:* Control objects 308 contain information that allows the DRM engine to make decisions regarding whether certain actions on content should be permitted when requested by the Host Application. The rules that govern the use of content keys are encoded in Control objects 308 as control byte

code. Control objects 308 also contain unique IDs so that they can be referenced by Controller objects 310. In a preferred embodiment, Control objects 308 are signed, so that the DRM engine can verify that the control byte code is valid and trusted before it is used to make decisions. The validity of a Control object 308 can also optionally be derived through the verification of a secure hash contained in a Controller object 310, when that information is available.

[0667] *Controller Object 310:* Controller objects 310 contain information that allows the DRM engine to find out which Control object(s) 308 govern the use of one or more keys represented by ContentKey objects 304. In a preferred embodiment, Controller objects 310 also contain hash values for the key data contained in the ContentKey objects 304 that they reference, so that the binding between the key data and the ContentKey object 304 cannot be easily tampered with. Controller objects 310 may also contain information that facilitates the verification of the integrity of the Control objects 308 with which they are associated. Controller objects 310 are preferably signed, so that the DRM engine can trust the validity of the binding between the ContentKey 304 and the Control object 308 that governs it, as well as the validity of the binding between the ContentKey ID and the actual key data. In embodiments where the hash of the Control object 308 referenced by the Controller object 310 is included in the Controller object 310, the validity of the Control object 308 can be derived without having to separately verify the signature of the Control object 308.

[0668] As previously indicated in connection with FIG. 1, in a preferred embodiment, the DRM engine utilizes a graph-based algorithm to verify relationships between entities in a DRM value chain. As shown in FIG. 1, the graph comprises a

collection of nodes, connected by links. Node objects represent entities in a DRM Profile. In a preferred embodiment, the DRM engine does not have implicit or explicit semantics for what the Node objects represent.

[0669] A given deployment (DRM Profile) of a system using the DRM engine will define what types of principals exist, and what roles and identities different Node objects represent. As shown in FIG. 4, that semantic information is typically expressed using Attributes of the Node object. Link objects represent relationships between nodes. Link objects can also optionally contain some cryptographic data that allows the DRM engine to use the Link for ContentKey derivation computations. Just as for Nodes, in a preferred embodiment the DRM engine does not have implicit or explicit semantics for what a Link relationship means. Depending on what the *from* and *to* Nodes of the link represent in a given DRM Profile, the meaning of the link relationship can express membership, ownership, association, and/or many other types of relationships. For example, in a typical DRM Profile, some Node objects might represent users, other Nodes might represent devices, and other Nodes might represent user groups. In such a context, links between devices and users might represent ownership relationships, and links between users and user groups might represent membership relationships. FIG. 4 shows examples of node and link objects.

[0670] *Node 402:* Node objects represent entities in the system. A Node object's Attributes define certain aspects of what the Node object represents, such as the role or identity represented by the Node object in the context of a DRM Profile. In a preferred embodiment, a Node object also has a Confidentiality asymmetric key pair that is used for targeting confidential information to the subsystems that have access to the confidential parts of the Node object (typically, the entity represented by

the Node, or some entity that is responsible for managing the Node). Confidential information targeted at a Node is preferably encrypted with that Node's Confidentiality Public Key. The Content Protection asymmetric key pair and the Content Protection symmetric key are used in conjunction with Link objects in embodiments that use the optional ContentKey derivation system for ContentKey distribution.

[0671] *Link 404:* In a preferred embodiment, link objects 404 are signed assertions that there exists a directed edge between the '*To*' and '*From*' nodes in the graph whose vertices are the node objects. For a given set of Nodes and Links, there is a *Path* between a Node X and a Node Y if there is a directed path between the Node X vertex and the Node Y vertex in the graph. When there is a Path between Node X and Node Y, Node Y is said to be *Reachable* from Node X. The assertions represented by Link objects are thus used to express which Nodes are reachable from other Nodes. The Controls that govern Content objects can use as conditions for allowing certain Actions to take place the Reachability of certain Target Nodes from the Node that represents the entity that is requesting the Action on the Content object. For example, if Node D represents a device that wants to perform the 'Play' Action on a Content object, a Control that governs this Content object can test if a certain Node U representing a certain User is Reachable. To determine if Node U is Reachable, the DRM engine checks if there exists a set of Link objects that can establish a Path between Node D and Node U.

[0672] The DRM engine preferably verifies Link objects before using them to ascertain the existence of Paths in the Node graph. Depending on the specific features of the certificate system used to sign Link objects (e.g., x509v3), Link objects can be

given limited lifetimes, be revoked, etc. Also, in a preferred embodiment the policies that govern which entities can sign Link objects, which Link objects can be created, and the lifetime of Link objects are not directly handled by the DRM engine. Those policies exist outside the DRM engine, and will typically leverage the Node's Attribute information. For example, an entity that has been granted the ability to create Link objects that link Device Nodes and User Nodes under a certain policy will probably check that it only creates links between Node objects that have attributes that indicate that they are indeed representing a Device, and Nodes that have attributes indicating that they represent a User.

[0673] Finally, the Link object can contain cryptographic data that enables the use of the Node's Content Protection Keys for key distribution. In some embodiments, that cryptographic data contains, in addition to metadata, the Private and/or Symmetric Content Protection keys of the 'From' Node, encrypted with the Content Protection public key and/or the Content Protection symmetric key of the 'To' Node. Additional information on the relationship and operation of nodes and links is provided below under the heading "Nodes, Links, and Content Protection."

[0674] Since, in a preferred embodiment, the DRM engine does not specify, implicitly or explicitly, the semantics attached to these basic objects, systems that use the DRM engine will participate in one or more usage patterns for those objects, which will put those objects in a semantic context. These semantic contexts will be referred to as *DRM Profiles*. In that sense, the DRM engine Objects can be considered to be DRM building blocks used to construct DRM Systems.

[0675] The following paragraphs illustrate typical types of DRM Profiles that can be built using the DRM engine Objects.

[0676] Example 1: Users, PCs and Devices

[0677] In this DRM Profile example, the following entities are defined: Users (individuals who are consuming content), PCs (computers running software applications that use content) and Devices (hardware and software combinations that use (e.g., play) content). In this example profile, any User can own an unlimited number of Devices, and can be associated with up to a predefined number of PCs (e.g., 4 PCs). Content can be governed by rules that allow the content to be bound to a User (e.g., a rule can check that the PC or the Device that plays the content belongs to, or is associated with, a given User).

[0678] Users, PCs, and Devices are each assigned a Node object. These Node objects have 'Type' Attributes that indicate whether they represent a User, a PC or a Device.

[0679] A back-end system manages User identities, and creates User Nodes as well as Link objects between PC Nodes and User Nodes. Before a User Node object is sent to a PC software instance, the back-end server will enforce the "4 PCs per User" policy by looking up the User in a User information database and determining if the User that wished to be associated with a given PC has reached the limit of 4. In some embodiments, the system may provide a way for users to remove existing PC-User associations so that they can change the set of PCs with which they are associated as time passes. The PC software will keep the PC-to-User Link objects, and can use them for as long as they remain valid.

[0680] The PC software is given the ability to associate Devices to Users. This means that the PC software is given the responsibility to manage the User Node object in addition to the PC Node object. That software has access to the confidential part of

those nodes. When a User wants to be associated with a Device (because she wants to play her content on this Device), a Link object from the Device Node and the User Node needs to exist. If that Link does not exist, the PC software will create it, and make it available to the Device. The Device will keep that Link object and use it for as long as it remains valid.

[0681] To bind content to a User, a Packager application will choose an ID, or use an existing ID for the content. The packager creates an encryption key and an associated ContentKey object, as well as a Protector object to bind the Content object and the ContentKey object. The packager then creates a Control object with a Control Program (preferably compiled in Control VM Byte Code) that will allow the "Play" action to take place if and only if the User Node is Reachable from the PC or Device Node that is requesting the performance of the Action. Typically, the Control, Controller, Protector and ContentKey objects will be embedded in the packaged content if appropriate, so that the PCs and Devices will not need to obtain them separately.

[0682] When a Device or a PC wants to play content, the DRM engine will find the Protector object for the Content ID of the content, then the ContentKey object referenced by that Protector, then the Controller object that references that ContentKey object, and finally the Control object referenced by that Controller. The DRM engine will execute the Control Program of the Control object, which will check if the User Node is Reachable. If the Device or PC Node has the necessary Link objects to verify that there exists a path between their Node and the User Node, then the Control Program will allow the use of the key represented in the ContentKey

object, and the Media Rendering engine of the Device or PC will decrypt and render (e.g., play) the content.

[0683] Example 2: Temporary Login

[0684] A second example makes use of a system almost identical to that used in connection with the previous example, with the addition of a new feature: the policy that governs the creation of Link objects between PC Nodes and User Nodes allows for a temporary login of no more than, e.g., 12 hours, as long as the User does not already have a temporary login on another PC. The purpose of this feature is to allow, for instance, a User to take his content to a friend's PC, Login to that PC for a period of time, and play his content on the friend's PC. When the User wishes to log in to his friend's PC, he enters his login credentials (this could be a username/password, a mobile phone authentication, a smartcard, or any authentication system allowed under the policy of that system), and the back-end system checks the attributes of the User Node and PC Node for which the Link is requested and verifies that there is no active temporary login Link that is still valid. If these conditions are met, the back-end service creates a Link object linking that PC and the User, with a validity period limited to the requested login duration (e.g., less than 12 hours) to comply with the policy. Having that Link now allows the PC to play the User's content until the Link expires.

[0685] ***Control Virtual Machine***

[0686] The Control virtual machine (or "Control VM") is a virtual machine used by a preferred embodiment of the DRM engine to execute Control Programs that govern access to content. The following is a description of where the Control VM fits

into the DRM engine architecture, as well as some of the basic elements of the VM, followed by more details about the Memory Model and Instruction Set.

[0687] In a preferred embodiment, the Control VM is a small-footprint virtual machine that is designed to be easy to implement using various programming languages. In a preferred embodiment, the Control VM is based on a stack-oriented instruction set that is designed to be minimalist in nature, without much concern for execution speed or code density. However, it will be appreciated that if execution speed and/or code density were issues in a given application, conventional techniques (e.g., data compression) could be used to improve performance.

[0688] In a preferred embodiment, the Control VM is suitable as a target for low or high level programming languages, and supports languages such as assembler, C and FORTH. Compilers for other languages, such as Java or custom languages, could also be implemented with relative ease.

[0689] In a preferred embodiment, the Control VM is designed to be hosted within a Host Environment, as opposed to being run directly on a processor or in silicon. The natural Host Environment for the Control VM is the DRM engine. FIG. 5 illustrates how the Control VM is integrated with the DRM engine in preferred embodiments of the present invention.

[0690] The Control VM runs within the context of its Host Environment, which implements some of the functions needed by the VM as it executes programs. Typically, the Control VM runs within the DRM engine, which implements its Host Environment.

[0691] The VM runs programs by executing instructions stored in Code Modules. Some of these instructions can make calls to functions implemented outside

of the program itself by making a System Call. System Calls are either implemented by the Control VM itself, or delegated to the Host Environment.

[0692] Some basic elements of preferred embodiments of the Control VM will now be described.

[0693] *Execution Model*

[0694] The Control VM executes instructions stored in Code Modules as a stream of Byte Code loaded in memory. The VM maintains a virtual register called the Program Counter (PC), which is incremented as instructions are executed. The VM executes each instruction, in sequence, until the `OP_STOP` instruction is encountered, an `OP_RET` instruction is encountered with an empty call stack, or an exception occurs. Jumps are specified either as a relative jump (specified as a byte offset from the current value of PC), or as an absolute address.

[0695] *Memory Model*

[0696] In preferred embodiments, the Control VM has a relatively simple memory model. The VM memory is separated into a Data Segment (DS) and a Code Segment (CS). The Data Segment is a single, flat, contiguous memory space, starting at address 0. The Data Segment is typically an array of bytes allocated within the heap memory of the Host Application or Host Environment. For a given VM implementation, the size of the memory space is preferably fixed to a maximum, and attempts to access memory outside of that space will cause faults and terminate program execution. The Data Segment is potentially shared between several Code Modules concurrently loaded by the VM. The memory in the Data Segment can be accessed by memory-access instructions, which can be either 32-bit or 8-bit accesses. 32-bit memory accesses are done using the big-endian byte order. No assumptions

are made with regard to alignment between the VM-visible memory and the host-managed memory (host CPU virtual or physical memory).

[0697] In one embodiment, the Code Segment is a flat, contiguous memory space, starting at address 0. The Code Segment is typically an array of bytes allocated within the heap memory of the Host Application or Host Environment.

[0698] In a preferred embodiment, the VM may load several code modules, and all of the code modules may share the same Data Segment (each module's data is preferably loaded at a different address), but each has its own Code Segment (e.g., it is preferably not possible for a jump instruction from one Code Module to cause a jump directly to code in another Code Module).

[0699] *Data Stack*

[0700] In a preferred embodiment, the VM has a notion of a data stack, which represents 32-bit data cells stored in the Data Segment. The VM maintains a virtual register call the Stack Pointer (SP). After reset, SP points to the end of the Data Segment, and the stack grows downward (when data is pushed onto the data stack, the SP registers is decremented). The 32-bit values on the stack are interpreted either as 32-bit addressed, or 32-bit signed integers, depending on the instruction referencing the stack data.

[0701] *Call Stack*

[0702] In a preferred embodiment, the VM manages a Call Stack for making nested subroutine calls. The values pushed on this stack cannot be read or written directly by the memory-access instructions, but are used indirectly by the VM when executing OP_JSR and OP_RET instructions. For a given VM Profile, the size of

this return address stack is preferably fixed to a maximum, which will allow a certain number of nested calls that cannot be exceeded.

[0703] *Instruction Set*

[0704] In a preferred embodiment, the Control VM uses a very simple instruction set. Even with a limited number of instructions; however, it is still possible to express simple programs. In a preferred embodiment, the instruction set is stack-based: except for the OP_PUSH instruction, none of the instructions have direct operands. Operands are read from the data stack, and results are pushed onto the data stack. In a preferred embodiment, the VM is a 32-bit VM: all the instructions operate on 32-bit stack operands, representing either memory addresses or signed integers. Signed integers are represented using a 2s complement binary encoding.

[0705] An illustrative instruction set used in a preferred embodiment is shown below:

OP_CODE	Name	Operands	Description
OP_PUSH	Push Constant	N (direct)	Push a constant on the stack
OP_DROP	Drop		Remove top of stack
OP_DUP	Duplicate		Duplicate top of stack
OP_SWAP	Swap		Swap top two stack elements
OP_ADD	Add	A, B	Push the sum of A and B (A+B)
OP_MUL	Multiply	A, B	Push the product of A and B (A*B)
OP_SUB	Subtract	A, B	Push the difference between A and B (A-B)
OP_DIV	Divide	A, B	Push the division of A by B (A/B)
OP_MOD	Modulo	A, B	Push A modulo B (A%B)
OP_NEG	Negate	A	Push the 2s complement negation of A (-A)
OP_CMP	Compare	A	Push -1 if A negative, 0 if A is 0, and 1 is a positive
OP_AND	And	A, B	Push bit-wise AND of A and B (A & B)

OP_OR	Or	A, B	Push the bit-wise OR of A and B (A | B)
OP_XOR	Exclusive Or	A, B	Push the bit-wise eXclusive OR or A and B (A ^ B)
OP_NOT	Logical Negate	A	Push the logical negation of A (1 if A is 0, and 0 if A is not 0)
OP_SHL	Shift Left	A, B	Push A logically shifted left by B bits (A << B)
OP_SHR	Shift Right	A, B	Push A logically shifted right by B bits (A >> B)
OP_JSR	Jump to Subroutine	A	Jump to subroutine at absolute address A
OP_JSRR	Jump to Subroutine (Relative)	A	Jump to subroutine at PC+A
OP_RET	Return from Subroutine		Return from subroutine
OP_BRA	Branch Always	A	Jump to PC + A
OP_BRP	Branch if Positive	A, B	Jump to PC+A if B > 0
OP_BRN	Branch if Negative	A, B	Jump to PC+A if B < 0
OP_BRZ	Branch if Zero	A, B	Jump to PC+A if B is 0
OP_JMP	Jump	A	Jump to A
OP_PEEK	Peek	A	Push the 32-bit value at address A
OP_POKE	Poke	A, B	Store the 32-bit value B at address A
OP_PEEKB	Peek Byte	A	Push the 8-bit value at address A
OP_POKEB	Poke Byte	A, B	Store the least significant bits of B at address A
OP_PUSHSP	Push Stack Pointer		Push the value of SP
OP_POPSP	Pop Stack Pointer	A	Set the value of SP to A
OP_CALL	System Call	A	Perform System Call with index A
OP_STOP	Stop		Terminate Execution

[0706] *Module Format*

[0707] In a preferred embodiment, code modules are stored in an Atom-based format that is essentially equivalent to the Atom structure used in the MPEG-4 File Format. An atom consists of 32 bits, stored as 4-octets in big-endian byte order, followed by a 4-octet type (usually octets that correspond to ASCII values of letters of the alphabet), followed by the payload of the Atom (size-8 octets).

[0708] An exemplary code module is shown in FIG. 6, and described below.

[0709] *pkCM Atom:* The pkCM Atom is the top-level Code Module Atom. It contains a sequence of sub-atoms.

[0710] *pkDS Atom:* The pkDS Atom contains a memory image that can be loaded into the Data Segment. The payload of the Atom is a raw sequence of octet values.

[0711] *pkCS Atom:* The pkCS Atom contains a memory image that can be loaded into the Code Segment. The payload of the Atom is a raw sequence of octet values.

[0712] *pkEX Atom:* The pkEX Atom contains a list of export entries. Each export entry consists of a name, encoded as an 8-bit name size followed by the characters of the name, including a terminating 0, followed by a 32-bit integer representing the byte offset of the named entry point (this is an offset from the start of the data stored in the pkCS Atom).

[0713] *Module Loader*

[0714] In a preferred embodiment, the Control VM is responsible for loading Code Modules. When a Code Module is loaded, the memory image encoded in the pkDS Atom is loaded at a memory address in the Data Segment. That address is chosen by the VM Loader, and is stored in the DS pseudo-register. The memory

image encoded in the pkCS Atom is loaded at a memory address in the Code Segment. That address is chosen by the VM Loader, and is stored in the CS pseudo-register.

[0715] *System Calls*

[0716] Control VM Programs can call functions implemented outside of their Code Module's Code Segment. This is done through the use of the `OP_CALL` instruction, that takes an integer stack operand specifying the System Call Number to call. Depending on the System Call, the implementation can be a Control VM Byte Code routine in a different Code Module (for instance, a library of utility functions), directly by the VM in the VM's native implementation format, or delegated to an external software module, such as the VM's Host Environment.

[0717] In a preferred embodiment, there are several System Call Numbers that are specified:

[0718] *SYS_NOP = 0*: This call is a no-operation call. It just returns (does nothing else). It is used primarily for testing the VM.

[0719] *SYS_DEBUG_PRINT = 1*: Prints a string of text to a debug output. This call expects a single stack argument, specifying the address of the memory location containing the null-terminated string to print.

[0720] *SYS_FIND_SYSCALL_BY_NAME = 2:* Determines whether the VM implements a named System Call. If it does, the System Call Number is returned on the stack, otherwise the value -1 is returned. This call expects a single stack argument, specifying the address of the memory location containing the null-terminated System Call name that is being requested.

[0721] *System Call Numbers Allocation*

[0722] In a preferred embodiment, the Control VM reserves System Call Numbers 0 to 1023 for mandatory System Calls (System Calls that have to be implemented by all profiles of the VM).

[0723] System Call Numbers 16384 to 32767 are available for the VM to assign dynamically (for example, the System Call Numbers returned by SYS_FIND_SYSCALL_BY_NAME can be allocated dynamically by the VM, and do not have to be the same numbers on all VM implementations).

[0724] *Standard System Calls*

[0725] In a preferred embodiment, several standard System Calls are provided to facilitate writing Control Programs. Such standard system calls may include a call to obtain a time stamp from the host, a call to determine if a node is reachable, and/or the like. System calls preferably have dynamically determined numbers (e.g., their System Call Number can be retrieved by calling the SYS_FIND_SYSCALL_BY_NAME System Call with their name passed as the argument).

[0726] ***Links, Rules, and Content Encryption***

[0727] The following provides additional information on the relationship between principals, groups, rules, assertions, and key management in preferred embodiments of a DRM engine and/or system.

[0728] *Links and Nodes*

[0729] As previously described, a DRM system can be conceptualized as having a set of *nodes* and *links*. The nodes can represent different entities, but typically represent groups of identities (with some groups being just one identity), and

can be connected by links. A link can be viewed as an assertion that a node "belongs to" the group represented by the node to which it is linked.

[0730] The links are assertions, typically made by entities other than the user and/or DRM engine creator under governance of a policy. How the assertions are made and governed can be done in any suitable manner; however, the links are preferably trusted. In practice, this will generally mean that the representation of the assertion that a link exists will be signed so that the parts of the system that need to make decisions based on that link can do so in a trusted manner.

[0731] The relationship implied by the existence of a link between nodes is transitive. Links can be represented by directed edges in a graph of nodes, and one can say that node Na belongs to the group node Nx if there is a directed path between Na and Nx in the graph.

[0732] The following paragraphs explain how these graphs of nodes can be used to express Rule Conditions, as well as to manage encryption keys.

[0733] *Rules and Membership Conditions*

[0734] The system can be used to govern content through rules, expressed as control programs, that decide whether certain requested actions (or intents) are allowed or not, and, if they are, what the consequences and/or obligations are. In a preferred embodiment, rules are small control programs that take a few parameters as input (such as content identification information, time/date, the value of certain counters, ID of the principal requesting the action, etc.), and make decisions as to whether to allow certain actions to take place. If an action is allowed to take place, the program can also carry out the rule's stated obligations (e.g., updating a counter).

In some embodiments, the types of inputs and outputs that the program can deal with may be relatively limited.

[0735] A rule can express the fact that a condition for a principal being able to perform an action is the existence of a path (or set of paths) from the node representing the principal requesting the action to a target node, or set of target nodes, one for each path (e.g., to express a rule such as "the principal can play the content if it is a member of the group of subscribers to this content"). A link between two nodes Na and Nb is a trustable assertion that Na belongs to the group represented by Nb. If, at the point of enforcement of the rule (e.g., where the rule program is being run), it is possible to obtain all the links necessary to traverse the graph to all the required target nodes, then the condition is satisfied. This graph of nodes will sometimes be referred to as a *Membership Graph*, and links in that graph *Membership Links*.

[0736] *Content Encryption*

[0737] Compliant elements designed to work with the system will be able to understand and enforce the appropriate rules when they want to perform an action on a protected piece of content. However, non-compliant systems generally cannot be trusted to use content in this manner. Thus, it is desirable to protect the content using cryptographic techniques. In a preferred embodiment, content is encrypted with a content key Kc, typically unique to the content being protected. When a principal wants to perform an action on the content, it will first run the rule program, and get an answer yes/no. If the rule allows the principal to perform the action, the principal will need to obtain the content key Kc. In a preferred embodiment, there is a relationship between the rule system, the node/links graph, and the content protection system that

enables the principal to obtain the key in an efficient manner, without an external entity explicitly "targeting" Kc to the requesting node.

[0738] In a preferred embodiment, this is accomplished by conceptualizing the relationship between principals/objects as a graph of nodes. Some nodes may be the same nodes as the ones used in the Membership Graph, but that is not necessary. In simple systems, however, it is expected that this graph will be a subset of the Membership Graph. This graph will be called the *Protection Graph*.

[0739] Each node N in the system can have, amongst other attributes, an associated *Content Protection Key* Kc[N], or Key Pair. Other attributes might include a public/private key pair for signing, and potentially targeting information to the node. For purposes of illustration, the following example shows the use of a single key, however, it should be appreciated that the same principles would work with a key pair. In this graph, a link between two nodes Na and Nb is a trustable assertion that anyone who has Kc[Na] can compute Kc[Nb]. Typically, this means that the link will contain an expression of Kc[Nb] encrypted by Kc[Na]. This allows one to chain links between nodes and compute a target node's Kc.

[0740] The graph used to create paths to keys may be different from the graph used for expressing membership conditions, but in simple cases, the same graphs and links can be reused, thus creating a "parallel" or alternative use of the node graph: the same path that was traversed for checking a rule's condition is now traversed to compute a content protection key. Any node Na that has access to links that can be followed to reach node Nb can compute Kc[Nb], even if no explicit targeting of Kc[Nb] to Na has ever been made.

[0741] These keys will typically be used in the following manner: content Ci is encrypted with a content key K[Ci]. One or more rules can be created, which can check, amongst other things, the existence of paths between the node requesting an action and a list of "condition" target nodes representing groups. The rule(s) may include a program that performs the checking of the conditions (and possibly also the side effects of obligations), as well as a copy of K[Ci] protected by successive encryption with each of the Kc[Ni] node content encryption keys of the nodes to which the requesting node must belong.

[0742] When a node Na wants to perform an action on the content Ci, it will run the rule program that governs the content for that action. If the rule requires the membership of Na in one or more groups Np, Nq, etc., it will check that the necessary links to Np, Nq, etc. exist, and produce a decision as to whether the action is permitted or not. If the action is permitted, the system will determine if it has a set of links in the Protection Graph to nodes Nm, Nn, etc. whose content protection keys Kc[Nm], Kc[Nn], etc. need to be used in sequence to obtain the content key K[Ci], of which an encrypted copy can be found in the rule. K[Ci] will be obtained by decrypting this encrypted content key with the nodes' content encryption keys Kc[Nm], Kc[Nn], etc. in the order specified by the protection metadata for the content key. Again, in simple systems, the set of nodes Np, Nq, etc. will often be the same as Nm, Nn, etc.

[0743] FIG. 7 shows an example scenario that illustrates the dual use of the node graph. Referring to FIG. 7, iPod1 is a content playback device. Nip1 is the node that represents this device. Kip1 is the content encryption key associated with

Nip1. Gilles is the owner of iPod1, and Ng is the node that represents Gilles. Kg is the content encryption key associated with Ng.

[0744] PubLib is a Public Library. Npl represents the members of this library, and Kpl is the content encryption key associated with Npl. ACME represents all the ACME-manufactured Music Players. Namp represents that class of devices, and Kamp is the content encryption key associated with this group.

[0745] L1 is a link from Nip1 to Ng, which means that iPod1 belongs to Gilles. L2 is a link from Ng to Npl, which means that Gilles is a member of the Public Library. L3 is a link from Nip1 to Namp, which means that iPod1 is an ACME device.

[0746] C1 is a movie file that the Public Library makes available to its members. Kc1 is a key used to encrypt C1. GB[C1] is the governance information for C1 (e.g., rules and associated information used for governing access to the content). E(a,b) means 'b' encrypted with key 'a'.

[0747] For purposes of illustration, assume that it is desired to set a rule that says that a device can play the content C1 as long as (a) the device belongs to someone who is a member of the library and (b) the device is manufactured by ACME.

[0748] The content C1 is encrypted with Kc1. The rules program is created, as well as the encrypted content key RK[C1] = E(Kamp, E(Kpl, Kc1)). Both the rules program and RK[C1] can be included in the governance block for the content, GB[C1].

[0749] iPod1 receives C1 and GB[C1]. For example, both might be packaged in the same file, or received separately. iPod1 received L1 when Gilles first installed

his device after buying it. iPod1 received L2 when Gilles paid his subscription fee to the Public Library. iPod1 received L3 when it was manufactured (e.g., L3 was built in).

[0750] From L1, L2 and L3, iPod1 is able to check that Nip1 has a graph path to Ng (L1), Npl (L1+L2), and Namp (L3). iPod1 wants to play C1. iPod1 runs the rule found in GB[C1]. The rule can check that Nip1 is indeed an ACME device (path to Namp) and belongs to a member of the public library (path to Npl). Thus, the rule returns "yes", and the ordered list (Namp, Npl).

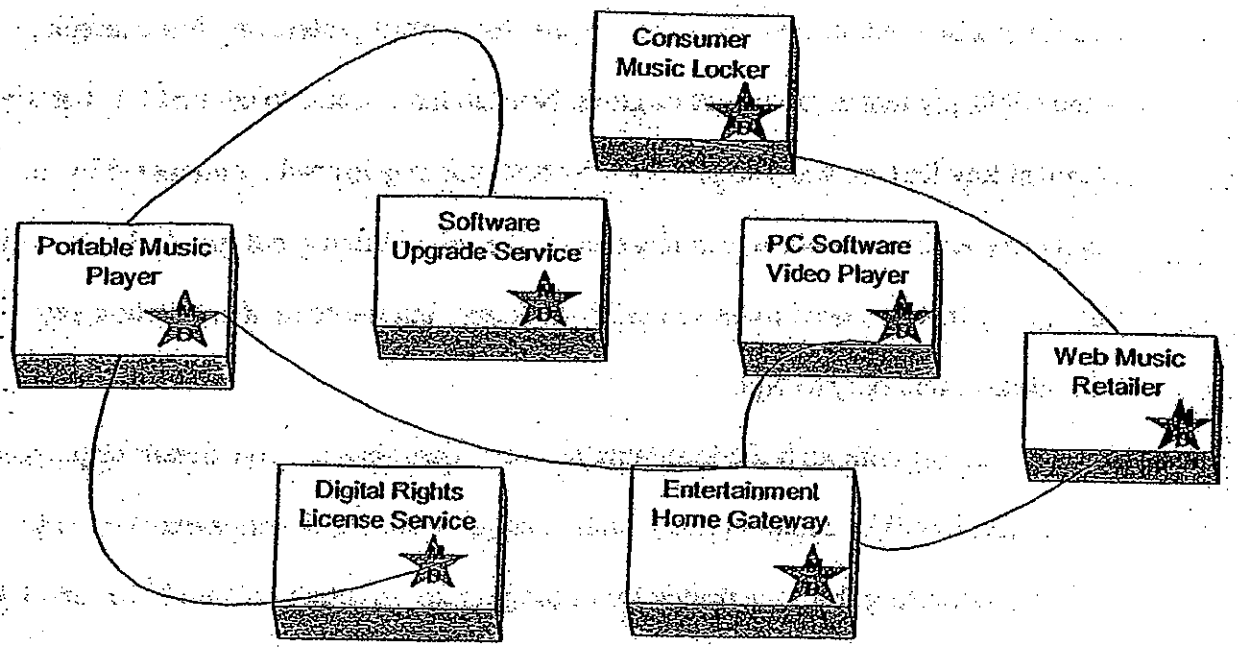
[0751] iPod1 uses L1 to compute Kg, and then L2 to compute Kpl from Kg. iPod1 also uses L3 to compute Kamp. iPod1 applies Kpl and Kamp to RK[C1] found in GB[C1], and computes Kc1. It then uses Kc1 to decrypt and play C1.

[0752] When node keys are symmetric keys, as in the previous examples, the content packager needs to have access to the keys of the nodes to which it wishes to "bind" the content. This can be achieved by creating a node that represents the packager, and a link between that node and the nodes to which it wishes to bind rules. This can also be achieved "out of band" through a service, for instance. But in some situations, it may not be possible, or practical to use symmetric keys. In that case, it is possible to assign a key pair to the nodes to which a binding is needed without shared knowledge. In that case, the packager would bind a content key to a node by encrypting the content key with the target node's public key. To obtain the key for decryption, the client would have access to the node's private key via a link to that node.

[0753] In the most general case, the nodes used for the rules and the nodes used for computing content encryption keys need not be the same. It is natural to use

the same nodes, since there is a strong relationship between a rule that governs content and the key used to encrypt it, but it is not necessary. In some systems, some nodes may be used for content protection keys that are not used for expressing membership conditions, and vice versa, and in some situations two different graphs of nodes can be used, one for the rules and one for content protection. For example, a rule could say that all members of group Npl can have access to content C1, but the content key Kc1 may not be protected by Kpl, but may instead by protected by the node key Kf of node Nf, which represents all public libraries, not just Npl. Or a rule could say that you need to be a member of Namp, but the content encryption key could be bound only to Npl.

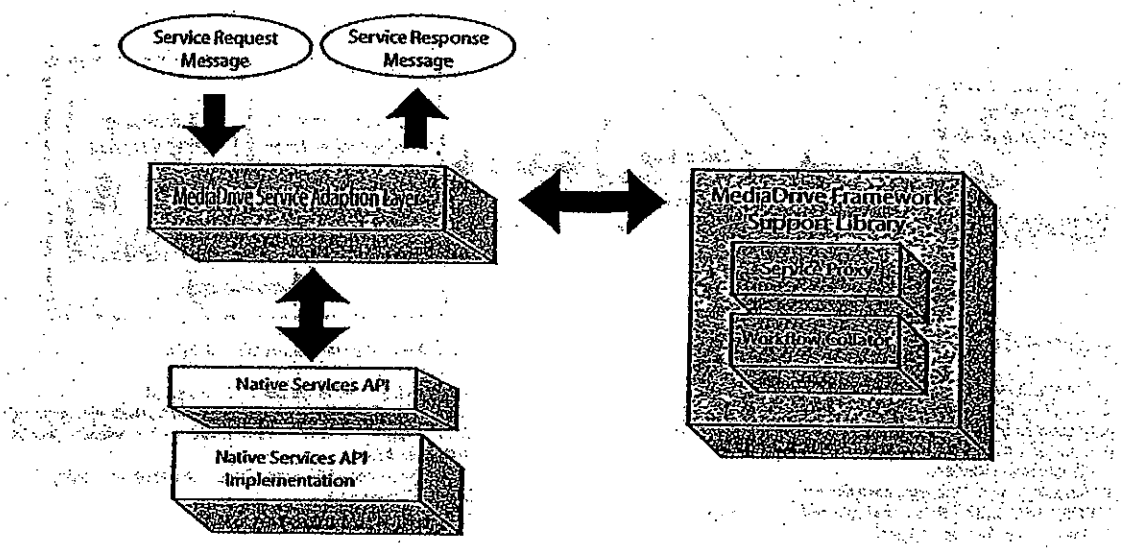
[0754] Although the foregoing has been described in some detail for purposes of clarity, it will be apparent that certain changes and modifications may be made without departing from the principles of the present invention. It should be noted that there are many alternative ways of implementing both the processes and apparatuses of the present invention. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the specific details given herein.



Example Embodiment of MediaDrive Framework

Fig. 1

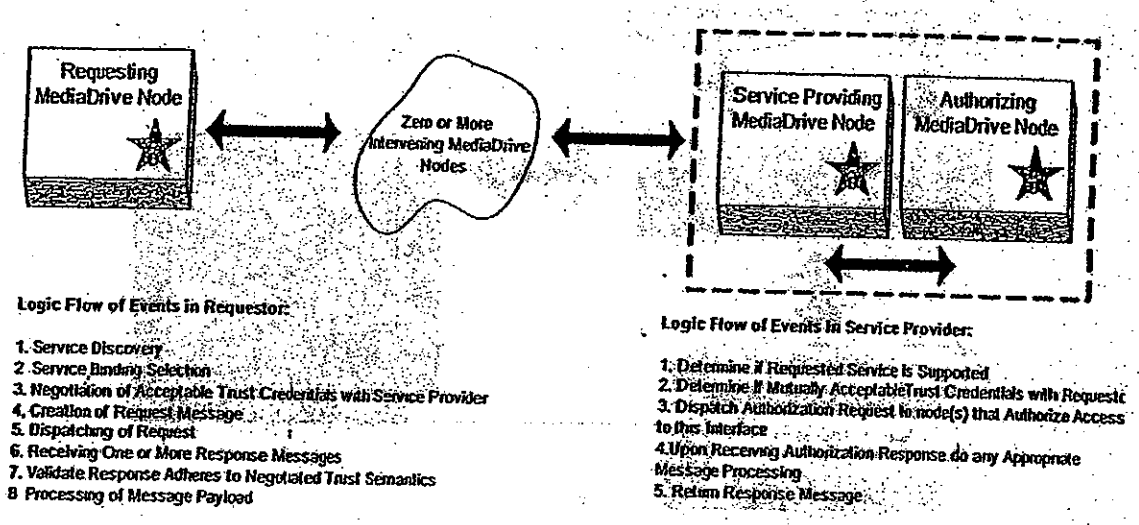
(part of Appendix B)



Conceptual View of a MediaDrive Node

Fig. 2

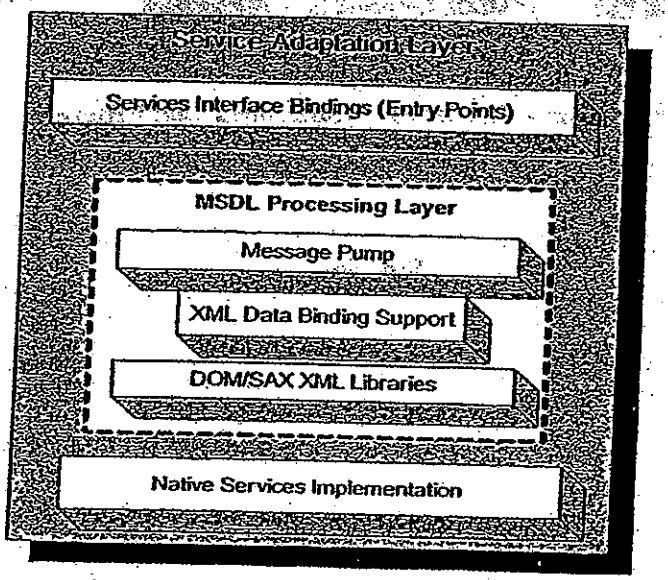
(part of Appendix B)



Illustrative Generic Interaction Pattern

Fig. 3

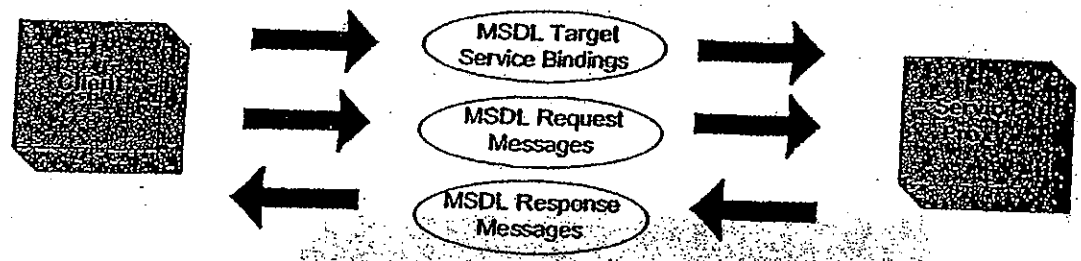
(part of Appendix B)



Service Adaptation Layer

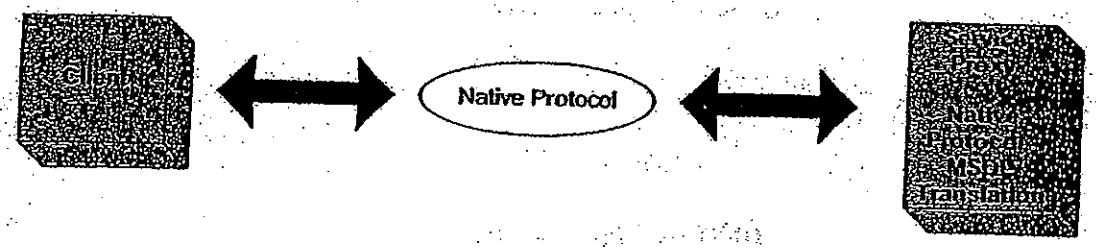
Fig. 4

(part of Appendix B)



Service Proxy – Client Side MSDL Interaction Pattern

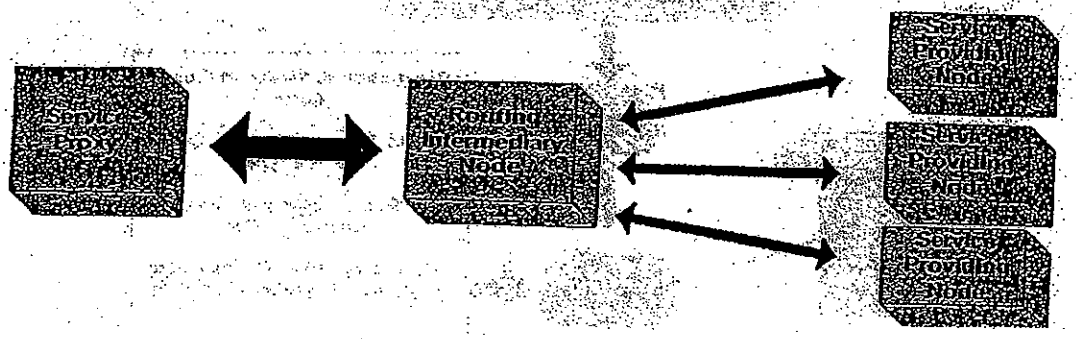
Fig. 5a



Service Proxy – Client-Side Native Interaction Pattern

Fig. 5b

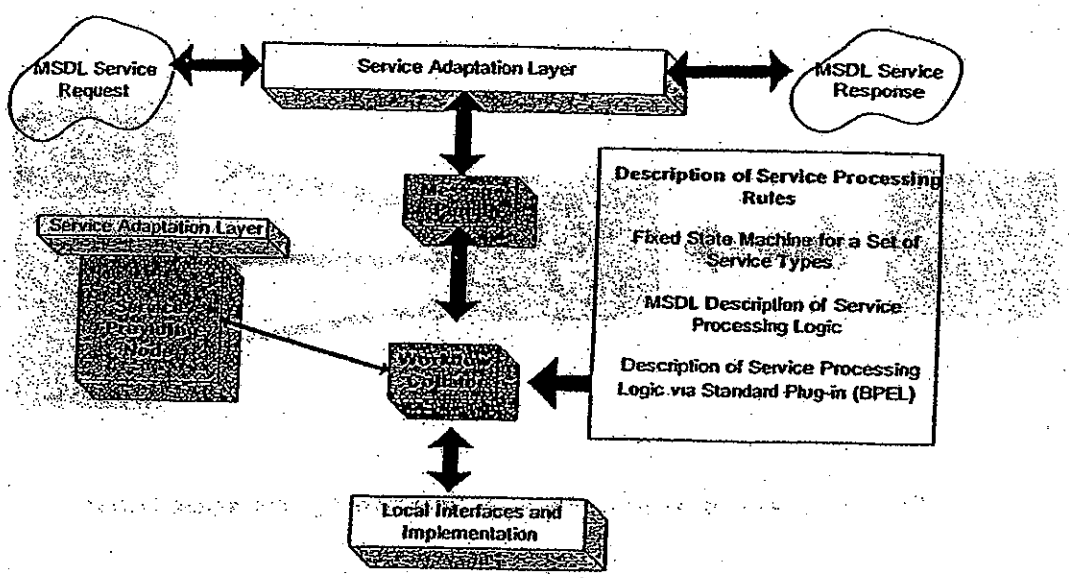
(part of Appendix B)



Service Proxy – Service-Side Point-to-Intermediary Interaction Pattern

Fig. 6

(part of Appendix B)



Interaction Pattern of MediaDrive Workflow Collator

Fig. 7

(part of Appendix B)

APPENDIX C

Service Definitions and Profile Schemas

Definitions

element nsdlc:Base

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

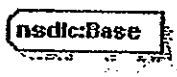
type nsdlc:Base

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:Base

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

used by

element nsdlc:Base

complexTypes CryptoKeyInfo CryptoKeyInfoPair DRMInfo License MembershipToken nsdlc:Evidence nsdlc:InterfaceBinding nsdlc:Node nsdlc:NodeIdentityInfo nsdlc:Policy nsdlc:ServiceAtt nsdlc:ServiceAttributeValue nsdlc:ServiceInfo nsdlc:ServiceMessage nsdlc:ServicePayl nsdlc:Status nsdlc:TargetCriteria OctopusNode Personality UDDIKeyedReference

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:Evidence

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

used by

elements ServiceInterfaceAuthorizationRequest/evidence nsdlc:InterfaceBinding/givenEvidence nsdlc:InterfaceBinding/requiredEvidence

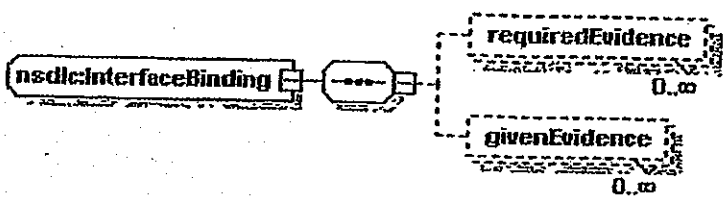
complexTypes SAMLAssertionEvidence WSPolicyAssertionEvidence

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:InterfaceBinding

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

children requiredEvidence givenEvidence

used by element nsdlc:ServiceInfo/interface
complexType WebServiceInterfaceBinding

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element nsdlc:InterfaceBinding/requiredEvidence

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type nsdlc:Evidence

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element nsdlc:InterfaceBinding/givenEvidence

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

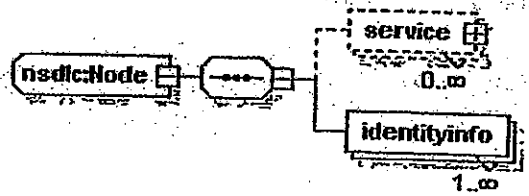
type nsdlc:Evidence

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:Node

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

children service identityinfo

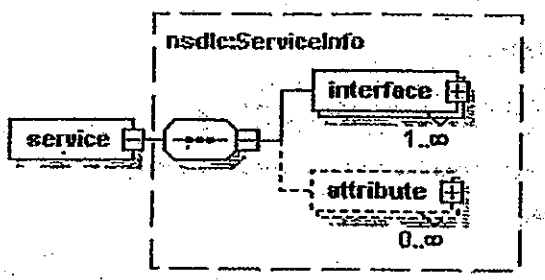
used by complexType SimpleNamedNode

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
domain	xsd:anyURI	required			

element nsdlc:Node/service

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type nsdlc:ServiceInfo

children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

element nsdlc:Node/identityinfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type nsdlc:NodeIdentityInfo

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:NodeIdentityInfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

used by
elements nsdlc:Node/identityinfo NodeIdentityInfoTargetCriteria/identityinfo
complexTypes ReferenceNodeIdentityInfo SimpleIdIdentityInfo SimpleSerialNumberNodeIdentityInfo X509CertificateIdentityInfo

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:Policy

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

used by
element ServiceInterfaceAuthorizationResponse/access_policy
complexTypes SAMLAssertionPolicy WSPolicyAssertionPolicy

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType nsdlc:ServiceAttribute

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

children value

used by elements nsdlc:ServiceInfo/attribute AttributeBasedServiceDiscoveryRequest/attribute

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	type	xsd:anyURI				

element nsdlc:ServiceAttribute/value

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type nsdlc:ServiceAttributeValue

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType nsdlc:ServiceAttributeValue

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

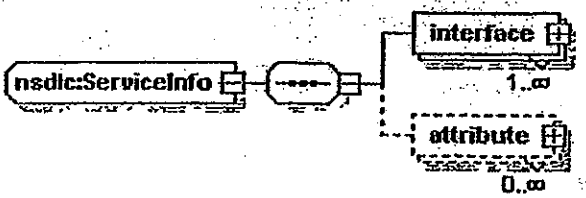
type extension of nsdlc:Base

used by element nsdlc:ServiceAttribute/value
complexType StringServiceAttributeValue

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType nsdlc:ServiceInfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

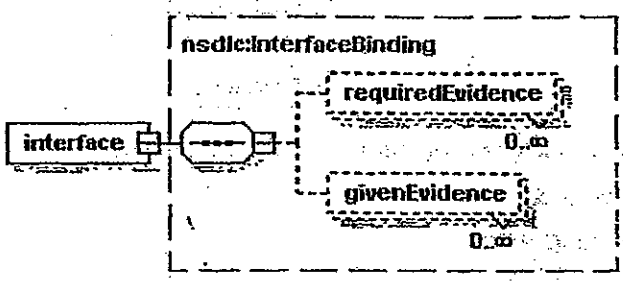
children interface attribute

used by elements nsdlc:Node/service nsdlc:ServiceMessage/service
UDDIBasedServiceDiscoveryResponse/service
AttributeBasedServiceDiscoveryResponse/service

	complexTypes	ServiceInterfaceAuthorizationRequest/serviceinfo ServiceInterfaceAuthorizationResponse/serviceinfo Authorization LicenseAcquisition LicenseTranslation MembershipTokenAcquisition NotPeerDiscovery Personalization ServiceDiscovery				
attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	service_profile	xsd:anyURI	required			
	service_category	xsd:anyURI	required			
	service_type	xsd:anyURI	required			

element nsdlc:ServiceInfo/interface

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

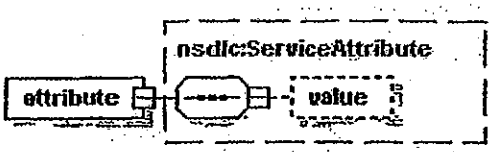
type nsdlc:InterfaceBinding

children requiredEvidence givenEvidence

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element nsdlc:ServiceInfo/attribute

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

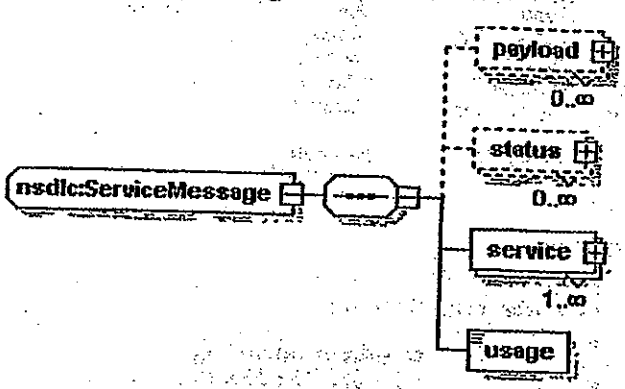
type nsdlc:ServiceAttribute

children value

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	type	xsd:anyURI				

complexType nsdlc:ServiceMessage

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

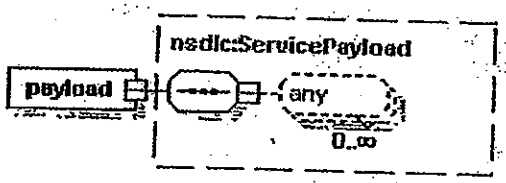
children payload status service usage

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element nsdlc:ServiceMessage/payload

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

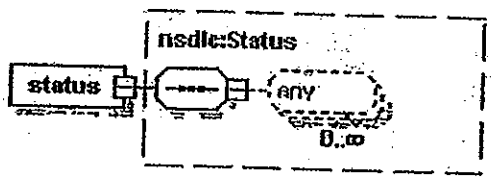
type nsdlc:ServicePayload

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element nsdlc:ServiceMessage/status

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

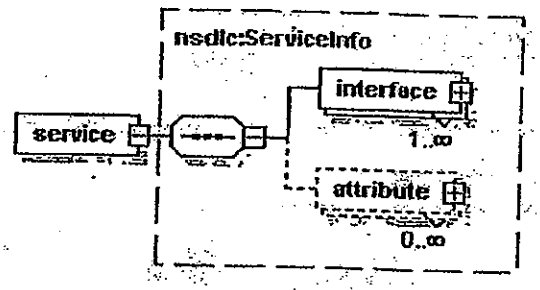
type nsdlc:Status

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
majorCode	xsd:anyURI				
minorCode	xsd:anyURI				

element nsdlc:ServiceMessage/service

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type nsdlc:ServiceInfo

children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

element nsdlc:ServiceMessage/usage

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

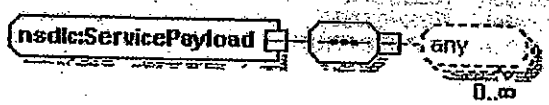
type nsdlc:MessageUsage

facets

enumeration	REQUEST
enumeration	RESPONSE
enumeration	NOTIFICATION

complexType nsdlc:ServicePayload

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

used by

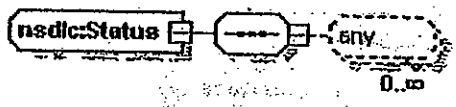
element	nsdlc:ServiceMessage/payload
complexTypes	AttributeBasedServiceDiscoveryRequest AttributeBasedServiceDiscoveryResponse OctopusLicenseAcquisitionRequest OctopusLicenseAcquisitionResponse OctopusLicenseTranslationRequest OctopusLicenseTranslationResponse OctopusPersonalizationRequest OctopusPersonalizationResponse ServiceInterfaceAuthorizationRequest ServiceInterfaceAuthorizationResponse UDDIBasedServiceDiscoveryRequest UDDIBasedServiceDiscoveryResponse UPnPPeerDiscoveryRequest UPnPPeerDiscoveryResponse

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType nsdlc:Status

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

used by element nsdlc:ServiceMessage/status

attributes

Name	Type	Use	Default	Fixed	Annota
id	xsd:anyURI	optional			
description	xsd:string	optional			
majorCode	xsd:anyURI				
minorCode	xsd:anyURI				

complexType nsdlc:TargetCriteria

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type extension of nsdlc:Base

used by complexTypes NodeIdentityInfoTargetCriteria SimplePropertyTypeTargetCriteria SimpleServiceTypeTargetCriteria

attributes

Name	Type	Use	Default	Fixed	Annotat
id	xsd:anyURI	optional			
description	xsd:string	optional			

simpleType nsdlc:MessageUsage

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

type restriction of xsd:string

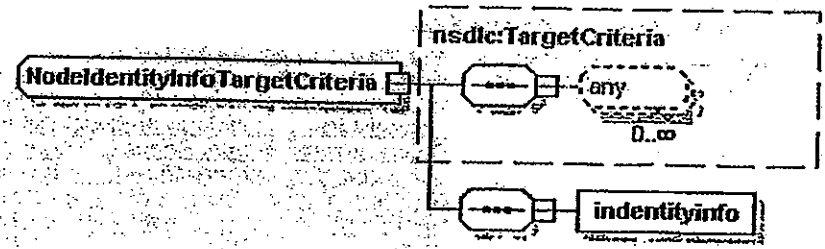
used by element nsdlc:ServiceMessage/usage

facets

enumeration	REQUEST
enumeration	RESPONSE
enumeration	NOTIFICATION

complexType NodeIdentityInfoTargetCriteria

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:TargetCriteria

children indentityinfo

attributes

Name	Type	Use	Default	Fixed	Annotatic
id	xsd:anyURI	optional			
description	xsd:string	optional			

element NodeIdentityInfoTargetCriteria/indentityinfo

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type nsdlc:NodeIdentityInfo

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType ReferenceNodeIdentityInfo

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:NodeIdentityInfo

children reference

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

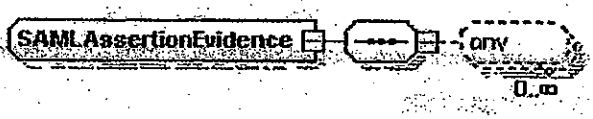
element ReferenceNodeIdentityInfo/reference

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

complexType SAMLAssertionEvidence

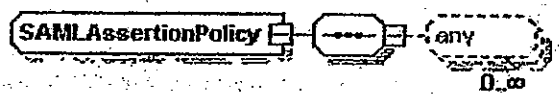
diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:Evidence

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType SAMLAssertionPolicy

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:Policy

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType **SimpleIdIdentityInfo**

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:NodeIdentityInfo

children simpleid

attributes	Name	Type	Use	Default	Fixed	Annotati
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element **SimpleIdIdentityInfo/simpleid**

diagram

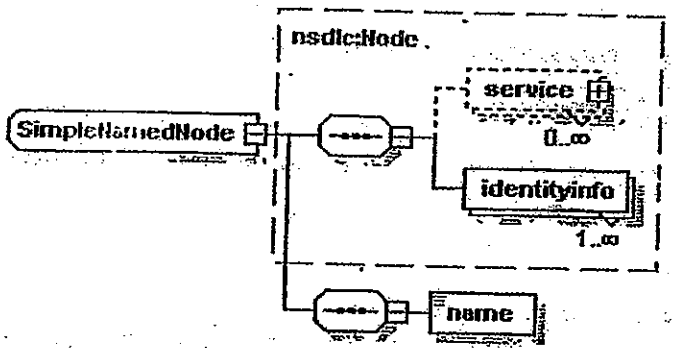


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

ype xsd:anyURI

complexType **SimpleNamedNode**

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:Node

children service identityinfo name

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	domain	xsd:anyURI	required			

element **SimpleNamedNode/name**

diagram

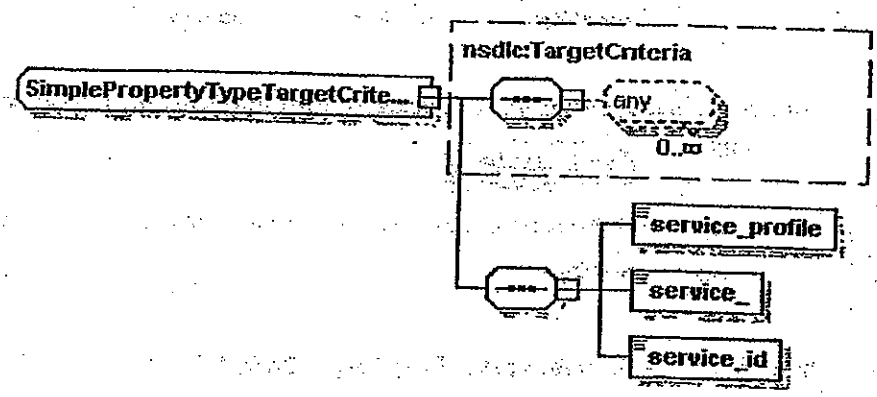


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

complexType SimplePropertyTypeTargetCriteria

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:TargetCriteria

children service_profile service_ service_id

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element SimplePropertyTypeTargetCriteria/service_profile

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

element SimplePropertyTypeTargetCriteria/service_

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

element SimplePropertyTypeTargetCriteria/service_id

diagram

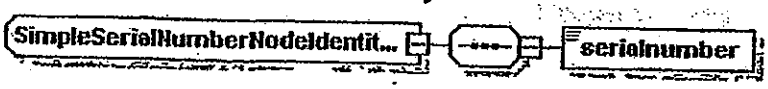


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

complexType SimpleSerialNumberNodeIdentityInfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:NodeIdentityInfo

children serialnumber

attributes	Name	Type	Use	Default	Fixed	Annota
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

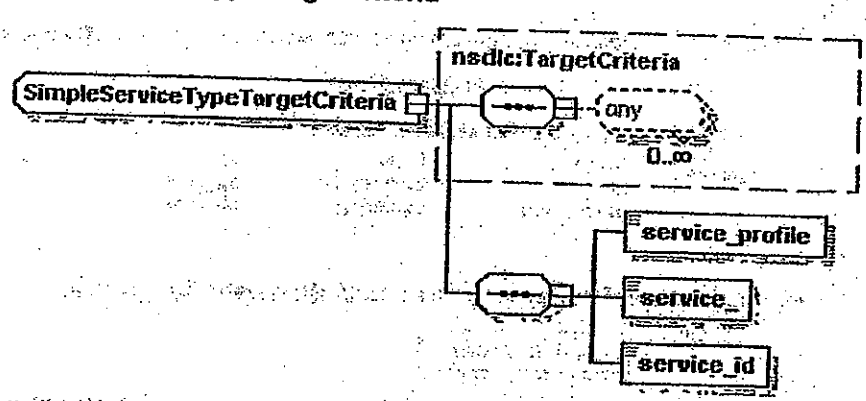
element SimpleSerialNumberNodeIdentityInfo/serialnumber

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:string

complexType SimpleServiceTypeTargetCriteria

diagram 

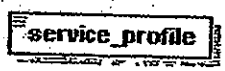
namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:TargetCriteria

children service_profile service_ service_id

attributes	Name	Type	Use	Default	Fixed	Annotatic
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element SimpleServiceTypeTargetCriteria/service_profile

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

element SimpleServiceTypeTargetCriteria/service_

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

element SimpleServiceTypeTargetCriteria/service_id

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

complexType StringServiceAttributeValue

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:ServiceAttributeValue

children stringvalue

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

element StringServiceAttributeValue/stringvalue

diagram

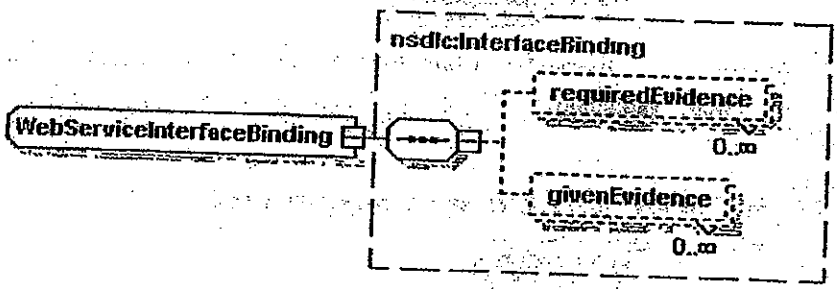


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:string

complexType WebServiceInterfaceBinding

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:InterfaceBinding

children requiredEvidence givenEvidence

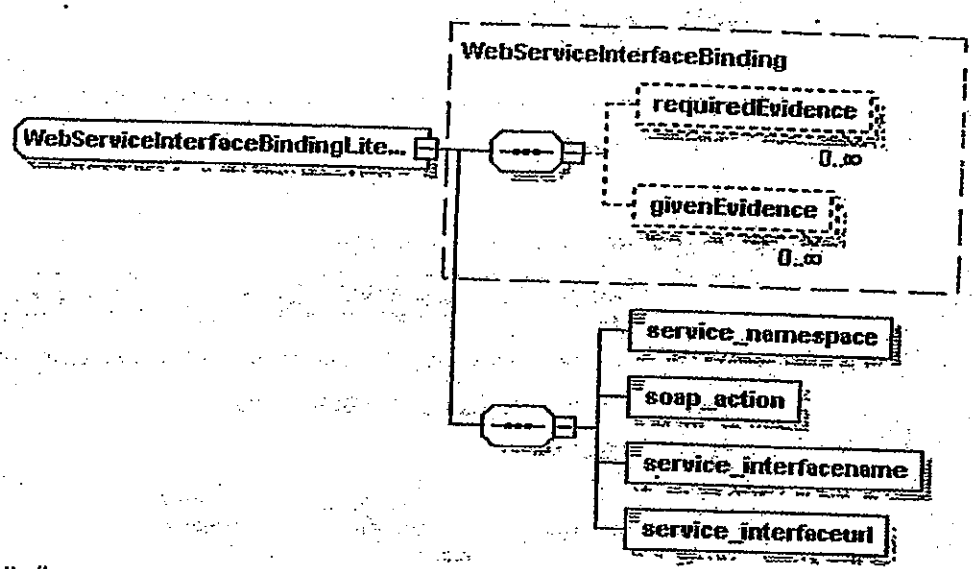
used by complexTypes WebServiceInterfaceBindingLiteral WebServiceInterfaceBindingWSDL

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType WebServiceInterfaceBindingLiteral

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of WebServiceInterfaceBinding

children requiredEvidence givenEvidence service_namespace soap_action service_interfacename service_inter

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

element WebServiceInterfaceBindingLiteral/service_namespace

diagram

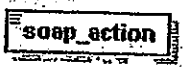


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:string

element WebServiceInterfaceBindingLiteral/soap_action

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:string

element WebServiceInterfaceBindingLiteral/service_interfacename

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:string

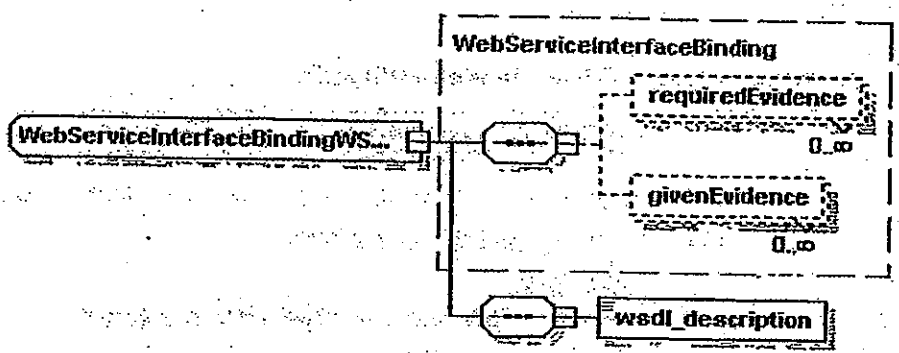
element WebServiceInterfaceBindingLiteral/service_interfaceurl

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:string

complexType WebServiceInterfaceBindingWSDL

diagram 

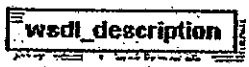
namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of WebServiceInterfaceBinding

children requiredEvidence givenEvidence wsdl_description

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

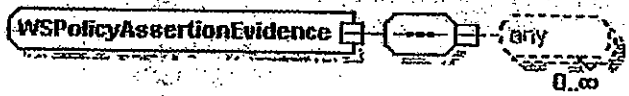
element WebServiceInterfaceBindingWSDL/wsdl_description

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:anyURI

complexType WSPolicyAssertionEvidence

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:Evidence

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType WSPolicyAssertionPolicy

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:Policy

attributes

Name	Type	Use	Default	Fixed	Annotatic
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType X509CertificateIdentityInfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type extension of nsdlc:NodeIdentityInfo

children x509certdata

used by element OctopusPersonalizationRequest/identity

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

element X509CertificateIdentityInfo/x509certdata

diagram

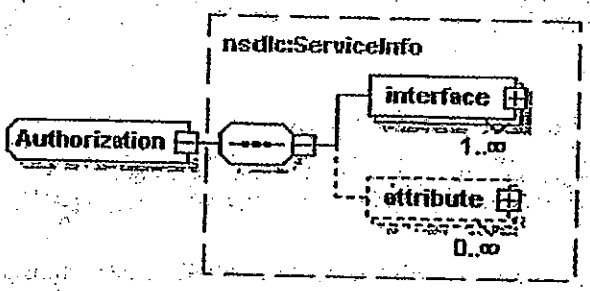


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

type xsd:base64Binary

complexType Authorization

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc

type extension of nsdlc:ServiceInfo

children interface attribute

used by complexType ServiceInterfaceAuthorization

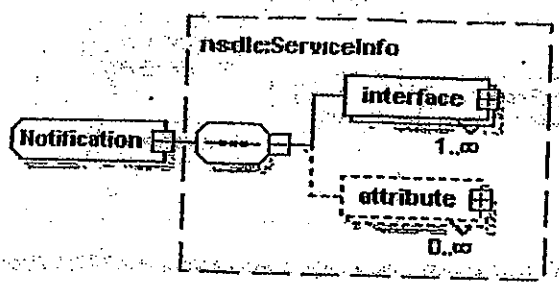
attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor	xsd:anyURI	required			

y		
service_type	xsd:anyURI	required

complexType Notification

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc

type extension of nsdlc:ServiceInfo

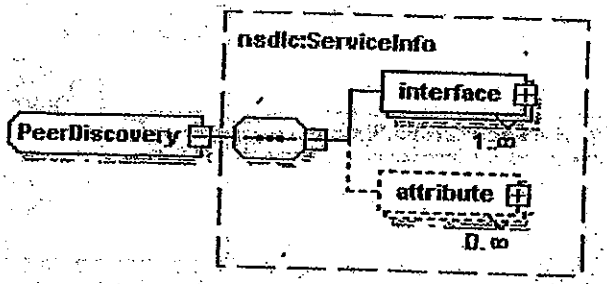
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType PeerDiscovery

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc

type extension of nsdlc:ServiceInfo

children interface attribute

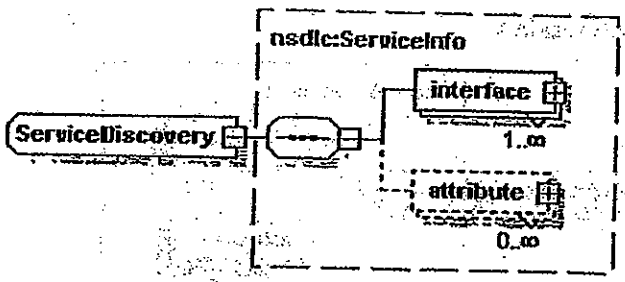
used by complexType UPnPPeerDiscovery

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType ServiceDiscovery

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc

type extension of nsdlc:ServiceInfo

children interface attribute

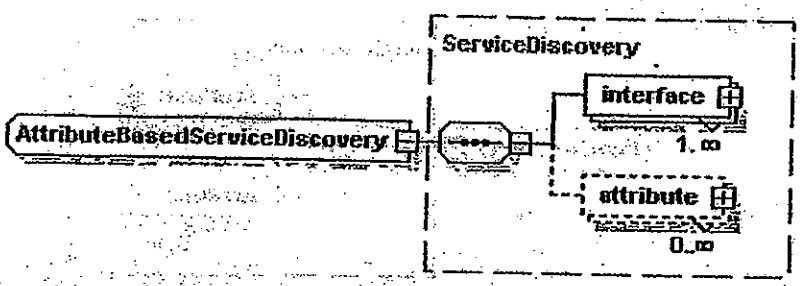
used by complexTypes AttributeBasedServiceDiscovery UDDIBasedServiceDiscovery

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_category	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType AttributeBasedServiceDiscovery

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of ServiceDiscovery

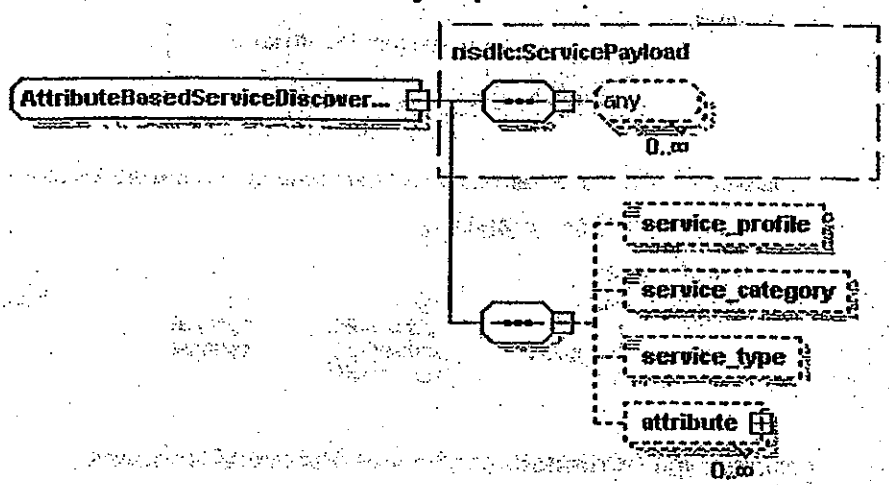
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_category	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType AttributeBasedServiceDiscoveryRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of nsdlc:ServicePayload

children service_profile service_category service_type attribute

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element AttributeBasedServiceDiscoveryRequest/service_profile

diagram service_profile

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:anyURI

element AttributeBasedServiceDiscoveryRequest/service_category

diagram service_category

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:anyURI

element AttributeBasedServiceDiscoveryRequest/service_type

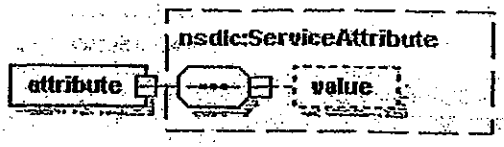
diagram service_type

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:anyURI

element AttributeBasedServiceDiscoveryRequest/attribute

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

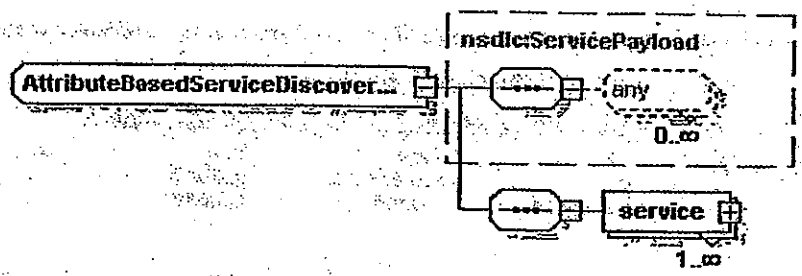
type nsdlc:ServiceAttribute

children value

attributes	Name	Type	Use	Default	Fixed	Annotati
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	type	xsd:anyURI				

complexType AttributeBasedServiceDiscoveryResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

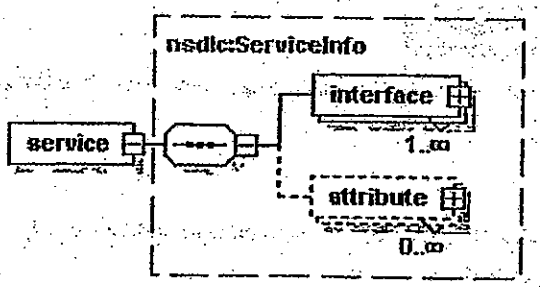
type extension of nsdlc:ServicePayload

children service

attributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element AttributeBasedServiceDiscoveryResponse/service

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type nsdlc:ServiceInfo

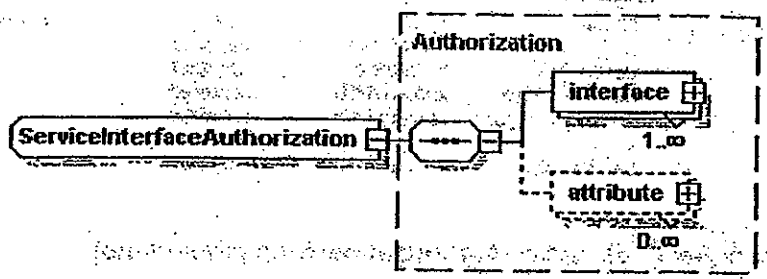
children interface attribute

attributes	Name	Type	Use	Default	Fixed	Annotatior
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	service_profile	xsd:anyURI	required			
	service_categor y	xsd:anyURI	required			

service_type	xsd:anyURI	required

complexType ServiceInterfaceAuthorization

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of Authorization

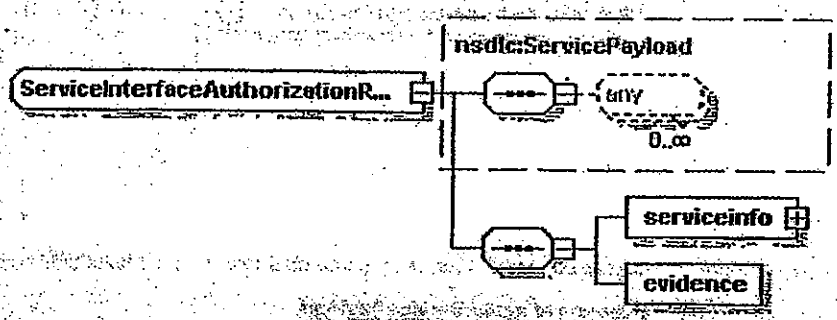
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType ServiceInterfaceAuthorizationRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of nsdlc:ServicePayload

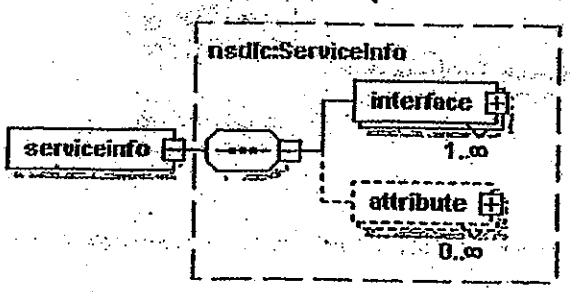
children serviceInfo evidence

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element ServiceInterfaceAuthorizationRequest/serviceinfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type nsdlc:ServiceInfo

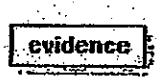
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotat
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

element ServiceInterfaceAuthorizationRequest/evidence

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

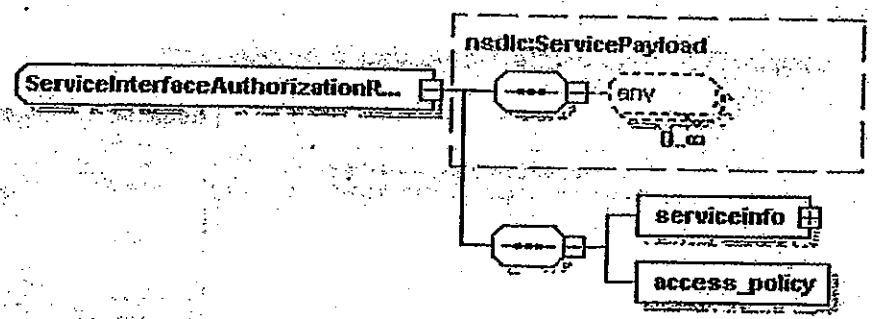
type nsdlc:Evidence

attributes

Name	Type	Use	Default	Fixed	Annotati
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType ServiceInterfaceAuthorizationResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of nsdlc:ServicePayload

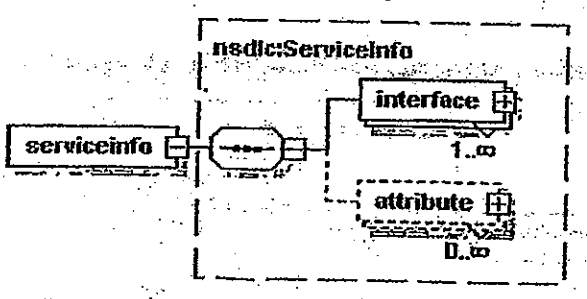
children serviceinfo access_policy

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

element ServiceInterfaceAuthorizationResponse/serviceinfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

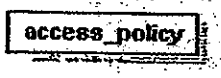
type nsdlc:ServiceInfo

children interface attribute

attributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	service_profile	xsd:anyURI	required			
	service_categor y	xsd:anyURI	required			
	service_type	xsd:anyURI	required			

element ServiceInterfaceAuthorizationResponse/access_policy

diagram



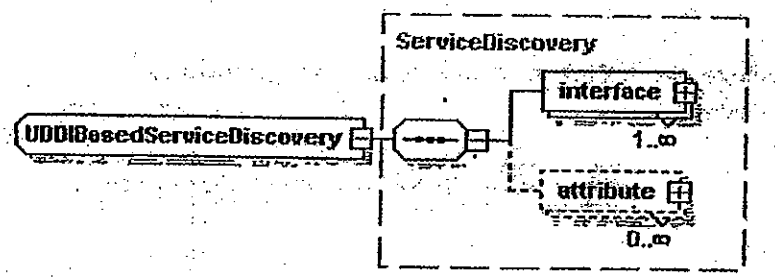
namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type nsdlc:Policy

attributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType UDDIBasedServiceDiscovery

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

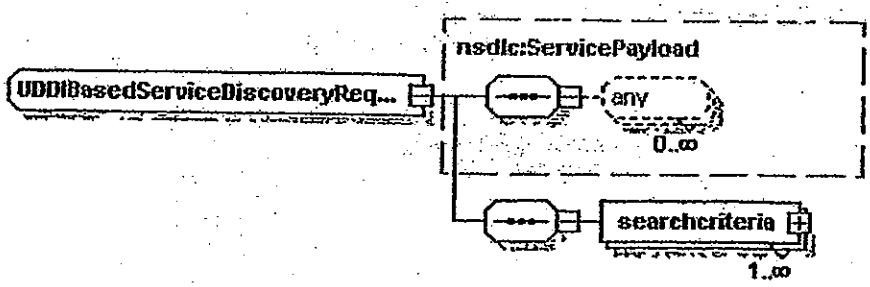
type extension of ServiceDiscovery

children Interface attribute

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	service_profile	xsd:anyURI	required			
	service_categor y	xsd:anyURI	required			
	service_type	xsd:anyURI	required			

complexType UDDIBasedServiceDiscoveryRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

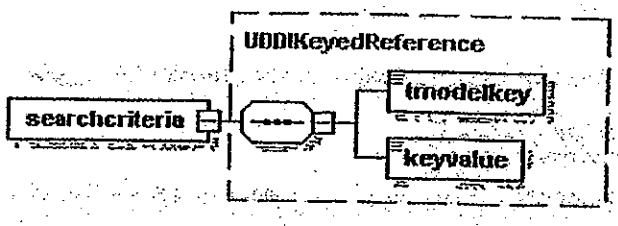
type extension of nsdlc:ServicePayload

children searchcriteria

attributes	Name	Type	Use	Default	Fixed	Annota
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element UDDIBasedServiceDiscoveryRequest/searchcriteria

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

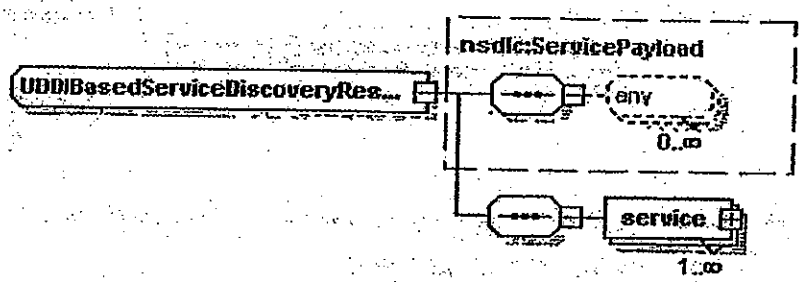
type UDDIKeyedReference

children tmodelkey keyvalue

attributes	Name	Type	Use	Default	Fixed	Annotati
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType UDDIBasedServiceDiscoveryResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

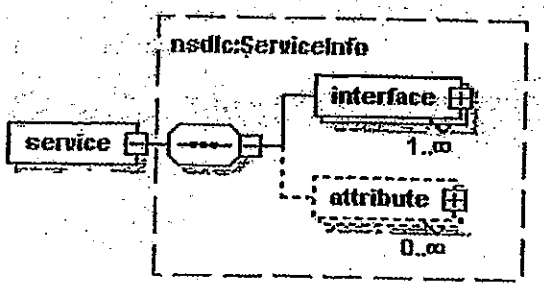
type extension of nsdlc:ServicePayload

children service

attributes	Name	Type	Use	Default	Fixed	Annotatic
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element UDDIBasedServiceDiscoveryResponse/service

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

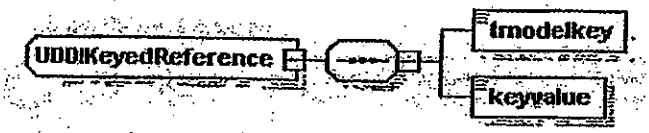
type nsdlc:ServiceInfo

children interface attribute

attributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	service_profile	xsd:anyURI	required			
	service_categor y	xsd:anyURI	required			
	service_type	xsd:anyURI	required			

complexType UDDIKeyedReference

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of nsdlc:Base

children tmodelkey keyvalue

used by element UDDIBasedServiceDiscoveryRequest/searchcriteria

attributes	Name	Type	Use	Default	Fixed	Annotatior
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element UDDIKeyedReference/tmodelkey

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:string

element UDDIKeyedReference/keyvalue

diagram

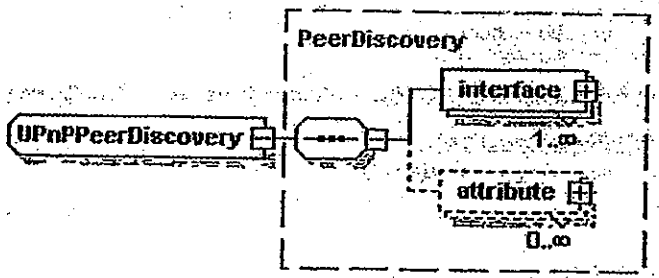


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:string

complexType UPnPPeerDiscovery

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

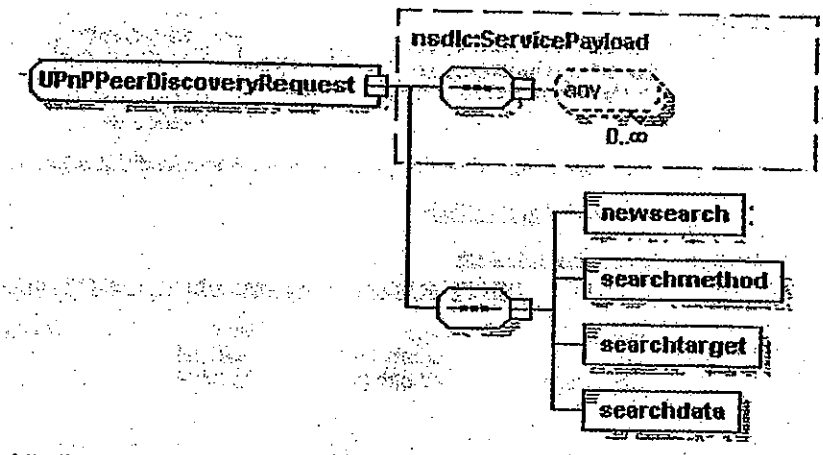
type extension of PeerDiscovery

children interface attribute

attributes	Name	Type	Use	Default	Fixed	Annota
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	service_profile	xsd:anyURI	required			
	service_categor y	xsd:anyURI	required			
	service_type	xsd:anyURI	required			

complexType UPnPPeerDiscoveryRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

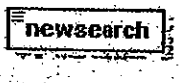
type extension of nsdlc:ServicePayload

children newsearch searchmethod searchtarget searchdata

attributes	Name	Type	Use	Default	Fixed	Annotati
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element UPnPPeerDiscoveryRequest/newsearch

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:boolean

element UPnPPeerDiscoveryRequest/searchmethod

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type UPnPPeerSearchMethod

facets enumeration M-SEARCH

element UPnPPeerDiscoveryRequest/searchtarget

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type UPnPPeerSearchTarget

facets enumeration ALL
enumeration ROOTDEVICE
enumeration DEVICETYPE

element UPnPPeerDiscoveryRequest/searchdata

diagram

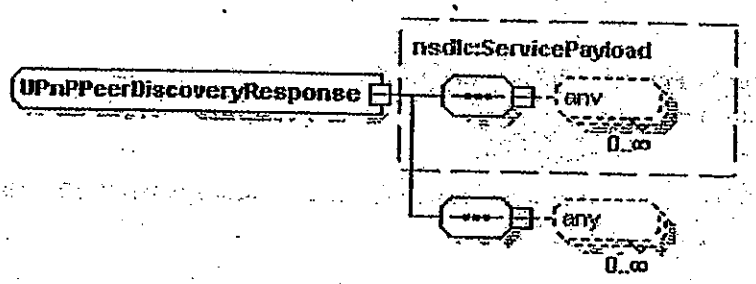


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type xsd:string

complexType UPnPPeerDiscoveryResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type extension of nsdlc:ServicePayload

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

simpleType UPnPPeerSearchMethod

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

type restriction of xsd:string

used by element UPnPPeerDiscoveryRequest/searchmethod

facets enumeration M-SEARCH

simpleType UPnPPeerSearchTarget

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

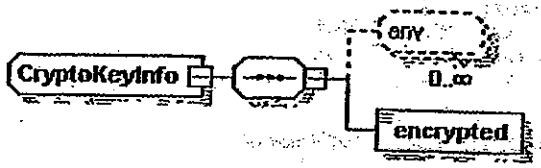
type restriction of xsd:string

used by element UPnPPeerDiscoveryRequest/searchtarget

facets enumeration ALL
enumeration ROOTDEVICE
enumeration DEVICETYPE

complexType CryptoKeyInfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type extension of nsdlc:Base

children encrypted

used by elements CryptoKeyInfoPair/privatekey CryptoKeyInfoPair/publickey

attributes

Name	Type	Use	Default	Fixed	Annotat
id	xsd:anyURI	optional			
description	xsd:string	optional			

element CryptoKeyInfo/encrypted

diagram

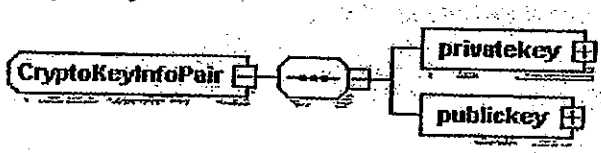


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type xsd:boolean

complexType CryptoKeyInfoPair

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type extension of nsdlc:Base

children privatekey publickey

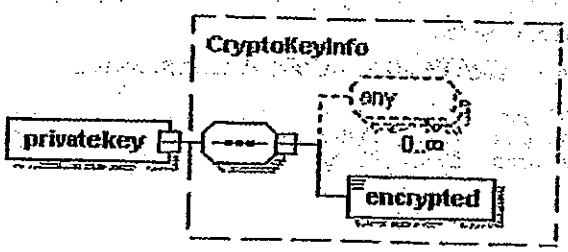
used by elements DeviceAOctopusNode/contentprotectionkey DeviceAOctopusNode/secretprotectionkey

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

element CryptoKeyInfoPair/privatekey

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

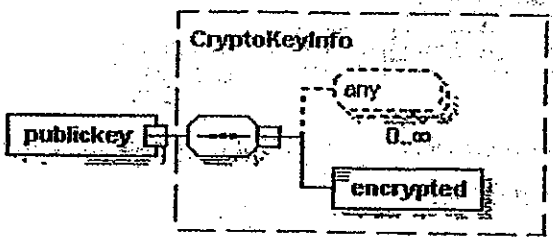
type CryptoKeyInfo

children encrypted

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element CryptoKeyInfoPair/publickey

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type CryptoKeyInfo

children encrypted

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType DRMInfo

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type extension of nsdlc:Base

used by elements Personality/drminfo License/drminfo MembershipToken/drminfo
complexType OctopusDRMInfo

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType License

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type extension of nsdlc:Base

children drminfo

used by complexType OctopusLicense

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element License/drminfo

diagram

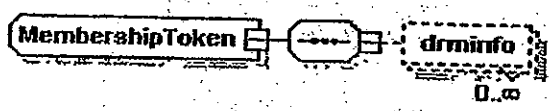


namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type DRMInfo

attributes	Name	Type	Use	Default	Fixed	Annotat
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType MembershipToken

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type extension of nsdlc:Base

children drminfo

used by complexType OctopusMembershipToken

attributes	Name	Type	Use	Default	Fixed	Annotatic
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element MembershipToken/drminfo

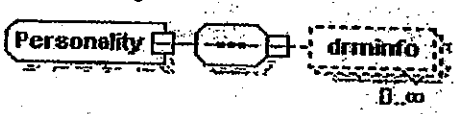
diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type DRMInfo

atributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType Personality

diagram 

namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type extension of nsdlc:Base

children drminfo

used by complexType OctopusPersonality

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

element Personality/drminfo

diagram 

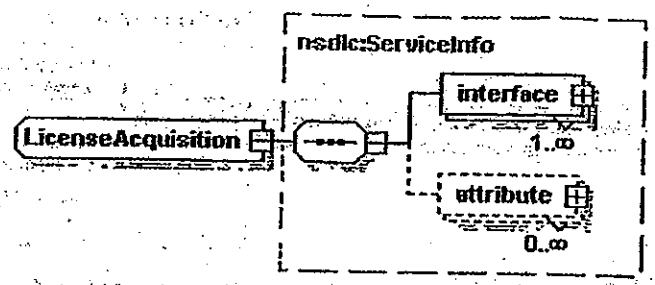
namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

type DRMInfo

attributes	Name	Type	Use	Defau t	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType LicenseAcquisition

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc

type extension of nsdlc:ServiceInfo

children interface attribute

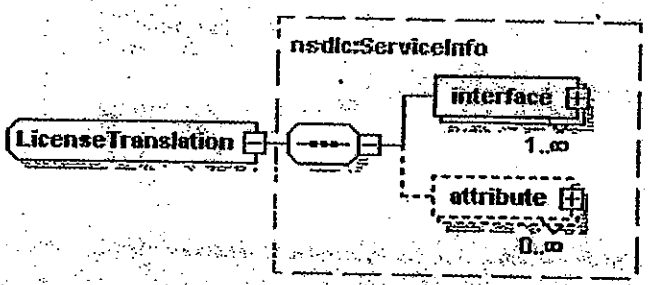
used by complexType OctopusLicenseAcquisition

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType LicenseTranslation

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc

type extension of nsdlc:ServiceInfo

children interface attribute

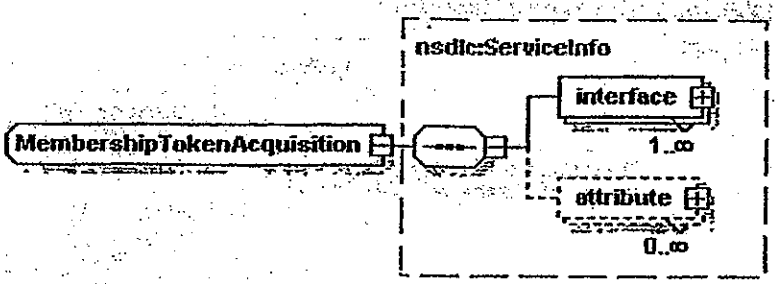
used by complexType OctopusLicenseTranslation

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType MembershipTokenAcquisition

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dnm/svc

type extension of nsdlc:ServiceInfo

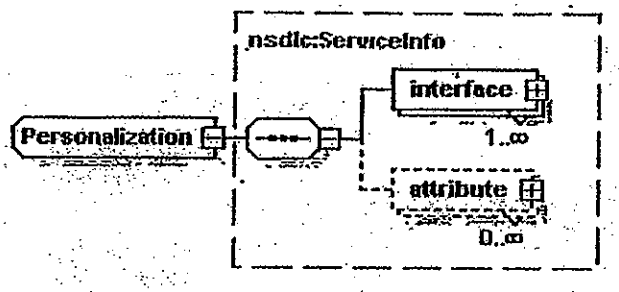
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotati
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType Personalization

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dnm/svc

type extension of nsdlc:ServiceInfo

children interface attribute

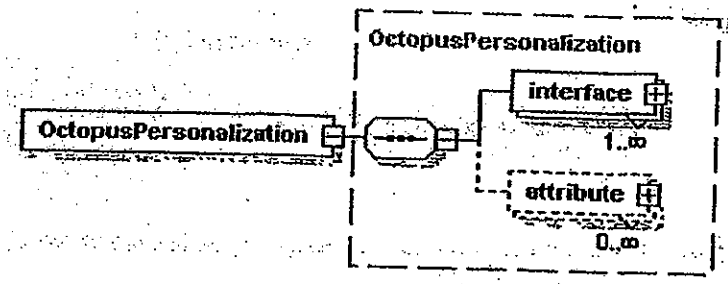
used by complexType OctopusPersonalization

attributes

Name	Type	Use	Default	Fixed	Annotati
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

element **OctopusPersonalization**

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type OctopusPersonalization

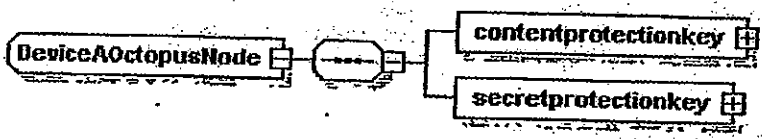
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotatic
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType **DeviceAOctopusNode**

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type extension of OctopusNode

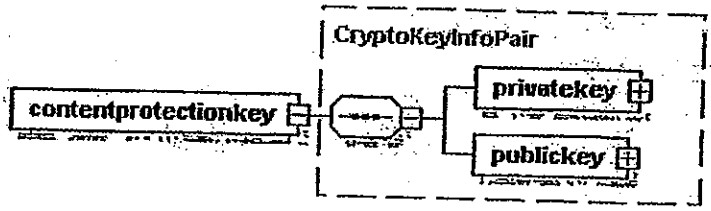
children contentprotectionkey secretprotectionkey

attributes

Name	Type	Use	Default	Fixed	Annotatior
id	xsd:anyURI	optional			
description	xsd:string	optional			
nodetype	xsd:anyURI				

element **DeviceAOctopusNode/contentprotectionkey**

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type CryptoKeyInfoPair

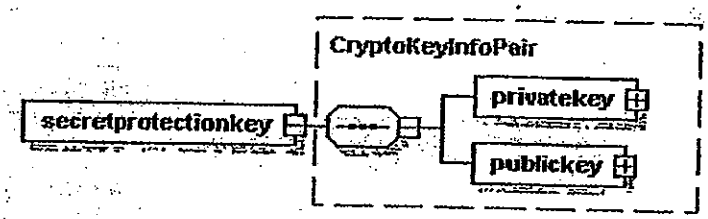
children privatekey publickey

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element DeviceAOctopusNode/secretprotectionkey

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

type CryptoKeyInfoPair

children privatekey publickey

attributes	Name	Type	Use	Default	Fixed	Annotati
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType OctopusDRMInfo

diagram



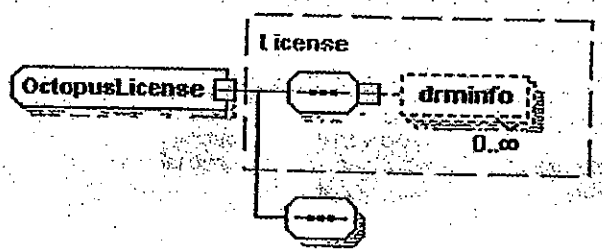
namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

type extension of DRMInfo

attributes	Name	Type	Use	Default	Fixed	Annotati
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType OctopusLicense

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

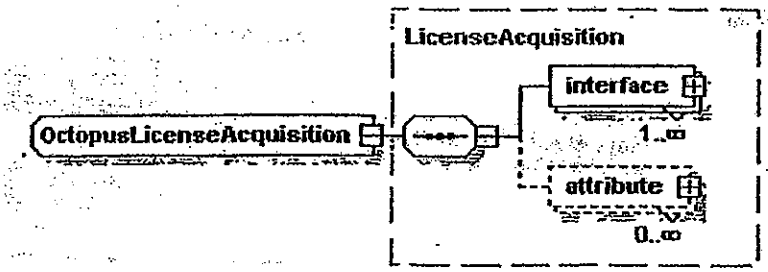
type extension of License

children drminfo

attributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType OctopusLicenseAcquisition

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

type extension of LicenseAcquisition

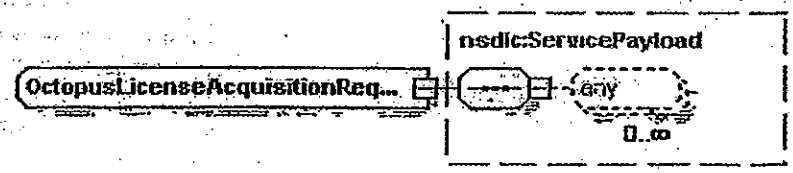
children Interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType OctopusLicenseAcquisitionRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

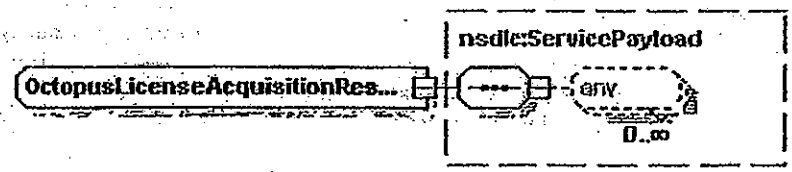
type extension of nsdlc:ServicePayload

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType OctopusLicenseAcquisitionResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

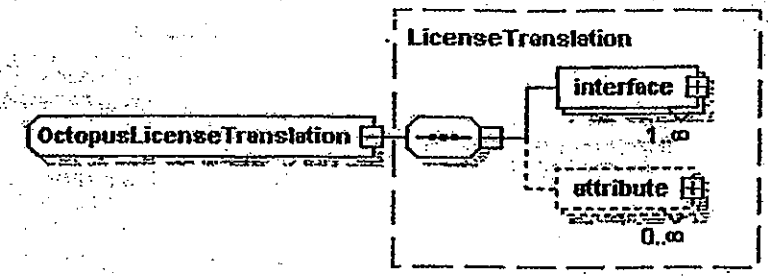
type extension of nsdlc:ServicePayload

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType OctopusLicenseTranslation

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type extension of LicenseTranslation

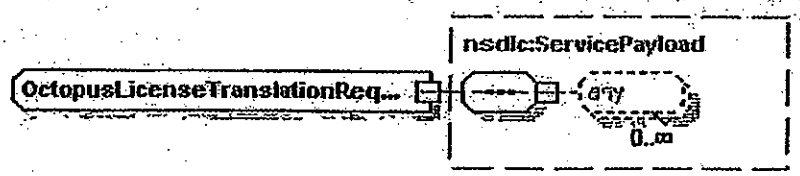
children interface attribute

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType OctopusLicenseTranslationRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

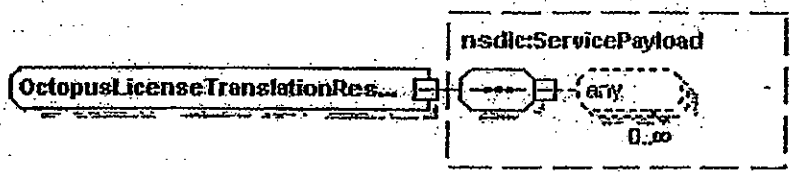
type extension of nsdlc:ServicePayload

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType OctopusLicenseTranslationResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

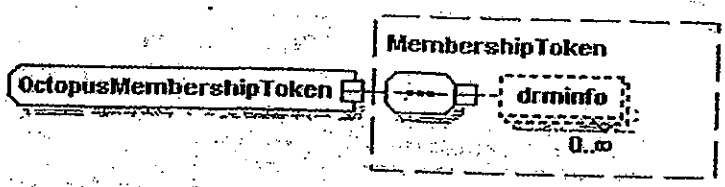
type extension of nsdlc:ServicePayload

attributes

Name	Type	Use	Default	Fixed	Annotatio
id	xsd:anyURI	optional			
description	xsd:string	optional			

complexType OctopusMembershipToken

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

type extension of **MembershipToken**

children **drminfo**

attributes	Name	Type	Use	Default	Fixed	Annotatio
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType OctopusNode

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

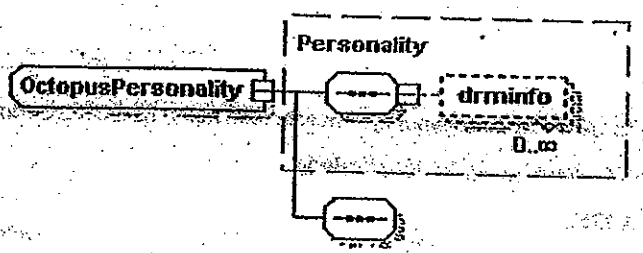
type extension of **nsdlc:Base**

used by element **OctopusPersonalizationResponse/personalItynode**
complexType **DeviceAOctopusNode**

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			
	nodetype	xsd:anyURI				

complexType OctopusPersonality

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

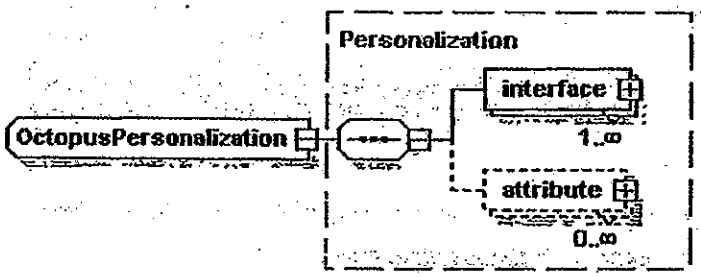
type extension of **Personality**

children **drminfo**

attributes	Name	Type	Use	Default	Fixed	Annotation
	id	xsd:anyURI	optional			
	description	xsd:string	optional			

complexType OctopusPersonalization

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type extension of Personalization

children interface attribute

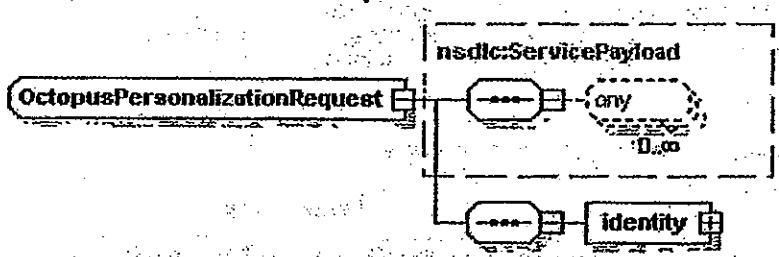
used by element OctopusPersonalization

attributes

Name	Type	Use	Default	Fixed	Annotatic
id	xsd:anyURI	optional			
description	xsd:string	optional			
service_profile	xsd:anyURI	required			
service_categor y	xsd:anyURI	required			
service_type	xsd:anyURI	required			

complexType OctopusPersonalizationRequest

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type extension of nsdlc:ServicePayload

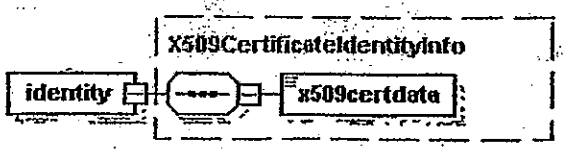
children identity

attributes

Name	Type	Use	Default	Fixed	Annotatic
id	xsd:anyURI	optional			
description	xsd:string	optional			

element OctopusPersonalizationRequest/identity

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type X509CertificateIdentityInfo

children x509certdata

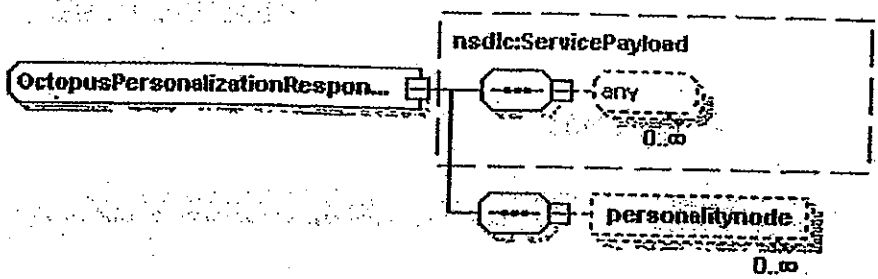
attributes

Name	Type	Use	Default	Fixed	Annotatic
id	xsd:anyURI	optional			

description	xsd:string	optional

complexType OctopusPersonalizationResponse

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type extension of nsdlc:ServicePayload

children personalitynode

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			

element OctopusPersonalizationResponse/personalitynode

diagram



namespace http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/dm/svc/ext/octopus

type OctopusNode

attributes

Name	Type	Use	Default	Fixed	Annotation
id	xsd:anyURI	optional			
description	xsd:string	optional			
nodetype	xsd:anyURI				

Profile Schemas

Core Profile

schema location: C:\ws\nemo\schema\nemo-schema-core-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core

Elements	Complex types	Simple types
Base	Base	MessageUsage
	Evidence	
	InterfaceBinding	
	Node	
	NodeIdentityInfo	
	Policy	
	ServiceAttribute	
	ServiceAttributeValue	
	ServiceInfo	
	ServiceMessage	
	ServicePayload	
	Status	
	TargetCriteria	

Core Profile Extension

schema location: C:\ws\nemo\schema\nemo-schema-core-extension-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/ext/01

Complex types
NodeIdentityInfoTargetCriteria
ReferenceNodeIdentityInfo
SAMLAssertionEvidence
SAMLAssertionPolicy
SimpleIdIdentityInfo
SimpleNamedNode
SimplePropertyTypeTargetCriteria
SimpleSerialNumberNodeIdentityInfo
SimpleServiceTypeTargetCriteria
StringServiceAttributeValue
WebServiceInterfaceBinding
WebServiceInterfaceBindingLiteral
WebServiceInterfaceBindingWSDL
WSPolicyAssertionEvidence
WSPolicyAssertionPolicy
X509CertificateIdentityInfo

Core Service Profile

schema location: C:\ws\nemo\schema\nemo-schema-core-service-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc

Complex types
Authorization
Notification
PeerDiscovery
ServiceDiscovery

Core Service Profile Extension

schema location: C:\ws\nemo\schema\nemo-schema-core-service-extension-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/core/svc/ext/01

Complex types
AttributeBasedServiceDiscovery
AttributeBasedServiceDiscoveryRequest
AttributeBasedServiceDiscoveryResponse
ServiceInterfaceAuthorization
ServiceInterfaceAuthorizationRequest
ServiceInterfaceAuthorizationResponse
UDDIBasedServiceDiscovery
UDDIBasedServiceDiscoveryRequest
UDDIBasedServiceDiscoveryResponse
UDDIKeyedReference
UPnPPeerDiscovery
UPnPPeerDiscoveryRequest
UPnPPeerDiscoveryResponse

Simple types
UPnPPeerSearchMethod
UPnPPeerSearchTarget

DRM Profile

schema location: C:\ws\nemo\schema\nemo-schema-drm-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm

Complex types
CryptoKeyInfo
CryptoKeyInfoPair
DRMInfo
License

MembershipToken
Personality

DRM Profile Extension

schema location: C:\ws\nemo\schema\nemo-schema-drm-extension-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/ext/01

DRM Service Profile

schema location: C:\ws\nemo\schema\nemo-schema-drm-service-01.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc

Complex types
LicenseAcquisition
LicenseTranslation
MembershipTokenAcquisition
Personalization

Octopus DRM Profile

schema location: C:\ws\nemo\schema\nemo-schema-drm-service-extension-octopus.xsd
targetNamespace: http://www.nemo-community.com/schemas/nemo/nsdl/2003/09/drm/svc/ext/octopus

Elements
OctopusPersonalization

Complex types
DeviceAOctopusNode
OctopusDRMInfo
OctopusLicense
OctopusLicenseAcquisition
OctopusLicenseAcquisitionRequest
OctopusLicenseAcquisitionResponse
OctopusLicenseTranslation
OctopusLicenseTranslationRequest
OctopusLicenseTranslationResponse
OctopusMembershipToken
OctopusNode
OctopusPersonality
OctopusPersonalization
OctopusPersonalizationRequest
OctopusPersonalizationResponse

[0755] Although the foregoing has been described in some detail for purposes of clarity, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. It should be noted that there are many alternative ways of implementing both the processes and apparatuses of the present inventions. Accordingly, the present embodiments are to be considered as illustrative and not restrictive, and the inventive body of work is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims.

[0756] WHAT IS CLAIMED IS:

CLAIMS

1. A system for orchestrating services provided by network peers with sufficient interoperability to enable the exchange of value through participation in distributed applications, the system comprising:

   a. a first service provider having a first service adaptation layer that exposes a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

   b. a first service consumer having a first service access point that accesses the first service interface exposed by the first service provider, discovers the first binding and invokes the first service by using that binding; and

   c. a first workflow collator that orchestrates the steps necessary to enable the first service provider to perform the first service for the first service consumer.

2. The system of claim 1 wherein the first service provides limited access to encrypted media content located on a remote Internet-based server, and the first workflow collator includes a first control program that verifies whether the first service consumer is authorized to access the content and, if so, locates the content and provides it to the first service consumer via the first service adaptation layer.

3. The system of claim 1 wherein the first service interface is specified in WSDL.

4. A system for orchestrating services provided by network peers, the system comprising:

a. a first service provider having a first service adaptation layer that exposes a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. a first service consumer having a first service access point that accesses the first service interface exposed by the first service provider, discovers the first binding and invokes the first service by using that binding;

c. a first workflow collator that orchestrates the steps necessary to enable the first service provider to perform the first service for the first service consumer; and

d. a second service provider having a second service adaptation layer that exposes a second service interface which enables network peers to access, via a second discoverable binding, a second service offered by the second service provider, the second service providing access to a service registry having a directory entry with information for locating and accessing the first service,

wherein the first service access point accesses the directory entry and uses the information to locate and invoke the first service.

5. The system of claim 4 wherein the service registry is a UDDI registry.

6. A system comprising:

a. a first service provider having a first service adaptation layer that exposes a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. a first service consumer having a first service access point that accesses the first service interface exposed by the first service provider, discovers the first binding and invokes the first service by using that binding;

c. a first workflow collator that orchestrates the steps necessary to enable the first service provider to perform the first service for the first service consumer;

d. a trust management certificate for certifying the first service for access only by authorized service consumers; and

e. a control program for providing limited access to the first service using the trust management certificate,

wherein the first access point uses the trust management certificate to validate the first service consumer for authorized access to the first service.

7. The system of claim 6 wherein the trust management certificate is an X.509 certificate.

8. The system of claim 6 wherein the trust management certificate is an SSL certificate.

9. A method for orchestrating services provided by network peers with sufficient interoperability to enable the exchange of value through participation in distributed applications, the method comprising:

a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. accessing, from a first service access point of a service consumer, the first service interface exposed by the first service provider, discovering the first binding and invoking the first service by using that binding; and

c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer.

10. The method of claim 9 wherein the first service provides limited access to encrypted media content located on a remote Internet-based server, and the first workflow collator includes a first control program that verifies whether the first service consumer is authorized to access the content and, if so, locates the content and provides it to the first service consumer via the first service adaptation layer.

11. The method of claim 9 wherein the first service interface is specified in WSDL.

12. A method for orchestrating services, the method comprising:

a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. accessing, from a first service access point of a service consumer, the first service interface exposed by the first service provider, discovering the first binding and invoking the first service by using that binding;

c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer; and

d. exposing, from a second adaptation layer of a second service provider, a second service interface which enables network peers to access, via a second discoverable binding, a second service offered by the second service provider, the second service providing access to a service registry having a directory entry with information for locating and accessing the first service,

wherein the first service access point accesses the directory entry and uses the information to locate and invoke the first service.

13. The method of claim 12 wherein the service registry is a UDDI registry.

14. A method comprising:

a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. accessing, from a first service access point of a service consumer, the first service interface exposed by the first service provider, discovering the first binding and invoking the first service by using that binding;

c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer;

d. certifying the first service, using a trust management certificate, for access only by authorized service consumers; and

e. executing a control program to provide limited access to the first service using the trust management certificate,

wherein the first access point uses the trust management certificate to validate the first service consumer for authorized access to the first service.

15. The method of claim 14 wherein the trust management certificate is an X.509 certificate.

16. The method of claim 14 wherein the trust management certificate is an SSL certificate.

17. A method for orchestrating services, the method comprising:

a. accessing, from a first service access point of a service consumer, a first service interface exposed by a first service provider, the first service interface being operable to enable network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. discovering the first discoverable binding; and

c. invoking the first service by using that binding,

wherein the steps necessary to perform the first service are orchestrated using a first workflow collator.

18. The method of claim 17 wherein the first service provides limited access to encrypted media content located on a remote Internet-based server, and the first workflow collator includes a first control program that verifies whether the first service consumer is authorized to access the content and, if so, locates the content and provides it to the first service consumer via a first service adaptation layer.

19. A method comprising:

a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. receiving a request for the first service from a first service access point of a first service consumer; and

c. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer.

20. A method comprising:

a. exposing, from a first service adaptation layer of a first service provider, a first service interface which enables network peers to access, via a first discoverable binding, a first service offered by the first service provider;

b. receiving a request for the first service from a first service access point of a first service consumer;

c. certifying the first service, using a trust management certificate, for access only by authorized service consumers;

d. executing a control program to provide limited access to the first service using the trust management certificate; and

e. orchestrating, using a first workflow collator, the steps necessary to enable the first service provider to perform the first service for the first service consumer.