ARTS

NRF's RETAIL TECHNOLOGY STANDARDS DIVISION

Best Practices for Services Implementation Using ARTS Standards
(Cloud, Enterprise, and Devices)

December 23, 2015

ACKNOWLEDGEMENTS

CONTENTS

1. INTRODUCTION

This technical report will help IT professionals in the retail industry adopt ARTS standards. Vendors, consultants, and retailers should use this paper to understand the benefits of a service-oriented-architecture (SOA), the technical considerations for implementing the ARTS standards, and reference examples set in the context of retail. Not only can these standards be used to simplify integrations, but they are also germane to today's requirements for cloud and mobile deployments.

The evolution of retail has put customers firmly in control of the buying process, so retailers must scramble to bolster the customer experience. The retail business is placing ever-more complex demands on IT to help deliver this customer experience, and it often relies on the integration of formerly separate systems. A key tenet of omnichannel retailing is the ease of moving between channels, merging the in-store and online experience. ARTS standards used in a service-oriented architecture that applies to traditional, cloud, and mobile applications can help achieve this goal.

This paper describes best practices for services implementation of ARTS standards. It should be used by IT professionals as a reference alongside the specific ARTS Technical Specification documents.

This technical report deals with the best practices for SOA implementation in a modern retail enterprise. It implies that services can potentially be deployed in virtualized environments such as public and/or private clouds. Therefore, they should be designed from the ground up not only to leverage the unique characteristics of the cloud but also to mitigate cloud-specific concerns. It is often referred to as cloud-first approach.

Potential business benefits of using SOA and Cloud Computing have been extensively covered in ARTS SOA Blueprint and Cloud Computing for Retail technical reports. Some of the major cloud adoption drivers are cost savings, scalability, and speed of deployment. Cloud technologies are essential in addressing crucial issues of modern retailing such as business agility, global customer reach, cross-channel integration, big data analytics, and providing backend services for mobile devices.

For many retailers the concept of cloud computing has been transitioning from a novel idea into a real and essential part of their IT. So, this technical report shifts the focus of the discussion from explaining what cloud computing is and how retailers can benefit from it, to advising on practical matters of implementation. It presents unbiased technical analysis dedicated to the best practices for implementing SOA strategy in the retail enterprise and is applicable to both public and private clouds. The goal is to provide retailers with the guidance on how to design and deploy complex distributed systems that run in modern environments with pervasive virtualization.

This technical report also considers recent developments in retail technology such as proliferation of mobile devices that consume RESTful APIs. Communications between devices in retail stores and modern APIs exposed by services located in the cloud are often performed using JSON data format. All this significantly impacts the development of the standards for the retail industry.

1.1 Overview

The report is primarily intended for ARTS work teams designing the standards but can be also used as a reference by technical team members that are responsible for architecting, developing, implementing and deploying services based on SOA principles within the retail enterprise. The audience that would most benefit from reading this technical report consists mainly of solution architects, software developers and IT professionals in the retail industry.

Section 1 INTRODUCTION discusses the contribution of prior ARTS technical reports on the development of this document. Readers who are interested in the historical significance of the previous ARTS papers such as the SOA Best Practice Technical Report, SOA Blueprint for Retail and Cloud Computing for Retail Technical Report should read this section.

Section 2 GENERAL SOA PRINCIPLES provides some background on the evolution of SOA concepts and introduces main SOA definitions. It discusses major SOA principals and classification of services.

Section 3 SERVICE INTERFACE DESIGN discusses the design aspects of a service interface. It describes different types of services interfaces and provides some examples.

Section 4 SERVICE IMPLEMENTATION discusses some important aspects of implementing services such as versioning discovery, etc.

Section 5 SERVICE SECURITY focuses on security aspects that are specific to the implementation of services such as transport security, authentication, and authorization.

Section 6 SERVICES INTEGRATION focuses on the services implementation scenarios in the context of integration within a retail enterprise.

1.2 Historical Background

In 2008 ARTS released two important guidance documents for the retail industry: SOA Best Practice Technical Report [1] and SOA Blueprint for Retail [2].

SOA Blueprint described the benefits that SOA could provide the retail community. At that time, service-orientation was a relatively new paradigm that had the promise of changing the way organizations fulfill business technology and application needs to achieve flexible, agile, and responsive IT architectures. The report contained retail-specific ideas that could contribute to successful implementations of SOA in the retail segment.

SOA Best Practice Technical Report identified what would constitute best practices involved with creating, maintaining, and interfacing SOA implementations. It specifically dealt with design of service interfaces, naming conventions, and details of organizing of XML schema and WSDL deliverables. Even though the primary audience for the report was ARTS workgroups, it was an extremely useful document for any SOA implementer. Many practical considerations such as granularity, versioning, extensibility, etc. were presented and discussed providing unbiased and pragmatic advice on design, development and deployment of services within a retail enterprise.

One year later, in 2009, ARTS released Cloud Computing for Retail Technical Report [3]. At that time ARTS recognized that Cloud Computing had the potential to change all the aspects of the retail value chain and create a dramatic shift inside IT departments. The report presented information about cloud computing with specific focus on retail community. It

identified areas in which a cloud-based solution would offer significant benefits to retailers as well as major obstacles to adopting cloud computing in retail.

At the time they were published those SOA and Cloud Computing technical reports provided tremendous help to the retail industry. They offered retailers valuable guidance in navigating the complex landscape of modern technology.

However, since the release of those technical reports there have been some significant technological innovations that produced a noticeable impact on how enterprise applications are implemented and deployed. The advancement of cloud computing drives retail enterprises into the world of the web scale.

This new Best Practices for Services Implementation Using ARTS Standards technical report builds upon the previously published technical reports providing more practical guidance on some important aspects of service-orientation and cloud computing like designing RESTful services, JSON serialization format, federated identity, etc. These topics had very limited coverage in the previous whitepapers but currently play an essential role in implementing integration inside a modern retail enterprise.

ARTS always influenced retail industry through thought leadership by producing technical reports that not only educated retailers and provided overview of the technology but also offered distinctive and innovative approaches. The following chapters will demonstrate how ARTS standards can be put into practice with specific focus on the implementation in the public and private clouds.

1.3 Service-Oriented Architecture

Service-orientation is an approach to designing distributed software systems as set of services. Thus, the concept of a service is the foundation of SOA. Analysis of the evolution of the main ideas and concepts behind SOA is very helpful to understanding the characteristics of services and services design.

SOA was first described by W. Roy Schulte and Yefim Natis from Gartner in 1996 [4]. It is probably fair to say that Gartner did not invent SOA but they recognized important design trends, presented them as a clear architectural concept, and gave it a name. The basic idea was to design a software system as a topology of loosely coupled components that can only be accessed through well-defined interfaces.

But it was not until early 2000s, that SOA started gaining momentum spurred by the development of powerful Web services technology. Even though there are many systems that could be reasonably called Web services at that time W3C Web Services Architecture Working Group published the following definition.

"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." [5]

This definition makes it clear that the concept of Web services is centered around technology specifications, whereas SOA is an architectural approach based on design principles. Of course, using Web services does not necessarily result in SOA implementation but Web services looked like a very good fit for execution of SOA strategy. Indeed, WSDL could

definitely be used to implement the key SOA concept of service interface and SOAP protocol provided a nice abstraction from the underling platform.

Cross-platform capabilities of Web Services and support from some major vendors made them a popular EAI technology, which helped to bring SOA concepts to mainstream users.

Unfortunately some people began to equate Web services with SOA. Many technology vendors contributed to this confusion by slapping Web services interfaces on top of their old products and declaring them to be service-oriented. These misconceptions and misuse undeservedly gave SOA a bad name.

The initial momentum behind Web services helped SOA to become one of the most overhyped technology buzzwords of the last decade. Then later, when many projects that used Web services technology began experiencing serious difficulties it created the perception that there was something wrong with the SOA approach in general.

Ironically, many Web services projects ran into difficulties because they failed to follow the major principles of SOA. Just a bunch of Web services does not constitute SOA. A lot of the industry experts at the time recognized the problem and even described JABOWS (Just a Bunch of Web Services) as a dangerous anti-pattern. Microsoft architect Nick Malik wrote in his blog back in 2008 that "JABOWS is the costly, time-consuming, valueless exercise that so many companies have taken upon themselves in the name of SOA." He also recognized that every failure of a Web services project had negative impact on the whole SOA approach. "We all lose when any one company kills their SOA initiative for lack of value. In the SOA community, we are all invested in the success of each company that has bought the hype." [6]

Critics of Web services often complained that the technology was too complex and mostly driven by large software vendors or integrators, rather than by the open community of developers. Indeed, different standards groups often led by large software vendors created numerous specifications, also known as WS-*, some of which overlapped and even competed with each other. This introduced additional complexity and hurt interoperability.

Problems with many Web services implementations and the complexity of WS-* stack pushed many developers and architects to use approach based on Representational state transfer (REST) that was described in 2000 by Roy Fielding in his doctoral dissertation [7]. REST is an architectural style for designing distributed hypermedia systems that is based on a set of architectural constraints. Even though, Roy Fielding defined REST in fairly abstract terms he had used the concepts behind this architectural style to design HTTP 1.1 and Uniform Resource Identifiers (URI), which provide the foundation for the Web. That is why the term REST is often used to describe systems that manipulate resources using standard Web technologies. Services that use HTTP protocol to manipulate URI-identified resources are referred to as RESTful services. Since SOA is technology agnostic it can be successfully implemented as a set of RESTful services.

The simplicity of REST won a lot of support in the development community. That simplicity and the adherence to fundamental principles of the Web made it much more portable across heterogeneous platforms. Many people argued that RESTful services are a much better fit for the implementation of SOA than Web services. SOA can be implemented using both REST-based and SOAP-based technologies but the choice of the approach does not change the most difficult task of designing the system as a set of well-defined interoperable services.

In January 2009 SOA experienced a significant setback when a well-known industry expert Anne Thomas Manes declared that it was dead. In her blog entitled "SOA is Dead; Long Live Services" she wrote the following damming conclusion.

"Once thought to be the savior of IT, SOA instead turned into a great failed experiment—at least for most organizations. SOA was supposed to reduce costs and increase agility on a massive scale. Except in rare situations, SOA has failed to deliver its promised benefits. After investing millions, IT systems are no better than before. In many organizations, things are worse: costs are higher, projects take longer, and systems are more fragile than ever" [8].

At the same time Anne suggested that it was not the problem with SOA approach but rather how it was executed by the technical community and product vendors. "People forgot what SOA stands for. They were too wrapped up in silly technology debates (e.g., "what's the best ESB?" or "WS-* vs. REST"), and they missed the important stuff: architecture" [8]. She went even further and re-iterated the value of SOA: "Although the word "SOA" is dead, the requirement for service-oriented architecture is stronger than ever." [8]. She suggested that the industry should stop talking about SOA and focus on services.

Obviously, Anne's post stirred up quite a bit of controversy. Some experts agreed that the term SOA was tarnished. Others disagreed stating that SOA was never about the technology but rather about architecture and therefore it never failed. The whole debate underscored that the concepts behind SOA had been suffering from lack of clarity. It also was a strong push towards creation SOA Manifesto [9] by a group of SOA experts including Anne Thomas Manes.

The SOA Manifesto was announced at the 2nd International SOA Symposium on October 23, 2009. Interestingly, at the same symposium just 9 month after she declared that SOA was dead Anne Thomas Manes gave presentation entitled "The Reincarnation of SOA". During her presentation she once again emphasized that SOA was crucial for the new cloud computing era.

Authors of the SOA Manifesto formulated 6 value statements:

1. **Business value** over technical strategy
2. **Strategic goals** over project-specific benefits
3. **Intrinsic interoperability** over custom integration
4. **Shared services** over specific-purpose implementations
5. **Flexibility** over optimization
6. **Evolutionary refinement** over pursuit of initial perfection

They also stated 14 guiding principles:

1. Respect the social and power structure of the organization.
2. Recognize that SOA ultimately demands change on many levels.
3. The scope of SOA adoption can vary. Keep efforts manageable and within meaningful boundaries.
4. Products and standards alone will neither give you SOA nor apply the service orientation paradigm for you.

5. SOA can be realized through a variety of technologies and standards.

6. Establish a uniform set of enterprise standards and policies based on industry, de facto, and community standards.

7. Pursue uniformity on the outside while allowing diversity on the inside.

8. Identify services through collaboration with business and technology stakeholders.

9. Maximize service usage by considering the current and future scope of utilization.

10. Verify that services satisfy business requirements and goals.

11. Evolve services and their organization in response to real use.

12. Separate the different aspects of a system that change at different rates.

13. Reduce implicit dependencies and publish all external dependencies to increase robustness and reduce the impact of change.

14. At every level of abstraction, organize each service around a cohesive and manageable unit of functionality.

A well-known SOA expert Thomas Erl wrote some interesting insights about the value statements and the principles above on the Annotated SOA Manifesto page [10].

Several SOA related specifications have been released in the recent years. In November 2011 the Open Group published "SOA Reference Architecture" standard [11]. Then, in December 2012, OASIS released "Reference Architecture Foundation for Service Oriented Architecture Version 1.0" [12]. In April 2014 the Open Group published "Service-Oriented Architecture Ontology Version 2.0" [13]. These releases prove the continued relevance of SOA to modern technology solutions.

It is important to note that SOA provides foundation for further innovation. Recently there has been a lot of buzz around the concept of "microservices". James Lewis and Martin Fowler wrote an article dedicated to the subject of Microservice Architecture [14]. Authors do not provide a formal definition of the microservices architectural style but they describe a set of common characteristics for microservice architectures.

1. Componentization via Services

2. Organized around Business Capabilities

3. Development teams own Products (microservices) they created rather than just participate in development Projects (Products not Projects)

4. Smart endpoints and dumb pipes

5. Decentralized Governance

6. Decentralized Data Management

7. Infrastructure Automation

8. Design for failure

9. Evolutionary Design

Microservices definitely go beyond just simple granularity considerations. The idea is to architect a complex system as a set of highly-cohesive services that can evolve independently over time. Still, a lot of experts believe that it is just an evolution of SOA. For example, Steve

Jones wrote that "Microservices is just a Service Oriented Delivery approach for a well architected SOA solution" [15]. Netflix that is often touted as an example of successful implementation of the microservices architecture uses the term "fine grained Service Oriented Architecture" [16] to describe the approach they employ.

Design of modern distributed systems and even novel architectures, like microservices or software-defined architecture [17], are based on solid SOA principles. Those principles need to be followed to produce systems that meet the requirements of a modern retail enterprise.

Recommendation 1.1 Follow SOA Principles

Recommendation	Services design and implementation should follow the major principles of service-orientation.
Rationale	Service orientation is a proven approach for building distributed software systems that helps to achieve enterprise agility. It facilitates quicker and more efficient response to changing business requirements and promotes reuse of service capabilities.

1.4 Cloud Computing

Cloud considerations are very important for successful implementation of contemporary services. Today, Cloud Computing is definitely one of the most disruptive technologies impacting modern retail enterprises. As such, it offers a lot of great opportunities and at the same time poses plenty of challenges.

Major aspects of Cloud Computing have been covered in ARTS' Cloud Computing for Retail Technical Report [3].

1.4.1 Cloud Computing Definition

For the purposes of this discussion we will use "The NIST Definition of Cloud Computing" [18] that was published as recommendation of the National Institute of Standards and Technology. This definition is composed of five essential characteristics, three service models, and four deployment options.

Essential Characteristics:

- *On-demand self-service.* A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.

- *Broad network access.* Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).

- *Resource pooling.* The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at

a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

- *Rapid elasticity.* Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.

- *Measured service.* Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

Service Models:

- *Software as a Service (SaaS).* The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

- *Platform as a Service (PaaS).* The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

- *Infrastructure as a Service (IaaS).* The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Deployment Models:

- *Private cloud.* The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

- *Community cloud.* The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

- *Public cloud.* The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

- *Hybrid cloud.* The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

1.4.2 Design for the Cloud

In the recent years Cloud Computing has been steadily replacing less flexible software licensing and deployment models. Advancements in technology open new markets and enable new products, services, and business processes. While many retail organizations have been gradually adopting cloud for several years, moving from development to production, the major shift to truly strategic adoption is just getting underway. Therefore, it is crucial that services are architected to be deployable in the cloud.

Even though cloud offers many potential benefits like greater scalability, utility-based pricing and global reach, it also brings forward cloud-specific challenges related to security, compliance, multitenancy, etc.

To make sure new services can be securely and effectively deployed inside cloud environment they should be designed from the very beginning with cloud back-ends in mind. The assumption should be that services will be delivered from a cloud infrastructure rather than running on a corporate server. That has significant implications for how services are designed since they have to deal with cloud-specific challenges. Another important point is that services designed with cloud environments in mind can be fairly easily deployed on-premises. Therefore, cloud-first approach to service design offers greater flexibility.

Recommendation 1.2 Design for the Cloud

Recommendation	Services should be designed so that they could be delivered from the cloud.
Rationale	If a service is designed with the cloud deployment model in mind it takes into account different cloud specific architectural considerations. That, in turn, facilitates efficient and secure functioning in a cloud environment. At the same time such service can also be deployed on-premises. Therefore, the approach of designing services for the cloud offers more flexibility.

It is important to note, that simply using cloud as the first deployment option will not automatically yield the desired result. The greatest benefit will be obtained from reorienting the focus, skills and architecture of the retail enterprise to change the way services are delivered, operated and consumed.

2. GENERAL SOA PRINCIPLES

Chapter 1 presented guiding SOA principles that were published as part of SOA Manifesto [9]. They deal with both technical and business aspects of SOA development. Design of enterprise SOA strategy should start with review of those principles and their implications in the context of a particular organization.

This chapter focuses on the technical SOA principles that can be used as objective criteria of what is considered a "truly" service-oriented solution.

2.1 Service

The key concept in SOA is the concept of service. ARTS SOA Best Practice Technical Report defined SOA Service as "a repeatable task within a discretely defined business process with a well-specified outcome and a standardized, published, discoverable interface" [1]. The technical report also cites four tenets of service-orientation that were originally formulated by Don Box [19] primarily in relation to Microsoft's implementation of WS-* stack called Windows Communication Foundation (WCF).

1. Boundaries are explicit
2. Services are autonomous
3. Services share schema and contract, not class
4. Service compatibility is based on policy

The Open Group has somewhat similar definition of the concept of service [20].

A service:

- Is a logical representation of a repeatable business activity that has a specified outcome (e.g., check customer credit, provide weather data, consolidate drilling reports)
- Is self-contained
- May be composed of other services
- Is a "black box" to consumers of the service

Oasis gives a more abstract definition of a service in its Reference Model for Service Oriented Architecture [21]. Service is defined as "the means by which the needs of a consumer are brought together with the capabilities of a provider," where capability is "a real-world effect that a service provider is able to provide to a service consumer". The SOA reference model also states that distributed capabilities "may be under the control of different ownership domains".

Thomas Erl, who authored several books on service orientation, simply defines service as a unit of solution logic to which service-orientation has been applied to a meaningful extent. Applying service-orientation means applying service-orientation design principles that he described in detail in his book SOA Principles of Service Design [22]. These principles are:

- Standardized service contract
- Service loose coupling
- Service abstraction
- Service reusability

- Service autonomy

- Service statelessness

- Service discoverability

- Service composability

It is important to note that all seemingly different service definitions have several common themes.

First, services communicate across boundaries. They can be spread over large geographical distances or as it was defined above "may be under the control of different ownership domains". It means that there are tangible costs associated with crossing these boundaries that have to be taken into account when designing service-oriented system. Also cross-boundary communications are much more likely to fail and services should be designed to deal with such contingencies.

Second, services should be loosely coupled. In software coupling is a measure of dependency between two modules. Making services autonomous (self-contained) and abstracting implementation behind the service interface (black box to consumers) are important considerations to reduce coupling. A distributed system that consists of loosely coupled services is more robust and easy to evolve since modifications to one of the services are less likely to significantly impact the rest of the system.

Third, services should be designed so that they could be easily reused and also could participate in service composition. Service reusability is a highly desirable characteristic. There are multiple aspects of service design that can facilitate potential reuse of service capabilities. In addition, some special considerations might be necessary to produce services that can be effective composition members.

2.2 Definitions

For the purposes of this technical report we will use the following definitions of the major SOA concepts.

Service is an autonomous self-contained unit of functionality that abstracts its implementation details behind a well-defined interface. Service can be thought of as a container of coherent service capabilities.

Service capability is an elemental piece of functionality with specified outcome (real-world effect). This outcome could be a change in data or producing certain response to a service consumer.

Service consumer is a software module that uses one or more of services capabilities. Service consumers access service capabilities via service interface.

Service interface represents the means for interacting with the service. Service interface is technology specific and provides detail information about protocols, message exchange patterns (MEP), means to invoke service capabilities and service information model.

Service information model is a detail description of the data that can be exchanged with the service. It includes structure and format of the information exchanged between the service and its consumers. It is important to note that consistent semantic interpretation of the exchanged data is very important especially if service interactions cross ownership boundaries.

2.3 Loose Coupling

In software engineering coupling is a measure of dependency between software modules. The concept of coupling is often discussed together with the concept of cohesion, which is a measure to which degree the elements of a software module belong together. High cohesion usually implies loose coupling and vice versa.

According to the last guiding principle from SOA manifesto services should represent cohesive units of functionality. Therefore, SOA system should consist of loosely coupled cohesive services. In other words, a service contains code that naturally belongs together but any two services should not have a high degree of dependency on each other.

Making services loosely coupled by reducing dependencies between them is crucial for successful implementation of an SOA solution. Some SOA practitioners believe that loose coupling is the key principle in service-orientation. For example, Ganesh Prasad wrote and interesting post "SOA as Dependency Oriented Thinking" [23]. In his opinion "SOA is the science of analysing and managing dependencies between systems, and "managing dependencies" means eliminating needless dependencies and formalising legitimate dependencies into readily-understood contracts."

In his article Ganesh considers four layers: business, application, information, and technology. To achieve loose coupling dependencies should be carefully managed at the every layer. As, at the business layer "the focus on dependencies forces us to rationalise processes and make them leaner" [23]. For example, is pre-authorization of a credit card necessary to accept a customer order? In many cases it is sufficient just to capture credit card information and then deal with infrequent payment issues at the time when the order is about to be fulfilled. In this case the customer order capture service would not need dependency on the tender authorization service.

Ganesh Prasad provided a detail described description of dependency-oriented approach to SOA in his two-volume book "Dependency-Oriented Thinking" [24], [25].

Many of the SOA principles listed above (autonomy, abstraction, discoverability, etc.) help to reduce service coupling.

Recommendation 2.1 Carefully Manage Dependencies

Recommendation	Services should be designed and implemented to be loosely coupled. It means that every effort should be made to minimize the dependencies between components inside a service-oriented solution. Service consumers should only depend on service interface and policies.
Rationale	It is difficult to evolve different parts of tightly coupled systems independently and it becomes more and more complex as the size of the system grows. Also at runtime failures in one part of a tightly coupled system can have ripple effect and essentially bring the whole system down. Also loose coupling facilitates the reuse and even may enable new behaviors that were not anticipated during the original design of the system.

Another common type of dependency is so-called temporal coupling between a service provider and a service consumer. If a system that consumes service capabilities is required to receive a response within a relatively small time interval and cannot tolerate latency then it is

said to have high temporal coupling with the service. Such distributed systems are brittle since any delay in processing of a request can potentially cause a systemic failure. Typically temporal coupling implies implementation using synchronous RPC-style communications (6.1) that do not offer a good support for offline scenario. Also, such systems are vulnerable to the effects of unexpected spikes in activity that usually result in delays in processing requests.

Temporal coupling often happens when human interaction with the user interface is involved or when the user have to interact with a UPOS device like, for example, insertion of a form into the "jaws" of POS printer. By contrast, printing a receipt can be done completely asynchronously.

The first tenet of service-orientation (explicit boundaries) underscores the importance of dealing with temporal coupling. Since service invocations cross boundaries, delays can occur and sometimes messages can get completely lost.

Recommendation 2.2 Use Asynchronous Communications

Recommendation	To avoid temporal coupling service consumers should communicate with services using asynchronous patterns.
Rationale	If service consumers communicate in a synchronous manner they cannot continue the execution before a response is received. Therefore, such service consumers implicitly assume that they will receive the response fairly fast. This assumption is not valid for cross-boundary communications. One way to avoid temporal coupling is to use a queuing system or some kind of Message-Oriented Middleware (MOM). Even if a response is required to continue the with the business process, an RPC style call can be performed in asynchronous fashion and response timeout should be properly handled as an error condition.

There are many other types of coupling: dependency of a particular technology, dependency on service implementation details, etc. In general, zero coupling is impossible to achieve and sometimes coupling can be justified as certain dependencies can be used to achieve significant performance gains. Nevertheless, SOA principles mandate that dependencies should be avoided and service consumers should only rely on service interfaces. Even when coupling is introduced to achieve some other design goals, it should be carefully managed.

2.4 Consideration for Reuse

Service interfaces should maximize flexibility of use. Because services are expected to be reused, the interfaces must accommodate the service usability at different levels of the solution. The best method to achieve this is to have a business-level abstraction at the interface that does not tie into a specific business process or implementation.

The availability of mature, non-proprietary technologies that can be used to implement service interfaces allows exposure of service capabilities to a wide variety of different clients. That dramatically increases the potential for services reuse. To realize this potential, services should be designed as business process agnostic enterprise resources.

2.5 Services Classification

Classification analysis of services is a very useful aspect of creating a consistent SOA vision. This topic was addressed in ARTS SOA Best Practices Technical Report [1]. The idea is to

define meaningful service categories based on well-understood underlying classification principles. This classification helps identify common characteristics of services that fall into a particular category. Therefore classification of services enables designing common approaches that can be applied to all services with similar characteristics.

Because services can be viewed as containers of capabilities related to a common functional context, it is possible to combine capabilities related to different classification categories under a single service umbrella. However, such design may not be optimal. Meaningful classification of services and their capabilities helps achieve higher degree cohesion. For example, combining a stateless capability with a service that has to manage state may have a negative impact on the potential scalability of that capability.

Another important consideration in SOA design is how services work differently with data and the underlying data stores. Some services use read-only data and can be easily scaled out. Other services could be constantly modifying the data in a data store and therefor have to deal with concurrency issues. Therefore, different categories of services may use different data access patterns and have different consistency requirements.

There are two main types of services: infrastructural services that address cross-cutting concerns and that are not part of core business logic and services that expose important business-oriented capabilities.

Infrastructure services, which are often referred to as *utility services*, provide generic business-agnostic capabilities. Because these services expose common functionality that is not associated with any particular business activity, they are highly reusable. Examples of functional areas typically addressed by utility services include security, discovery, and logging, to name a few.

Business services can be further divided into three major groups: entity services, task services, and process services.

Entity services are responsible for the maintenance of business entities (e.g., customer, item, address) that define the functional context of the service. For this reason, entity services are agnostic to business processes that might use them and thus have high degree of reusability. It is very common for entity services to support an entity-level create, read, update, and delete (CRUD) interface. Entity services are often used to abstract data stores and can be thought of as data-centric services.

Task services expose business-level capabilities that are used to make up an organization's business processes. They represent action-centric units of business logic. A few examples of task services include a transaction tax calculation service or a service that evaluates a customer's creditworthiness. Task services can be fairly agnostic to business processes that use them, resulting in high reusability. On the other hand, task services can be designed to address a quite specific concern of a particular business process. In this case, they might be difficult to reuse in a different business context. To improve reusability, task services should be created to encapsulate clear-purpose abstract units of business logic.

Process services represent end-to-end logic of a business process. They are often used to tie together the data-centric and action-centric units of business logic. Process services are commonly positioned as composition controllers, composing functionality offered by task services, entity services, and utility services. An example of a process service is a customer order-processing service. Depending on implementations, this service might coordinate a variety of business activities. It creates a basket of items (catalog service), loads customer

information (customer service), calculates the price (price service) and applicable taxes (tax service) of every item, determines shipping costs (shipping service), verifies customer's credit (credit-verification service), performs merchandise reservation (inventory service), secures payment (payment-processing service), schedules shipment (shipping service), and notifies the customer (customer-notification service) of the order status. The most important aspect of process services is that they typically manage the process state for the entire duration of a process. This means there is a certain level of correlation among different invocations of service capabilities. Despite the fact that process services encapsulate a particular business process, they still can be reused and can participate as a composition member in another process service that has a larger scope.

	Utility Services	Entity Services	Task Services	Process Services
Main Purpose	Provide generic infrastructural functionality	Expose and manage business entities	Implement a business task	Implement a business process
State Management	Stateless	Stateless	Stateless	Manage process state
Reusability	Highly reusable	Highly reusable	Reusable	Limited reusability
Interface	Group of infrastructural capabilities associated with a common functional context	Entity-level CRUD and other entity-related capabilities	Capabilities associated with a particular business task	Capabilities used to accomplish a particular business process
Data Access	Accesses different types of data stores depending on the nature of the service. Typically utility services read configuration information.	Data centric services that typically read and write entities data.	Typically read configuration and business rules that are necessary to complete the task, for example price derivation rules.	Typically read configuration information and business rules that are necessary to perform the business process. Often use highly available storage to manage the process state.

Example	Logging service	Customer service	Tax Calculation service	Retail Transaction service

3. SERVICE INTERFACE DESIGN

As was defined in the previous chapter, the service interface represents the means by which service consumers access service capabilities. It is one of the key concepts of service-orientation. This chapter discusses the design aspects of a service interface and how ARTS standards can be leveraged in this process.

Designing a well-crafted service interface requires taking into consideration a number of factors. Because the interface is what consumers use to interact with the service, its simplicity and convenience for consumers are a key criterion. Other factors include service architecture, selection of the most appropriate technology, using standards, etc.

The sections that follow highlight the chief considerations in designing a well-crafted service interface.

3.1 Different Types of Service Interfaces

A service interface represents the means for interacting with the service. It is technology specific and provides detailed information about protocols, MEP, and the service information model. A service interface fully describes how clients can consume service capabilities.

The ecosystem of APIs that can be used to build a distributed system is fairly complex. The figure below represents a simplified categorization of the distributed APIs.
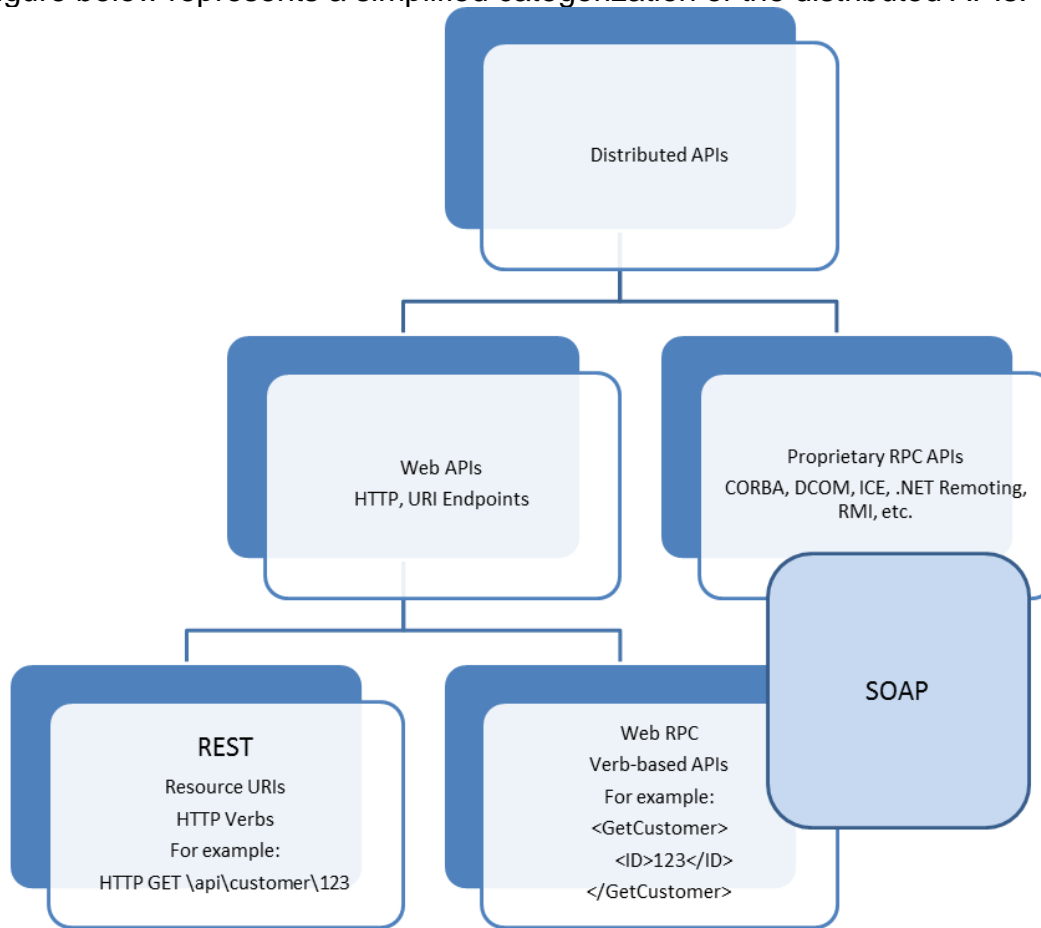
Figure 3.1 Ecosystem of Distributed APIs

Best Practices for Services Implementation Using ARTS Standards

Even though there are many different variations of service interfaces, there are two major approaches adopted by the industry: SOAP-based and RESTful.

In the past vast majority of implementations used SOAP-based services but recently RESTful APIs became the preferred approach to implementing service interfaces. The key factor is that they are simpler and provide greater interoperability.

Recommendation 3.1 Use RESTful API as Default Approach

Recommendation	Use RESTful Web APIs as the default option for implementing service interfaces. Use a SOAP-based RPC approach only if the service needs a certain capability that is impossible or very difficult to implement using RESTful approach. You should also consider SOAP if platform interoperability is a secondary concern and SOAP-based implementation offers some real benefits.
Rationale	Restful Web API approach offers the highest degree of interoperability, which is a key factor for building distributed systems. For some types of clients such as mobile devices SOAP protocol is too heavy, and could be difficult to implement.

3.1.1 SOAP Services

SOAP can be thought of as an evolution of the XML-RPC approach that used XML messages to communicate over HTTP as a transport mechanism. SOAP is an extensible protocol maintained by XML Protocol Working Group of W3C [26] that can operate over any lower level transport protocol like HTTP, TCP, UDP, etc. It is often used in combination with WSDL [27], an XML-based interface definition language that has been extensively utilized for describing functionality of SOAP-based web services.

SOAP was the first popular protocol for implementing platform agnostic interoperable web services. Every major software platform has tools to develop SOAP services. These powerful tools, on the one hand, can greatly simplify and accelerate the implementation, but on the other hand, they can be easily abused, which significantly hurts the interoperability. The problem is that software-generated WSDLs and the corresponding client proxies can be very complex, which results in interfaces that could potentially break down when communicating between different platforms. Also, it can be difficult to consume such services from mobile devices.

Support for different transport protocols made SOAP a good candidate for the implementation of enterprise services where interoperability was not the primary concern. SOAP services could take advantage of platform specific protocols and use interesting MEPs. For example, it is possible to implement a SOAP service that could take advantage of the fact that the service and its consumer are running on the same computer and use duplex MEP over very efficient platform-specific IPC protocol. On the other hand, it can have a detrimental impact on interoperability since such features may not be universally supported.

Recommendation 3.2 Use SOAP for Internal Platform-Specific Services

Recommendation	Use SOAP for the implementation of internal services when it is necessary to take advantage of more efficient binary bindings or MEP other than request-response.

Rationale	If communication with the service should be done over some platform specific (potentially more efficient) transport protocol, using SOAP services becomes a good implementation option since it still leaves the possibility of exposing service capabilities over HTTP binding to potential consumers on other platforms.

3.1.2 RESTful Services and Web APIs

REST (Representational State Transfer) is an architectural style for designing and building scalable distributed software systems that was formally described in Roy Fielding's doctoral dissertation [7]. It has gained a wide acceptance among implementers of web services as a simpler alternative to SOAP-based services. Because of its simplicity RESTful web services can be easily consumed from a variety of mobile devices and it was one of the major factors contributing to the popularity of REST.

In his dissertation Fielding defines REST architectural style based on a set of constraints. The idea is to understand forces that impact system behavior and then to identify constraints on the system design so that it works with those forces instead of against them. Here is the list of formal REST constraints as defined by Fielding.

- Client-Server

- Stateless

- Cache

- Interface / Uniform Contract

- Layered System

- Code-On-Demand

These constraints result in a set of practical guidelines that define how RESTful system should be built.

- Identify every resource with unique, global ID (URI).

- Use IDs to link resources. One resource can contain links to other resources. For example, order resource can have links (via URI) to customer and product resources.

- Use standard methods defined by HTTP protocol (GET, POST, PUT, DELETE, HEAD, etc.)

- Resources can have multiple representations. For example, GET request may be able to return representation of the resource in XML or JSON format depending on HTTP Content-type header.

- Use stateless communications. This means that after every request, the state should either be turned into a resource state or returned back to the client.

Not every type of API can be nicely represented using a resource-oriented approach. Some experts noted that a number of very successful internet companies have not been using "pure" REST APIs. William Vambenepe in one of his posts [28] pointed out that since Amazon's use of RPC over HTTP has not prevented them from becoming one of the most successful internet companies in the world, using REST is not a requirement to build a highly distributed and

robust internet system. He suggested that ultimately, simplicity of the API was more important than its RESTfulness.

Leonard Richardson developed a REST Maturity Model [29] that shows how a service can evolve into becoming a "real" RESTful service by adding certain RESTful characteristics.
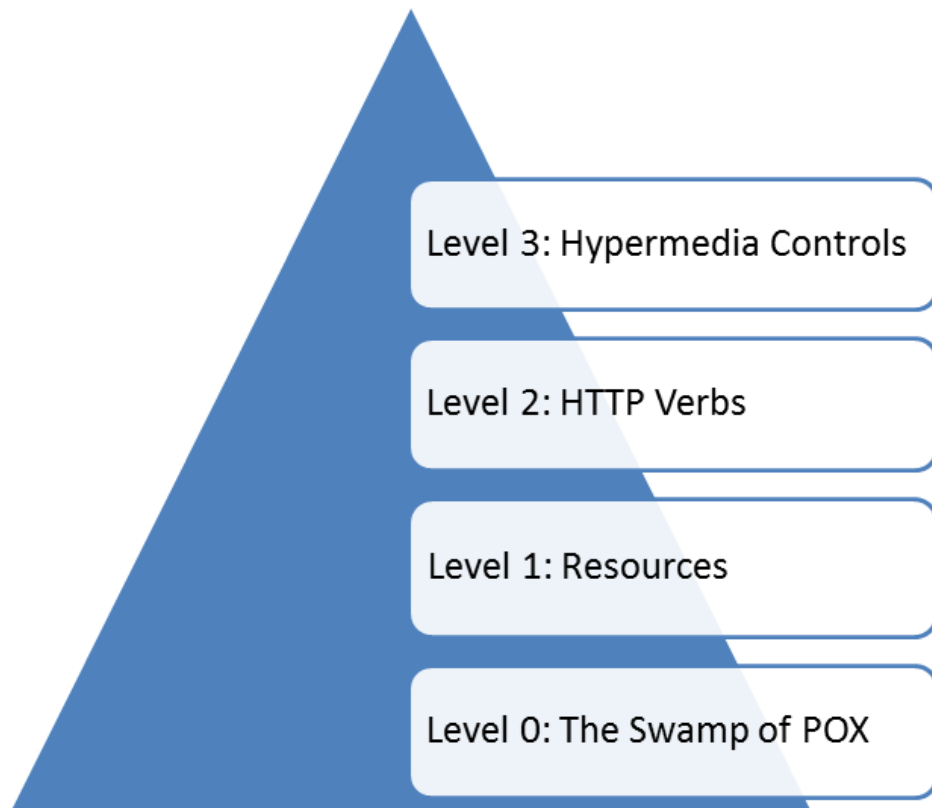
Figure 3.2 Richardson Maturity Model

From the maturity point of view, the lowest Level 0 is represented by Plain Old XML (POX) services that basically exchange XML messages over HTTP protocol. Conceptually it is very close to XML-RPC. At this level the service typically has a single URI and uses a single HTTP method like, for example, POST.

Level 1 adopts the resource-centric approach, which also implies identification of resources via URI. So at this level, the service typically has multiple URIs, but still uses a single HTTP method like Level 0.

Level 2 adds correct semantic usage of HTTP verbs. It means that services at this level use the unified HTTP interface. Now the service has multiple URIs and uses multiple HTTP verbs.

At the highest, Level 3 services use hypermedia to drive the application flow. It is sometimes referred to as HATEOAS (Hypermedia As The Engine Of Application State).

Recently the concept of pragmatic, as opposed to dogmatic, REST gained a lot of popularity. It suggests that even though it is a good idea to follow RESTful principals in general, sometimes it makes sense to deviate from strict REST rules to keep the API simple and easy to consume.

In recent years the concept of Web API (API that is reachable over internet) became very popular. Since Web APIs define the means that allow clients to consume service capabilities they, by definition, represent service interfaces.

Technically SOAP over HTTP is also a web API but it is not how the term is used in the industry. In fact, the term Web API strongly implies that it is not a SOAP-based interface. This is just another manifestation that the role of SOAP on the web continues to be diminished.

Simplicity is one of the key aspects of a Web API design. They are mostly based on a RESTful approach and can be thought of as pragmatic REST. Web APIs are typically resource-based and use HTTP verbs. They are usually implemented in a stateless manner and support content negotiation.

There is wide range of opinions about the usefulness of the "pure" REST approach. Web API design uses RESTful ideas to build simple interfaces reachable over the internet without necessarily following all the strict rules to be considered a "pure" RESTful interface. This pragmatic approach has gained a lot of support.

3.1.3 RESTful Web API vs SOAP-based RPC

To be able to provide a standardization approach that would be applicable to both SOAP and RESTful services, it is useful to compare them and try to identify potential commonality. The main challenge is that SOAP and RESTful service interfaces are defined very differently. The table below highlights some major differences.

SOAP RPC	RESTful Web API
Service contract is represented as a set of service operations. Operation is specified as a WS-Addressing action or as a wrapper element inside the message body.	Service contract is represented by the uniform interface. Method information is passed via HTTP header.
Operation semantics is defined out of band.	Method semantics in most cases is defined by the uniform interface.
Payload is defined by using XML schema and passed inside a SOAP envelope body.	Payload is tied to a media type and passed inside an HTTP message body.
Error semantics is defined out of band using SOAP faults that are also part of the interface.	Error semantics is defined by the uniform interface and returned back as standardized code.
Scoping information is defined as part of the payload.	Scoping information is defined by the URI. It could be a single resource or a collection of resources.

Best Practices for Services Implementation Using ARTS Standards

Single service endpoint exposes many operations.	Small set of HTTP methods is applied to resources referenced by multiple URIs.
Supports multiple transport protocols. If used with HTTP treats it just as a transport protocol.	Works over HTTP and treats it as application protocol.

The example below demonstrates the differences between SOAP and RESTful services. It is loosely based on simplified examples from ARTS XML Retail Transaction Interface Technical Specification [30]. This example shows the creation of a retail transaction that contains just one item and is paid for with cash.

The example consists of four interactions with services.

1. Begin transaction.

2. Add item to the transaction.

3. Total the transaction.

4. Pay with cash.

Step 1. Begin Transaction

SOAP Request:

POST /RTS HTTP/1.1

Host: www.example.org

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:TransactionBegin>
  <ARTSHeader>
   <MessageID>1234</MessageID>
   <DateTime>2015-01-17T09:30:47.0Z</DateTime>
   <BusinessUnit TypeCode="RetailStore">1001</BusinessUnit>
   <WorkstationID>Reg1</WorkstationID>
   <TillID>25</TillID>
  </ARTSHeader>
 </rts:TransactionBegin>
</soap:Body>

</soap:Envelope>
```

SOAP Response:

HTTP/1.1 200 OK

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
```

```
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:TransactionBeginResponse>
  <ARTSHeader>
   <MessageID>9876</MessageID>
   <DateTime>2015-01-17T09:30:47.0Z</DateTime>
   <Response>
    <RequestID>1234</RequestID>
   </Response>
  </ARTSHeader>
  <POSLog>
   <Transaction TypeCode="SaleTransaction"
     MajorVersion="4"
     MinorVersion="0"
     FixVersion="0" >
    <SequenceNumber>1011</SequenceNumber>
    <POSLogDateTime>2015-01-17T09:30:47.0Z</POSLogDateTime>
    <CustomerOrderTransaction TransactionStatus="InProcess"/>
   </Transaction>
  </POSLog>
 </rts:TransactionBeginResponse>
</soap:Body>

</soap:Envelope>
```

REST Request:

POST /RTS/transaction HTTP/1.1

Host: www.example.org

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<ARTSHeader>
 <MessageID>1234</MessageID>
 <DateTime>2015-01-17T09:30:47.0Z</DateTime>
 <BusinessUnit TypeCode="RetailStore">1001</BusinessUnit>
 <WorkstationID> Reg1</WorkstationID>
 <TillID>25</TillID>
</ARTSHeader>
```

REST Response:

201 Created

Location: http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<POSLog>
 <Transaction TypeCode="SaleTransaction"
   MajorVersion="4"
   MinorVersion="0"
   FixVersion="0" >
  <SequenceNumber>1011</SequenceNumber>
  <POSLogDateTime>2015-01-17T09:30:47.0Z</POSLogDateTime>
  <CustomerOrderTransaction TransactionStatus="InProcess"/>
 </Transaction>
</POSLog>
```

The SOAP request looks like a procedure call with parameters. The service creates the transaction and returns it back to the client.

Best Practices for Services Implementation Using ARTS Standards

The RESTful request is a request to create a transaction resource. So, it is done via POST to the transaction URI (www.example.org/RTS/transaction). Retail transaction service creates the transaction resource and then responds to the client with the location of this new resource in the Location HTTP header. For convenience, the service also puts the representation of the newly created resource in the response body. Any client can now view the transaction resource by sending GET request to the resource URI (http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011). It could be even achieved using an internet browser.

Step 2. Add Item to Transaction

<u>SOAP Request:</u>

POST /RTS HTTP/1.1

Host: www.example.org

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:LineItemAdd>
  <ARTSHeader />
  <POSLog>
   <Transaction TypeCode="SaleTransaction"
     MajorVersion="4"
     MinorVersion="0"
     FixVersion="0" >
    <SequenceNumber>1011</SequenceNumber>
    <CustomerOrderTransaction>
     <LineItem Action="Add">
      <SequenceNumber>1</SequenceNumber>
      <Sale>
       <ItemID>0430020006</ItemID>
       <Quantity>1</Quantity>
      </Sale>
     </LineItem>
    </CustomerOrderTransaction>
   </Transaction>
  </POSLog>
 </rts:LineItemAdd>
</soap:Body>

</soap:Envelope>
```

<u>SOAP Response:</u>

HTTP/1.1 200 OK

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0" encoding="UTF-8"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:LineItemAddResponse>
  <ARTSHeader />
  <POSLog>
   <Transaction TypeCode="SaleTransaction"
```

```
      MajorVersion="4"
      MinorVersion="0"
      FixVersion="0" >
    <SequenceNumber>1011</SequenceNumber>
    <POSLogDateTime>2015-01-17T09:30:47.0Z</POSLogDateTime>
    <CustomerOrderTransaction TransactionStatus="InProcess">
     <ItemCount>1</ItemCount>
     <LineItem>
      <SequenceNumber>1</SequenceNumber>
      <Sale TaxableFlag="true">
       <ItemID>0430020006</ItemID>
       <Description>Milk</Description>
       <RegularSalesUnitPrice>2.49</RegularSalesUnitPrice>
       <ExtendedAmount>2.49</ExtendedAmount>
       <Quantity>1</Quantity>
      </Sale>
     </LineItem>
    </CustomerOrderTransaction>
   </Transaction>
  </POSLog>
 </rts:LineItemAddResponse>
</soap:Body>

</soap:Envelope>
```

REST Request:

POST /RTS/transaction/2015-01-17-1001-Reg1-1011/item HTTP/1.1

Host: www.example.org

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0" encoding="UTF-8"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<Sale>
 <ItemID>0430020006</ItemID>
 <Quantity>1</Quantity>
</Sale>
```

REST Response:

201 Created

Location: http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/item/1

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<LineItem>
 <SequenceNumber>1</SequenceNumber>
 <Sale TaxableFlag="true">
  <ItemID>0430020006</ItemID>
  <Description>Milk</Description>
  <RegularSalesUnitPrice>2.49</RegularSalesUnitPrice>
  <ExtendedAmount>2.49</ExtendedAmount>
  <Quantity>1</Quantity>
 </Sale>
</LineItem>
```

Again, the SOAP request looks like a procedure call with parameters. The service adds a new item to the transaction and then returns the whole transaction back to the client. It is important to note that the POST request is sent to the same endpoint as in Step 1, which is not the case for the RESTful service.

Best Practices for Services Implementation Using ARTS Standards

The RESTful request is a request to create a new item resource under the transaction. So, it is done via a POST to the item under that specific transaction URI (http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/item). The retail transaction service creates the line item resource and then responds to the client with the location of this new resource in the Location HTTP header (http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/item/1). Again, the representation of the newly created resource is placed in the response body. The assumption here is that every line item is a separate resource and therefore it has its own URI and can be manipulated independently. If client need the whole transaction it can send a GET request to the transaction URI (http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011).

Step 3. Get Transaction Total

SOAP Request:

POST /RTS HTTP/1.1

Host: www.example.org

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:TransactionTotal>
  <ARTSHeader />
  <POSLog>
    <Transaction TypeCode="SaleTransaction"
      MajorVersion="4"
      MinorVersion="0"
      FixVersion="0" >
     <SequenceNumber>1011</SequenceNumber>
    </Transaction>
  </POSLog>
 </rts:TransactionTotal>
</soap:Body>

</soap:Envelope>
```

SOAP Response:

HTTP/1.1 200 OK

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0" encoding="UTF-8"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:LineItemAddResponse>
  <ARTSHeader />
  <POSLog>
  <Transaction TypeCode="SaleTransaction"
      MajorVersion="4"
      MinorVersion="0"
      FixVersion="0" >
   <SequenceNumber>1011</SequenceNumber>
   <POSLogDateTime>2015-01-17T09:30:47.0Z</POSLogDateTime>
   <CustomerOrderTransaction TransactionStatus="Totaled">
```

```
     <ItemCount>1</ItemCount>
     <LineItem>
      <SequenceNumber>1</SequenceNumber>
      <Sale TaxableFlag="true">
       <ItemID>0430020006</ItemID>
       <Description>Milk</Description>
       <RegularSalesUnitPrice>2.49</RegularSalesUnitPrice>
       <ExtendedAmount>2.49</ExtendedAmount>
       <Quantity>1</Quantity>
       <Tax TaxType="Sale">
        <SequenceNumber>1</SequenceNumber>
        <TaxableAmount
          TaxIncludedInAmountFlag="false">2.49</TaxableAmount>
        <Amount>0.12</Amount>
        <Percent>5.00</Percent>
       </Tax>
      </Sale>
     </LineItem>
     <Total TotalType="TransactionNetAmount">2.49</Total>
     <Total TotalType="TransactionTaxAmount">0.12</Total>
     <Total TotalType="TransactionGrandAmount">2.61</Total>
    </CustomerOrderTransaction>
   </Transaction>
   </POSLog>
 </rts:LineItemAddResponse>
</soap:Body>

</soap:Envelope>
```

REST Request:

POST /RTS/transaction/2015-01-17-1001-Reg1-1011 HTTP/1.1

Host: www.example.org

Content-Type: application/xml; charset=utf-8

Content-Length: nnn0

```
<?xml version="1.0"?
 xmlns="http://www.example.org/retailtransaction">
<rts:TransactionTotal />
```

REST Response:

HTTP/1.1 200 OK

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```
<POSLog>
<Transaction TypeCode="SaleTransaction"
    MajorVersion="4"
    MinorVersion="0"
    FixVersion="0" >
 <SequenceNumber>1011</SequenceNumber>
 <POSLogDateTime>2015-01-17T09:30:47.0Z</POSLogDateTime>
 <CustomerOrderTransaction TransactionStatus="Totaled">
  <ItemCount>1</ItemCount>
  <LineItem>
   <SequenceNumber>1</SequenceNumber>
   <Sale TaxableFlag="true">
    <ItemID>0430020006</ItemID>
    <Description>Milk</Description>
    <RegularSalesUnitPrice>2.49</RegularSalesUnitPrice>
    <ExtendedAmount>2.49</ExtendedAmount>
    <Quantity>1</Quantity>
    <Tax TaxType="Sale">
     <SequenceNumber>1</SequenceNumber>
     <TaxableAmount
       TaxIncludedInAmountFlag="false">2.49</TaxableAmount>
```

```
      <Amount>0.12</Amount>
      <Percent>5.00</Percent>
     </Tax>
    </Sale>
   </LineItem>
   <Total TotalType="TransactionNetAmount">2.49</Total>
   <Total TotalType="TransactionTaxAmount">0.12</Total>
   <Total TotalType="TransactionGrandAmount">2.61</Total>
  </CustomerOrderTransaction>
 </Transaction>

</POSLog>
```

There is no difference in the structure of the SOAP message. It is again a POST request to the same endpoint and the only thing that changes is the SOAP body that represents the invocation of a remote procedure.

The RESTful request looks very different for Step 3. The totaling transaction is an interesting operation. If it was just an update of the TransactionStatus field then it would be possible to use PUT or PATCH requests. However, it is really a much more complex change of the transaction state. Therefore to be consistent with the REST methodology we have to use POST request passing with it the type of action to be performed.

It might be useful to note that if we just wanted to get (calculate) totals on the transaction we could treat totals as a resource that was created automatically with the creation of the transaction. Then we could use a GET request to obtain the totals resource.

REST Request:

GET /RTS/transaction/2015-01-17-1001-Reg1-1011/total HTTP/1.1

Host: www.example.org

Content-Type: application/xml; charset=utf-8

Content-Length: 0

REST Response:

HTTP/1.1 200 OK

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```
<?xml version="1.0"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<Totals>
 <Total TotalType="TransactionNetAmount">2.49</Total>
 <Total TotalType="TransactionTaxAmount">0.12</Total>
 <Total TotalType="TransactionGrandAmount">2.61</Total>
</Totals>
```

If we treat totals as a resource then clients can send GET request to get totals any time after the transaction resource was created. The totals resource is represented by the following URI http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/total. It is also possible to add URIs for specific types of total:

http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/total/TransactionNetAmount

These examples show the power of the RESTful approach and how natural it is for exposing some types of service capabilities. For example, transaction status request from the RTI specification could be implemented in the same manner by just sending a GET request to the following URL http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/status.

Step 4. Pay in Cash

SOAP Request:

POST /RTS HTTP/1.1

Host: www.example.org

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:TenderAdd>
  <ARTSHeader />
  <POSLog>
   <Transaction TypeCode="SaleTransaction"
     MajorVersion="4"
     MinorVersion="0"
     FixVersion="0" >
    <SequenceNumber>1011</SequenceNumber>
    <CustomerOrderTransaction>
     <LineItem>
      <SequenceNumber>2</SequenceNumber>
      <Tender TenderType="Cash">
       <Amount>5.00</Amount>
      </Tender>
     </LineItem>
    </CustomerOrderTransaction>
   </Transaction>
  </POSLog>
 </rts:TenderAdd>
</soap:Body>

</soap:Envelope>
```

SOAP Response:

HTTP/1.1 200 OK

Content-Type: application/soap+xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0" encoding="UTF-8"?>
<soap:Envelope
 xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
<soap:Body xmlns:rts="http://www.example.org/retailtransaction"
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
 <rts:TenderAddResponse>
  <ARTSHeader />
  <POSLog>
  <Transaction TypeCode="SaleTransaction"
     MajorVersion="4"
     MinorVersion="0"
     FixVersion="0" >
    <SequenceNumber>1011</SequenceNumber>
    <POSLogDateTime>2015-01-17T09:30:47.0Z</POSLogDateTime>
    <CustomerOrderTransaction TransactionStatus="InProcess">
     <ItemCount>1</ItemCount>
```

```xml
        <LineItem>
          <SequenceNumber>1</SequenceNumber>
          <Sale TaxableFlag="true">
            <ItemID>0430020006</ItemID>
            <Description>Milk</Description>
            <RegularSalesUnitPrice>2.49</RegularSalesUnitPrice>
            <ExtendedAmount>2.49</ExtendedAmount>
            <Quantity>1</Quantity>
          </Sale>
        </LineItem>
        <LineItem>
          <SequenceNumber>2</SequenceNumber>
          <Tender TenderType="Cash">
            <Amount>5.00</Amount>
            <TenderChange>
              <Amount>2.39</Amount>
            </TenderChange>
          </Tender>
        </LineItem>
        <Total TotalType="TransactionNetAmount">2.49</Total>
        <Total TotalType="TransactionTaxAmount">0.12</Total>
        <Total TotalType="TransactionGrandAmount">2.61</Total>
      </CustomerOrderTransaction>
    </Transaction>
   </POSLog>
  </rts:TenderddResponse>
</soap:Body>

</soap:Envelope>
```

REST Request:

POST /RTS/transaction/2015-01-17-1001-Reg1-1011/tender HTTP/1.1

Host: www.example.org

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0" encoding="UTF-8"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<Tender TenderType="Cash">
 <Amount>5.00</Amount>

</Tender>
```

REST Response:

201 Created

Location: http://www.example.org/RTS/transaction/2015-01-17-1001-Reg1-1011/tender/1

Content-Type: application/xml; charset=utf-8

Content-Length: nnn

```xml
<?xml version="1.0"?
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/">
<LineItem>
 <SequenceNumber>2</SequenceNumber>
 <Tender TenderType="Cash">
  <Amount>5.00</Amount>
  <TenderChange>
   <Amount>2.39</Amount>
  </TenderChange>
 </Tender>
</LineItem>
```

Adding a tender is very similar to adding an item but there is an important nuance. When a cash payment is made in an amount exceeding the total amount of the transaction a very important change takes place. The transaction gets settled and that means that its transition from the modifiable customer order stage to an immutable retail transaction is completed. The

settled retrial transaction represents exchange of goods and/or services for a tender. So, from the RESTful point of view, creation of a new tender resource caused a fundamental change in the parent resource status. In fact, one could argue that a completely new resource was created.

Those four steps complete the transaction. The messages that were used to accomplish this very simple business process, using RESTful and SOAP-based approaches, have only one thing in common. They share common data structures representing different parts of the POSLog standard. In the SOAP-based interface, POSLog is one of the RPC parameters. In the RESTful approach, it is a top level resource. However, the data structures that are used in both approaches are pretty much the same. This data structure can be also used for POX and Queuing approaches.

Another important point is that it is not only a common piece among different approaches, but it is also the largest and the most complex one. It means that if standard representations of nouns (resources) are available and well-designed then moving from one type of interface to another is not that difficult and some services may choose to expose their capabilities using several different styles.

Recommendation 3.3 Create Standard Representations of Nouns

Recommendation	The focus of standardization efforts should be the creation of standard representations of nouns that can be passed as parameters using SOAP web services or used as a representation of resources for RESTful APIs. These standard representations should include clear and unambiguous definitions of all data elements to ensure a consistent interpretation of the data.
Rationale	Nouns are a common part for communicating with any type of service interface. They are also the largest and the most complex part of the messages that are exchanged with services. Standardized nouns significantly lower integration costs. It is much easier to switch from one type of service to another if the structure of the payload mostly stays the same than to deal with significantly different structures of messages or semantic inconsistencies.

3.1.4 Queues

Modern computer systems extensively use message queues to implement asynchronous communications between different sub-systems. In the context of this paper, message queues represent another type of interface to pass messages to a service.

This type of interface is very popular for certain types of services, especially where the immediate response is not required. Queues offer a number of interesting features.

First of all, due to their asynchronous nature they provide temporal decoupling between services and message senders. That reduces dependencies among the components of the distributed system. Also, many message queuing systems can ensure that messages do not get lost in the event of a system failure. It is achieved by securing a message in some kind of persistent storage. This behavior is critical for many retail systems that should provide guaranty that transactions or customer orders would not get lost even if the system experiences a failure.

Queues are often used as a buffer between message senders and a service in order to level load spikes that could potentially overwhelm the service. Since queues effectively decouple the service from the message senders, the service can process the messages at its own pace irrespective of the rate the messages are placed into a queue.
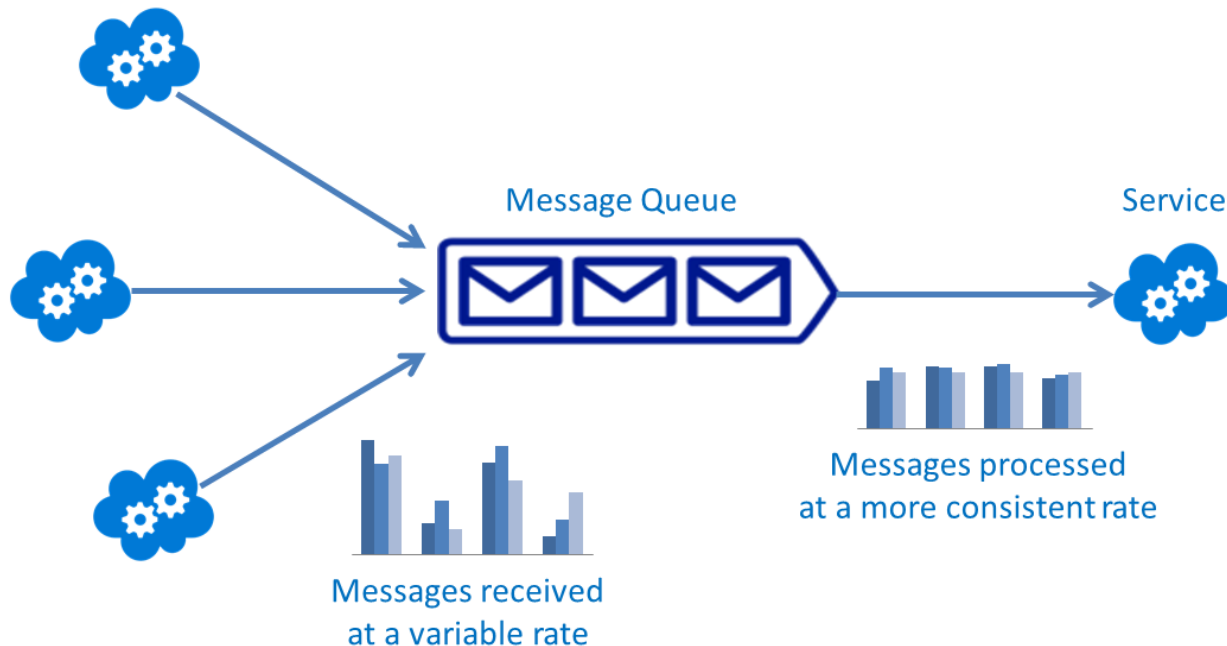
Figure 3.3 Queue-Based Load Leveling

Even though technically SOAP can use a message queuing system as the underlying communication protocol it is most commonly used to implement the request-response MEP. The RESTful APIs utilize HTTP and therefore they also employ the request-response MEP. Due to the nature of request-response it implies point-to-point communications. Queuing systems can be also used for point-to-point communications but they are also able to implement a publish-subscribe pattern.

This pattern is commonly used to propagate events in a distributed system in a loosely coupled manner. Many message queuing systems provide topic-based publishing and subscribing where publishers can associate each message with a topic addressing them to recipients that subscribed to the topic. This type of messaging architecture allows sending messages only to the subscribers that are interested in receiving the messages without having to know their identities.

Queuing systems often play a central role inside different MOM (Message-Oriented Middleware) products, such as services buses or message brokers.

Traditionally queuing systems have used proprietary protocols. Consequently the integration between queuing systems developed by different vendors was a challenge. To overcome these complexities AMQP (Advanced Message Queuing Protocol) [31] was designed and approved as an OASIS standard in 2012.

AMQP 1.0 is an open, standard, application layer protocol for MOM systems. It is an efficient, binary standard that can support a wide variety of messaging applications and enables interoperation and communication and sharing of resources. AMQP supports different broker architectures, which may be used to receive, queue, route, and deliver messages.

Features of AMQP 1.0 include:

- Efficient Wire Protocol

- Supports Multiple Broker Architectures

- Message Security

- Global Addressing

- Extensible Layering

- Support for Multiple Messaging Systems

AMQP can provide different message delivery guarantees, such as *at-most-once*, *at-least-once*, and *exactly-once*.

The approach based on standard representations of nouns is a very good fit for queue-based integration. Indeed, queue messages often represent certain business entities like, for example, new customer or retail transactions. They are also often used to propagate important business events like CustomerModified or PromotionAdded, where nouns are the biggest and the most important part of the payload. Therefore, standardized representations of nouns are crucial for achieving seamless queue-based integration between different components of the distributed system.

3.2 Service Information Model

Service information model is a detail description of the data that can be exchanged with the service. It includes structure and format of the information exchanged between the service and its consumers. It is important to note that consistent semantic interpretation of the exchanged data is very important especially if service interactions cross ownership boundaries.

The information model specifies all the data that is necessary for successful interactions with the service. It includes the structure of the information, its semantics, actions that can be performed against the service, and sometimes even dependencies among different service invocations.

As was demonstrated earlier in this chapter certain parts of the service information model depend on the type of service interface. On the other hand, the data structures and their business domain semantics can often be used consistently by different types of service interfaces. This data represents the exchange of business information that is necessary to complete a certain capability of the service.

In RESTful Web APIs these data structures correspond to the resources, whereas in SOAP-based RPC services they represent the parameters and the return data of the RPC.

Recommendation 3.3 suggests that to create standards that can be used with different styles of services the focus should be on nouns and their semantic definitions.

It is important to note that data structures can be represented using different formats. In the past ARTS used XML schema language to define standard service interface data structures.

This approach limits data representation format to XML. To be able to support other data representation formats a new methodology to specify standard data structures is required.

3.2.1 Data Serialization Formats

For years XML was the primary message format for communicating information between applications, and between applications and devices. ARTS has built an extensive library of XML schemas for standardizing these communication interfaces for retail industry. But with the advent of mobile and the limited bandwidth for communicating with mobile devices, it was decided that a lighter weight communication model was needed.

Before XML a CSV (Comma-Separated Values) format was widely used in things like EDI. The problem with CSV format is understanding what was meant by the data between each comma. Besides, CSV data was difficult to use to represent hierarchical structures. JSON (JavaScript Object Notation) came about as a compromise. It contains some of the benefits of XML and the lighter weight formatting of CSV.

3.2.2 XML

ARTS started creating XML schemas in 1999. This was two years before the XML Schema format became an approved W3C recommendation in May 2001. Since then ARTS has produced over 20 different XML schemas covering almost every interface on the entire operational side of retail. These standards were created in collaboration with over 1000 retailers and vendors. They have been downloaded in almost every industrialized country in the world by both large and small companies. On top of that, there are very successful products and infrastructures built solely on the breadth and depth of these standards.

When ARTS started this work, retail was a lagging technology innovator. The cost of changing systems was just too prohibitive to respond quickly to any changes in technology. In part due to breaking the tight coupling between systems through ARTS standards, today retail is shifting to a leading technology innovator.

The beauty of XML is its self-describing characteristic. That is both sides of the communication channel can understand each other without knowing who originated the message. This is strengthened by the use of XML schemas to properly define the message format. The following is a list of current XML schemas:

ARTS Standard	Description
Associate Management	Defines information necessary to manage an associate in scheduling their participation.
Change 4 Charity	Identifies the data need to communicate real-time charity contributions.
Comparison Shopping Engine	Communicates with CSE's to request information about a particular item(s).
Compliance Audit Interchange	Shares industry audits of factories concerning their compliance with human, safety, etc. requirements.

Customer	Publishes information about a customer contact, demographic, loyalty and target marking potentials.
Digital Receipt	One of ARTS first standards, electronically communicates in effect the traditional paper receipt .
Fresh Item Management (Traceability)	Tracks and communicates information about recalls on either food or items.
Inventory	Reports inventory positions.
Item Maintenance	Contains an extensive list of retail item attributes. It synchronizes with GS1's item definitions for communicating B2B item information such as pallets. This allows following an item from the factory to the customer.
Kitchen	Shares information about kitchen equipment (grills, fryers, etc.) to the operational side.
Location	Helps one identify the location for virtually all items within the store.
Product Content Management	Reports image relates information about items. This is the visual equivalent of the Item Maintenance schema.
POSLog	The heart of retail. This is virtually all information about the sale/return of items. In effect a superset of a Retail Transaction. This takes the Digital Receipt schema and adds all kinds of additional information around the sales process, including issuing POS events and reconciling the POS at the end of a period.
Price	The interface to a price engine. It allows the sending of price rules to the price lookup unit and returning the calculated price of the item during the sales process.
Remote Equipment Monitoring and Control	Allows remote monitoring of equipment such as the POS.
Stored Value	Used to communicate information about any stored value card such as debit cards, gift cards, etc.
Time Punch	The second schema used to manage associates by reporting their punching in and out.
Transaction Tax	There are two components to this schema. One is used to populate the tax engine with appropriate tax rules. The second is the calculation of the taxes related to a particular sell. This is one

	of many ARTS schemas which support the omnichannels sales models.
Video Analytics	This schema is used to provide information that can be used by the analytics engine to evaluate performance issues, such as inventory management, loss prevention, sales analysis, etc.
XMLPOS	There are 72 schemas in this set. They convert the UnifiedPOS API's into XML for remote communication of information and control around the 36 devices connected to the POS.

3.2.3 JSON

Just a decade ago XML was definitely the primary data interchange format. But in the last several years the simpler and bandwidth-non-intensive JSON format emerged as an attractive alternative to XML. JSON is getting more and more popular to the point that some analysts begin to believe that it may eventually become XML's successor. Beside pure technical factors that contributed to JSON's popularity there are also certain historical forces that are working against XML.

First of all, JSON is the format of choice for many popular APIs, some of which no longer offer support for XML. Second, JSON often is used a as data representation format for NoSQL databases. And finally, the IoT (Internet of Things) and mobile devices are mostly using JSON format to exchange data. All these factors contribute to the growing popularity of the JSON format.

It is important to note that ARTS standards could be expressed in terms of JSON. ARTS will use two different approaches to provide support for JSON. First, the mapping-based approach will be used to enable JSON format for previously developed XML schemas. Second, for all new standards, we recommend a new design process which includes the development of an abstract information model which will later be used to derive representation of this information into different formats like XML, JSON, and relational.

3.2.3.1 Dealing with Existing ARTS XML Standards

Since ARTS already has an extensive library of standard XML schemas developed over many years, it would be very difficult to go back and start adding JSON definitions to all the existing standards. A more pragmatic approach to adding JSON support would be to provide some kind of mapping between standard XML documents and their JSON representation.

In general the task of mapping arbitrary XML to JSON can be fairly complex since there are two conflicting goals: round-tripping and simplicity of the resulting JSON. There are some other challenges like data types mismatch, dealing with collection, etc.

Round-tripping means that converting a document from one format to another and then performing a reverse conversion results in exactly the same document as the original. It is an important consideration since it implies that there is no loss or distortion of information during the conversion process. To enable generic round-tripping (XML => JSON => XML) requires storing additional metadata as part of the JSON document. This extra metadata makes JSON documents more difficult to read and to use.

Best Practices for Services Implementation Using ARTS Standards

On the other hand, our goal is to be able to represent standard ARTS XML documents using JSON format. Since we can assume that the XML schema of the ARTS XML document is known we can easily leverage that schema rather than trying to embed all the metadata inside the JSON document itself. This means that we could produce fairly simple JSON documents that can be interpreted correctly with the help of the corresponding standard XML schemas.

Below is a set of rules to convert a standard ARTS XML document to the JSON format. Simplified examples from the existing ARTS technical specifications are used for illustration purposes.

Rule 1: XML element names become JSON keys.

XML:

```
<RetailStoreID>HighStreet</RetailStoreID>
```

JSON:

```
"RetailStoreID": "HighStreet"
```

Rule 2: XML attribute names also become JSON keys but are prefixed with @@ to preserve the information that they are attributes. If converter knows the XML structure then the @@ symbol becomes optional. It is still useful to be able to easily see which fields were attributes inside the original XML. The double @ prefix is used to avoid collision with many JSON frameworks that use the single @ character for special keys.

XML:

```
<Sale ItemType="Stock" />
```

JSON:

```
"Sale": { "@@ItemType": "Stock" }
```

Rule 3: To represent XML complex types with simple content use a special "keyword" #value. This keyword represents the value of the simple content.

XML:

```
<Quantity UnitOfMeasureCode="Each">1</Quantity>
```

JSON:

```
"Quantity": { "@@UnitOfMeasureCode": "Each",  "#value": 1 }
```

Rule 4: XML children elements become JSON object.

XML:

```
<Sale ItemType="Stock">
 <POSIdentity>
  <POSItemID>01234567890123</POSItemID>
 </POSIdentity>
</Sale>
```

JSON:

```
"Sale": {
 "@@ItemType": "Stock",
 "POSIdentity": {
  "POSItemID": "01234567890123"
 }
}
```

Best Practices for Services Implementation Using ARTS Standards

Rule 5: Multiple child XML elements with the same name become a JSON array. In other words, elements that are specified in XML schema with maxOccurs attribute set to more than 1, including "unbounded", should be represented as a JSON array.

XML:

```xml
<LineItem>
 <SequenceNumber>1</SequenceNumber>
 <Sale ItemType="Stock" />
</LineItem>
<LineItem>
 <SequenceNumber>2</SequenceNumber>
 <Tender TenderType="Cash" />

</LineItem>
```

JSON:

```json
"LineItem": [
{
 "SequenceNumber": 1,
 "Sale": {"@@ItemType": "Stock"}
},
{
 "SequenceNumber": 2,
 "Tender": {"@@TenderType": "Cash"}
}

]
```

Rule 6: The default XML namespace is omitted from JSON document. The ARTS standard namespace http://www.nrf-arts.org/IXRetail/namespace is implied. The non-default XML namespaces that are used as part of the ARTS extensibility approach can be specified inside the JSON extensibility object using the special #namespace "keyword".

XML:

```xml
<Sale ItemType="Stock">
 <ItemID>011111</ItemID>
 <RegularSalesUnitPrice>50.0000</RegularSalesUnitPrice>
 <ExtendedAmount>50.00</ExtendedAmount>
 <Quantity>1</Quantity>
 <Tax>
  <TaxAuthority>3</TaxAuthority>
  <TaxableAmount TaxIncludedInTaxableAmountFlag="false">50.00</TaxableAmount>
  <Amount>3.00</Amount>
  <Percent>6.000000</Percent>
 </Tax>
 <SaleExtension xmlns="http://www.mycompany.com/artsxml">
  <ComparePrice>50.0000</ComparePrice>
 </SaleExtension>

</Sale>
```

JSON:

```json
"Sale": {
   "@@ItemType" = "Stock",
   "RegularSalesUnitPrice" = 50.0000,
   "ExtendedAmount" = 50.00,
   "Quantity" = 1,
   "Tax" : {
     "TaxAuthority" : 3,
```

```
    "TaxableAmount" : {
        "TaxIncludedInTaxableAmountFlag" = false,
        "#value" = 50.00
    },
    "Amount" = 3.00,
    "Percent" = 6.000000
},
"SaleExtension" : {
    "#namespace" : "http://www.mycompany.com/artsxml",
    "ComparePrice" : 60.0000
  }
}
```

Rule 7: Simple XML data types should be represented in JSON document using the following mapping table.

XML Data Type	JSON Data Type
numeric types: float, double, decimal, integers	number
string-based types	string
date, time, and dateTime	string in ISO 8601 [32] format
boolean	boolean
xsi:nil = "true"	null

Some JSON processors may deserialize decimal data types into floating-point numbers, which might introduce rounding errors. In this case it is acceptable to represent decimal data types as strings. Then these strings should be converted to decimals inside the code that accepts the JSON document.

These seven simple mapping rules do not cover generic mapping between XML and JSON but can be used to produce representation of a standard ARTS XML document in the JSON format. Below is an example of applying these rules to a POSLog XML document.

XML:

```
<?xml version="1.0" encoding="utf-8"?>
<POSLog
 xmlns="http://www.nrf-arts.org/IXRetail/namespace/"
 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
 xsi:schemaLocation="http://www.nrf-arts.org/IXRetail/namespace/POSLog.xsd">
<Transaction>
 <RetailStoreID>HighStreet</RetailStoreID>
 <WorkstationID>POS5</WorkstationID>
 <SequenceNumber>7295</SequenceNumber>
 <OperatorID>John</OperatorID>
 <RetailTransaction Version="2.2">
  <LineItem>
   <SequenceNumber>1</SequenceNumber>
   <Sale ItemType="Stock">
    <POSIdentity>
     <POSItemID>01234567890123</POSItemID>
    </POSIdentity>
```

```xml
    <ExtendedAmount>4.89</ExtendedAmount>
    <Quantity UnitOfMeasureCode="Each">1</Quantity>
    </Sale>
    </LineItem>
    <LineItem>
    <SequenceNumber>2</SequenceNumber>
    <Tender TenderType="Cash" TypeCode="Sale">
    <Amount>4.89</Amount>
    </Tender>
    </LineItem>
 </RetailTransaction>
</Transaction>

</POSLog>
```

JSON:

```json
"POSLog": {
 "Transaction": {
 "RetailStoreID": "HighStreet",
 "WorkstationID": "POS5",
 "SequenceNumber": 7295,
 "OperatorID": "John",
 "RetailTransaction": {
 "@@Version": "2.2",
 "LineItem": [
 {
 "SequenceNumber": 1,
 "Sale": {
 "@@ItemType": "Stock",
 "POSIdentity": {
 "POSItemID": "01234567890123"
 },
 "ExtendedAmount": 4.89,
 "Quantity": {
 "@@UnitOfMeasureCode": "Each",
 "#value": 1
 }
 }
 },
 {
 "SequenceNumber": 2,
 "Tender": {
 "@@TenderType": "Cash",
 "@@TypeCode": "Sale",
 "Amount": 4.89
 }
 }
 ]
 }
 }
}
```

If the JSON-to-XML converter knows that the JSON document above represents a POSLog then it can create the corresponding POSLog XML document and add namespaces.

This mapping approach does not require creation of JSON schemas. The XML is treated as a primary standard and the JSON documents are produced by derivation by applying the mapping rules.

3.2.3.2 Dealing with New ARTS Standards

The approach proposed for the existing ARTS XML standards could work for the future ARTS projects as well, but treating JSON as a secondary standard fails to recognize important trends in the industry where the JSON format is becoming more and more popular. For this reason, ARTS needs a new approach creating the information model. This new approach will provide native support for both, JSON and XML formats.

3.2.4 Standard Data Structures

According to the Recommendation 3.3, defining standard data structures for nouns (business entities, business concepts, etc.) should be the major focus of ARTS standardization efforts. These structures should be defined in the context of a service that offers certain business capabilities. The context helps to provide clear semantic definitions of all the data elements and to validate them via use cases.

The main goal is to create data structures that can be easily represented using both XML and JSON formats, with the primary focus on interoperability. It means that structures should be simple enough so that their XML and JSON representations could be easily generated using different software development platforms.

To achieve this goal, data structures can be modeled in a representation-agnostic fashion and then mapped to JSON or XML schemas, or even relational data structures. Such an approach would guarantee consistency between different artifacts produced by a work team.

One of the popular representation-agnostic modeling formats is UML (Unified Modeling Language) [33]. UML supports modeling of data structures like Classes. These UML structures can be mapped to both, XML schema complex types and JSON schema objects. Therefore, UML is a good candidate for representation-agnostic modeling language.

David Carlson created a special UML profile for XML Schema that was described in his book Modeling XML Applications with UML [34]. It proves that UML is suitable for modeling different data structures.

Some modeling tools, for example, Enterprise Architect [35], support generation of XML schemas using the UML Profile for XML Schema.

The proposed modeling process consists of two steps. The first step is creation of UML diagrams that represent the information model. The second step is the generation of XML and/or JSON schemas that can be used for a formal description of the model created in the first step.

3.2.5 Reconciling Service Information Model with ARTS Data Model

There is a significant difference between message data structures that are part of the service information model and relational data structures that are used to store data in a database. Not everything in the message has to necessarily end up in the database. And vice versa, databases often contain a lot of information that is used to facilitate service functionality, but is not a part of the service interface.

Even data elements of a service interface that are serialized in a database can be persisted in data structures with a considerably different shape. For example, a message that creates a customer order can be stored as a customer order control transaction and also as a set of records representing the current state of the newly created customer order.

It is not uncommon that service interface information model and enterprise data model are specified at a different level of abstraction. For example, a layaway plan created by a special type of layaway service can be treated behind the scene as a particular type of a more abstract concept of a customer order.

To maintain data consistency and facilitate integration between different sub-systems it is crucial to analyze and reconcile any differences between the service interface data definitions and the ARTS ODM (Operational Data Model) [36]. Such analysis should be performed jointly by members of the work team creating the service information model and by members of the ARTS Data sub-committee.

The goal of the analysis is to guarantee the consistency in the following way.

1. If a data element has the same name in the service information model and the ODM then it should have exactly the same semantic meaning.

2. If a data element with the same semantic meaning has a different name in the service information model because that name is more descriptive and well-understood in the context of the service, such a discrepancy should be clearly documented. It can happen if, for example, a service business domain uses different terminology than more generic and abstract terminology in the ODM.

3. All service information model data elements that are not represented in the ODM should be clearly defined and documented with references to the corresponding use cases. This information is necessary for the data dictionary and the unambiguous interpretation of the message data.

3.3 Service Capabilities

A service capability is a unit of functionality exposed by a service. For example, a service can offer a capability to add a new customer or to calculate a transaction tax. Often, service capabilities require input data in order to provide the desired functionality. They also may return some data back to the service consumers.

The term "service capability" is technology agnostic. However, the way service capabilities are exposed to the service consumers depends on the particular implementation technology.

3.3.1 Designing Capabilities of SOAP Services

SOAP web services expose their capabilities as service operations. WSDL (Web Services Description Language) [27] has the operation element that contains input and output elements. The structure of the operation input and output is defined in terms of the XML schema.

Since interoperability is the primary goal of implementing standard interfaces the design of a SOAP service should follow the Web Services Interoperability (WS-I) [37] guidelines.

Recommendation 3.4 Use WS-I Guidelines for the Design of SOAP Services

Recommendation	When designing SOAP services use WS-I guidelines.
Rationale	WS-I provides a set of implementation and interoperability guidelines. These recommendations have gone through thorough testing and therefore increase the chance that the service will be accessible from a different platform.

3.3.1.1 SOAP Service Capabilities Naming Guidelines

Since an operation of a SOAP-based service typically represents a remote invocation of a method, the service API conceptually can be mapped to a class interface in a high level programing language like Java or C#.

So conceptually designing a SOAP-based API is similar to designing an interface for a class, where a service operation is similar to a method. One significant difference is that a single operation typically results in two distinct messages: request and response. The most common approach is to wrap the request parameters in an element that has the same name as the service operation and wrap the response data in an element with the name as a combination of the operation name and the Response suffix. The example below shows how an operation named ItemPriceGet is mapped to a SOAP request and response.

SOAP request:

```
POST /PriceLookup HTTP/1.1
Host: www.example.org
Content-Type: application/soap+xml; charset=utf-8
Content-Length: nnn

<?xml version="1.0"?>
<soap:Envelope
xmlns:soap="http://www.w3.org/2003/05/soap-envelope"
>

<soap:Body xmlns:m="http://www.example.org/price">
  <m:ItemPriceGet>
    <m:ItemID>02884562323433</m:ItemID>
  </m:ItemPriceGet>
</soap:Body>

</soap:Envelope>
```

SOAP response:

```
HTTP/1.1 200 OK
Content-Type: application/soap+xml; charset=utf-8
Content-Length: nnn

<?xml version="1.0"?>
<soap:Envelope
xmlns:soap="http://www.w3.org/2003/05/soap-envelope"
>

<soap:Body xmlns:m="http://www.example.org/price">
  <m:ItemPriceGetResponse>
    <m:Price>34.5</m:Price>
  </m:ItemPriceGetResponse>
</soap:Body>

</soap:Envelope>
```

Best Practices for Services Implementation Using ARTS Standards

The ARTS SOA Best Practices technical report [1] provided the following guidelines for naming service capabilities.

Guideline	Examples
Names should be composed of words in the English language, using primary English spelling in Oxford English Dictionary.	ItemColor
Names should be drawn from the following character set: a-z, A-Z, 0-9.	Track2Data2Get
SOAP service operation name should have the following structure: <ZeroOrMoreAdjectives>+<Noun>+<Verb>	PromotinalItemAdd DiscountCalculate
Readability is more important than length/	
Only commonly accepted abbreviations should be used and they should appear in all UPPERCASE. Some approved abbreviations: UCC, EAN, UPC, SKU, ID, GTIN, PLU, ISBN, ISSN, RFID, MICR, POS, PO, and ASN.	ItemPLUPriceGet

Therefore according to the SOA Best Practices document the names of capabilities of a SOAP service should end with a verb.

The table below contains the list of commonly used verbs in retail to ensure reuse and consistency.

Verb	Implication	Example
Add	Adding something to the document (e.g., adding an item to a transaction). Change in state indicated. Can also designate mathematical addition.	TransactionAdd TaxAmountAdd
Adjust	Modifies typically numerical data to achieve a desired or correct value. Change in state indicated.	ItemQuantityAdjust
Approve	Give/seek disposition to proceed	POApprove
Authorize	Get permission to perform certain action.	ItemReturnAuthorize

Begin	The start of a series of actions (e.g., starting a transaction). Change in state may be indicated.	TransactionBegin
Calculate	Process business algorithms or flow logic. No state is altered except for in memory view of the session state.	TaxCalculate
Cancel	Cancels a sequence before it is complete. Change in state may be indicated.	TransactionCancel
Complete	Reports a sequence finished. Change in state indicated. Indicates ACID transactional context.	TransactionComplete
Certify	Proof of eligibility.	CustomerEligibilityCertify
Confirm	Proof of validity of fact or assumption.	AvailabilityConfirm
Dispatch	Instructs that message associated with the noun should be sent to a destination. No change in state.	TransactionDispatch
Maintain (Create/Update/Delete)	Add, remove, or modify data. Change in state indicated. Indicates ACID transactional context.	AllocationCreate AllocationUpdate AllocationDelete
Obtain/Get/Read	Retrieves information associated with a noun. No change in state.	AllocationRead
Override	Change the value generated by the system. Typically requires authorization.	ItemPriceOverride
Perform	Instructs to perform an action. Change in state may occur. Frequently indicates a long-running transactional context with compensatory transactions needed to maintain integrity of the state.	PhysicalInvetoryPeform
Preview	Present data for review before going further with the business process. No state is altered except for in memory view of the session state.	BulkOrderPreview

Receive	Accept merchandise or funds. Change in state indicated.	StockItemReceive
Remove	Removes data element from a larger dataset. Change in state indicated.	PaymentInformationRemove
Reserve	Requests to reserve merchandise to insure successful completion of a business process. Change in state indicated.	MerchandiseReserve ItemReserve
Resume	Begin business activity again after it was suspended before.	RetailTransactionResume
Search/Lookup	Informational inquiry that allows open-ended browsing of information. No change in state. It is possible to return and empty set.	CatalogItemLookup
Send	Arrange for delivery of merchandise or mail.	CustomerEmailSend
Subtract	Takes away certain amount of value indicated by a noun.	LoyaltyPointsSubtract
Suspend	Stop business activity with the option to resume it lately.	RetailTransactionSuspend
Validate	Validates the data. No change in state.	AddressValidate
Void	Requests to reverse of a previously completed activity. Change in state indicated.	TransactionVoid

3.3.1.2 SOAP Service Error Handling

SOAP services use a special platform-independent mechanism to describe errors. When an error happens, a special message containing a SOAP Fault element is sent back to the service consumer.

SOAP Faults are defined in WSDL and therefore they are a part of the service interface. They should be designed to provide consumers with useful information about the error and at the same time abstract the consumers from implementation details. For example, if the service operation encounters a primary key violation when trying to insert an item into a database, that is not the error that should be communicated back to the client. It makes much more sense to inform the service consumer that the item already exists rather than to pass through the low level database exception.

If the ItemPriceGet operation cannot find the item with the specified ItemID then the operation could return a fault indicating that the item was not found.

SOAP request:

```
POST /PriceLookup HTTP/1.1
Host: www.example.org
Content-Type: application/soap+xml; charset=utf-8
Content-Length: nnn

<?xml version="1.0"?>
<soap:Envelope
xmlns:soap="http://www.w3.org/2001/12/soap-envelope"
>

<soap:Body xmlns:m="http://www.example.org/price">
  <m:ItemPriceGet>
    <m:ItemID>02884562323433</m:ItemID>
  </m:ItemPriceGet>
</soap:Body>
</soap:Envelope>
```

SOAP response:

```
HTTP/1.1 200 OK
Content-Type: application/soap+xml; charset=utf-8
Content-Length: nnn

<?xml version="1.0"?>
<soap:Envelope
xmlns:soap="http://www.w3.org/2003/05/soap-envelope"

xmlns:xml="http://www.w3.org/XML/1998/namespace"
>

<soap:Body xmlns:m="http://www.example.org/price">
  <soap:Fault>
    <soap:Code>
      <soap:Value>soap:Receiver</soap:Value>
      <soap:Subcode>
        <soap:Value>m:ErrorItemNotFound</soap:Value>
      </soap:Subcode>
    </soap:Code>
    <soap:Reason>
      <soap:Text xml:lang="en">Item 02884562323433 not found.</soap:Text >
    </soap:Reason>
  </soap:Fault>
</soap:Body>
```

More information about SOAP faults can be found in the SOAP Version 1.2 specifications [26].

3.3.1.3 SOAP Service Considerations

The SOAP service represents a group of capabilities that have a common functional context. It could be a service that maintains a certain business entity (for example, Customer

Maintenance Service) or a service that performs a certain task (for example, Tax Calculation Service). The service has to have a clearly defined functional scope.

A SOAP-based approach is more suitable for implementation of internal enterprise services where interoperability is not a primary concern. They can provide better performance due to the use of efficient proprietary bindings. Also, some bindings have additional features like, for example, duplex channels that can provide certain functionality that is not available over HTTP-based communications.

On the other hand SOAP services should not be used on the edge of the retail enterprise where the APIs are exposed to the general public or business partners. SOAP-based APIs are more difficult to consume and to integrate with.

3.3.1.4 SOAP Service Description

SOAP services use a special XML-based interface definition language known as WSDL [27]. The current version of WSDL is 2.0. It is a W3C recommendation. However, the support for version 2.0 is still poor and many tools support only WSDL 1.1.

Regardless of the version, WSDL is the only commonly accepted method for describing SOAP-based services.

3.3.1.5 SOAP Service Description Namespaces

The data structures representing payloads of SOAP services are described in WSDL using the XML schema language. All standard ARTS XML schemas have the targetNamespace attribute set to http://www.nrf-arts.org/IXRetail/namespace/. This approach simplifies the reuse of common data structures in different ARTS schemas.

However, the targetNamespace attribute inside the definition element inside WSDL should be set to http://www.nrf-arts.org/IXRetail/namespace/service_name. This namespace qualifies names of the elements inside the WSDL definition (message, portType, etc.). Therefore, the names of the service operations will be defined within the scope of a particular service.

Recommendation 3.5 Add Service Name to ARTS Namespace inside WSDL Definition

Recommendation	Target namespace for the WSDL definition should be http://www.nrf-arts.org/IXRetail/namespace/service_name/.
Rationale	This approach guarantees that every service operation name is unique among all services. For example, ItemAdd operation has completely different semantic meaning inside Shopping Basket and Item Maintenance services and using different namespaces provides two different contexts for correct interpretation of the operation name.

Even if two different services, for example Shopping Basket and Item Maintenance, have an operation named ItemAdd, different namespaces allow unambiguous interpretation of the SOAP messages.

Shopping Basket service ItemAdd SOAP request:

```xml
<?xml version="1.0"?>
<soap:Envelope
xmlns:soap="http://www.w3.org/2003/05/soap-envelope">
```

```
<soap:Body
  xmlns:svc=" http://www.nrf-arts.org/IXRetail/namespace/ShoppingBasket/"
  xmlns:data=" http://www.nrf-arts.org/IXRetail/namespace/"
>
  <svc:ItemAdd>
   <data:ItemID>02884562323433</data:ItemID>
  </svc:ItemAdd>
</soap:Body>

</soap:Envelope>
```

Item Maintenance service ItemAdd SOAP request:

```
<?xml version="1.0"?>
<soap:Envelope
xmlns:soap="http://www.w3.org/2003/05/soap-envelope"
>

<soap:Body
  xmlns:svc=" http://www.nrf-arts.org/IXRetail/namespace/ItemMaintenance/"
  xmlns=" http://www.nrf-arts.org/IXRetail/namespace/">
  <svc:ItemAdd>

   <Item>
    <ItemID>02884562323433</ItemID>

    <Description>Denim Jacket</Description>

    <Price>26.99</Price>

   <Item>
  </svc:ItemAdd>
</soap:Body>

</soap:Envelope>
```

3.3.2 Designing Capabilities of RESTful Services

The RESTful API typically exposes its capabilities via manipulation of resources. In SOAP-based services capabilities usually represent actions. According to the naming guidelines, the name of every SOAP service operation ends with a verb. On the other hand, RESTful APIs are centered around nouns.

3.3.2.1 Noun-Based API

Typically, for every resource (noun) there are two base URIs: one for a collection of resources and another for a specific resource in the collection. For example, if Customer is a resource than URI /Customer represents the collection of all customers and URI /Customer/92371 represent the specific customer with ID = 92371.

It is important to note that it is very common to use plural (Customers) inside the resource URIs. This document follows the same naming convention as the ARTS Data Model where all entities are presented in singular form unless the concept itself is plural.

In the RESTful approach resources can be operated on using HTTP verbs. The four major HTTP verbs are POST, GET, PUT, and DELETE. They can be loosely mapped to CRUD (Create-Read-Update-Delete).

These four verbs and two base resources represent an intuitive set of capabilities.

Resource	POST	GET	PUT	DELETE
/Customer	Create a new customer	Return all customers	Bulk update all customers	Delete all customers
/Customer/92371	Disallowed in most cases	Return a specific customer	Update a specific customer	Delete a specific customer

The simple convention described in the table above creates a consistent and intuitive way to manipulate the resources, which makes the REST APIs easy to understand and consume.

It is important to note that in most practical implementations "Delete all customers" action would be disallowed and return an error. It is presented here to demonstrate the completeness of the approach.

3.3.2.2 HTTP Status Codes

The status codes defined by HTTP protocol are important part of RESTful interface. SOAP services use HTTP only as transport while RESTful services rely on HTTP as application level protocol.

Most of the status codes are define as part of HTTP/1.1 standard [38]. There are over 70 different HTTP status codes. Most practical API implementation use less than 10. Here is the list of the most commonly used status codes:

Status Code	Reason-Phrase	Status Description
200	OK	The 200 status code indicates that the request has succeeded.
201	Created	The 201 status code indicates that the request has been fulfilled and has resulted in one or more new resources being created.
202	Accepted	The 202 status code indicates that the request has been accepted for processing, but the processing has not been completed.
302	Found	The 302 status code indicates that the target resource resides temporarily under a different URI.

304	Not Modified	The 304 status code indicates that a conditional GET or HEAD request has been received and would have resulted in a 200 (OK) response if it were not for the fact that the condition evaluated to false.
307	Temporary Redirect	The 307 status code indicates that the target resource resides temporarily under a different URI and the user agent MUST NOT change the request method if it performs an automatic redirection to that URI.
308	Permanent Redirect	The 308 status code indicates that the target resource has been assigned a new permanent URI and any future references to this resource ought to use one of the enclosed URIs.
400	Bad Request	The 400 status code indicates that the server cannot or will not process the request due to something that is perceived to be a client error (e.g., malformed request syntax, invalid request message framing, or deceptive request routing).
401	Unauthorized	The 401 status code indicates that the request has not been applied because it lacks valid authentication credentials for the target resource.
403	Forbidden	The 403 status code indicates that the server understood the request but refuses to authorize it.
404	Not Found	The 404 status code indicates that the origin server did not find a current representation for the target resource or is not willing to disclose that one exists.
405	Method Not Allowed	The 405 status code indicates that the method received in the request-line is known by the origin server but not supported by the target resource.
409	Conflict	The 409 status code indicates that the request could not be completed due to a conflict with the current state of the target resource.
500	Internal Server Error	The 500 status code indicates that the server encountered an unexpected condition that prevented it from fulfilling the request.

At the very minimum, any API should support three status codes: 200 indicating that everything is OK, 400 indicating that there was a client error, and 500 indicating that there was a server error. Most APIs support less than 10 different status codes since a large number of status codes makes API more difficult to consume.

The subset of status codes that makes sense for most APIs contains eight status codes (200, 201, 304, 400, 401, 403, 404, and 500). This set can be expanded based on particular requirements of the API.

If an error happened it is also a good idea to return in the payload more information about the problem. For example, the following JSON document could be placed in the body of the HTTP response:

```
{
  "ErrorMessage" : "Plain language error description to help people to understand the problem.",
  "ServiceErrorCode" : 98765,
  "MoreInfo": "http://www.example.com/errors/98765"

}
```

3.3.2.3 Relationship between Resources

Essentially, a RESTful API is comprised of a collection of URIs and HTTP calls to those URIs that take, as parameters and return back, some JSON or XML representations of resources. Many of the resources are conceptually related.

Since resources are a foundation of a RESTful API the relationship between resources may play a significant role in the API design.

There are two different types of relationship between resources. The first is a relationship between two resources that have their own identity. Such resources would map to the concept of an Entity in DDD (Domain-Driven Design) [39]. For example, both a customer order and a customer have their own identities but they also have a relationship. The second type of relationship is when one of the resources does not have its own identity and represents a child resource that can only be identified in the context of a parent resource. For example, an order item can only be identified in the context of an order. This is similar to the concept of a weak entity in a relational database. Defining order items as a separate resource allows for more granular RESTful API, especially if the order contains a lot of items.

The URI for customer number 9832 is:

www.example.org/api/customer/9832

Since there is a relationship between customers and customer orders, the following URI identifies customer orders related to the customer 9832.

www.example.org/api/customer/9832/order

GET www.example.org/api/customer/9832/order should return orders associated with the customer 9832. If customer order number 7799 is one of the orders that were placed by customer 9832 then the URI for that order is www.example.org/api/order/7799. This URI should be used for all operations with that order resource.

Therefore, to get items of that customer order 7799 instead of :

GET www.example.org/api/customer/9832/order/7799/item

the service consumers should use:

GET www.example.org/api/order/7799/item

In other words all manipulations of resources should be done using the direct resource URI (www.example.org/api/order/7799). Hierarchical URIs (www.example.org/api/customer/9832/order) should be only used to provide convenient syntax to get a list of resources in the context of a related resource. Hierarchical resources should go only one level deep.

For customer order items the situation is different since they are defined in the context of a customer order. Therefore, www.example.org/api/order/7799/item is the URI for the list of items in the customer order 7799. Therefore, the following URI

www.example.org/api/order/7799/item/1

is valid, but the similar URI

www.example.org/api/customer/9832/order/7799

should not be used.

Recommendation 3.6 User Shallow Resource Hierarchies

Recommendation	When resource1 and resourse2 are related and have their own identities, id1 and id2 correspondingly, use the following URI
	www.example.org/api/resource1/id1/resource2
	to get the list of resource2 in the context of specific resource1 identified by its id1. For all other manipulation of the resource2 use its URI
	www.example.org/api/resource2/id2
Rationale	The relationship between resources can be quite complex. There is no reason to build deep hierarchies. If a resource identifier is available it can always be accessed directly using its URI.

3.3.2.4 Non-Resource API Capabilities

Sometimes the API has to expose a capability that doesn't deal with the resources directly and is more functional in nature.

Expressing such capabilities in a RESTful way can make them more difficult to understand. A pragmatic approach allows adding such capabilities to the API in a more natural way as actions. For example, the following URI can be used to calculate tax:

www.example.org/api/TaxCalculate/?State=OH&ItemID=21344&Amount=99.99

If the functional capability requires more complex input it can be supplied in the body of the request as JSON or XML.

3.3.2.5 RESTful Service Description

There are several competing approaches to describing RESTful Web APIs. In 2009 WADL (Web Application Description Language) was submitted to W3C but the consortium has no plans to standardize it. WADL was designed as the RESTful equivalent of WSDL but never gained wide acceptance in the industry.

Currently, there are three popular RESTful API description languages on the market: API Blueprint [40], RAML (RESTful API Modeling Language) [41], and Swagger [42]. All three languages have open format to describe REST APIs coupled with tools, like web interface, for visualizing and sharing.

Originally Swagger took a different approach from the others because it did not have clear separation between design and implementation. So, Swagger was more of an API documentation tool where the documentation was hosted alongside the API. Such an approach guaranteed that the API description would always be up-to-date, but it also meant that before an API could be documented at least a skeleton of the API had to be implemented in code. Swagger 2.0 offered a new feature (Swagger Editor) that allows creating APIs using YAML (YAML Ain't Markup Language), which is a human-readable data serialization format [43].

RAML from the very beginning was designed as an API modeling language. That makes it more attractive for enterprises that still require more governance. To describe a RESTful API RAML uses YAML.

The key feature that makes RAML more suitable for the development process of ARTS work teams is its support for XML schemas. RAML interface definition can reference external XSD files. RAML offers support for both XML and JSON schemas and that gives work teams more flexibility and allows the reuse of the existing ARTS standards.

3.4 Service Interface Design Example

The best way to illustrate the concepts and the approach proposed in this chapter is to use a simple example.

This section describes the design of a simplified Gift Registry service that supports just a few capabilities.

3.4.1 Designing Service Information Model

The design of the service starts with the design of the information model. The subject area experts analyze the use cases and create a model of the data structures that are necessary to communicate with the service. The model is created in the context of the capabilities that the service is expected to provide. The capabilities are explored and validated via use cases.

The Gift Registry service will provide the following capabilities.

- Create a new gift registry

- Add an item to the gift registry

- Delete an item from the gift registry

- Update an item in the gift registry

- Get the gift registry data

- Get a certain item from the gift registry

- Get list of items with a certain status from the gift registry

- Close the gift registry

The subject area experts analyze a set of use cases that represent different aspects of the service capabilities and create UML diagram that represents the data to be exchanged with the service.

Service Capability	Required Data	Output Data
Create a new gift registry	Gift registry data. It may or may not contain any items.	None.
Add an item to the gift registry	Identification of the gift registry to be updated and item data.	None.

Delete an item from the gift registry	Identification of the gift registry to be updated and identification of the item within the gift registry.	None.
Update an item in the gift registry	Identification of the gift registry to be updated, identification of the item to be updated and new item data.	None.
Get the gift registry data	Identification of the gift registry to be returned.	Gift registry data.
Get a certain item from the gift registry	Identification of the gift registry that contain the desired item and identification of the item within the gift registry.	Gift registry item data.
Get list of items with a certain status from the gift registry	Identification of the gift registry that contain the items and the status of items to be returned.	List of gift registry items.
Close the gift registry	Identification of the gift registry to be closed.	None.

For the sake of simplicity, this example does not deal with error handling.
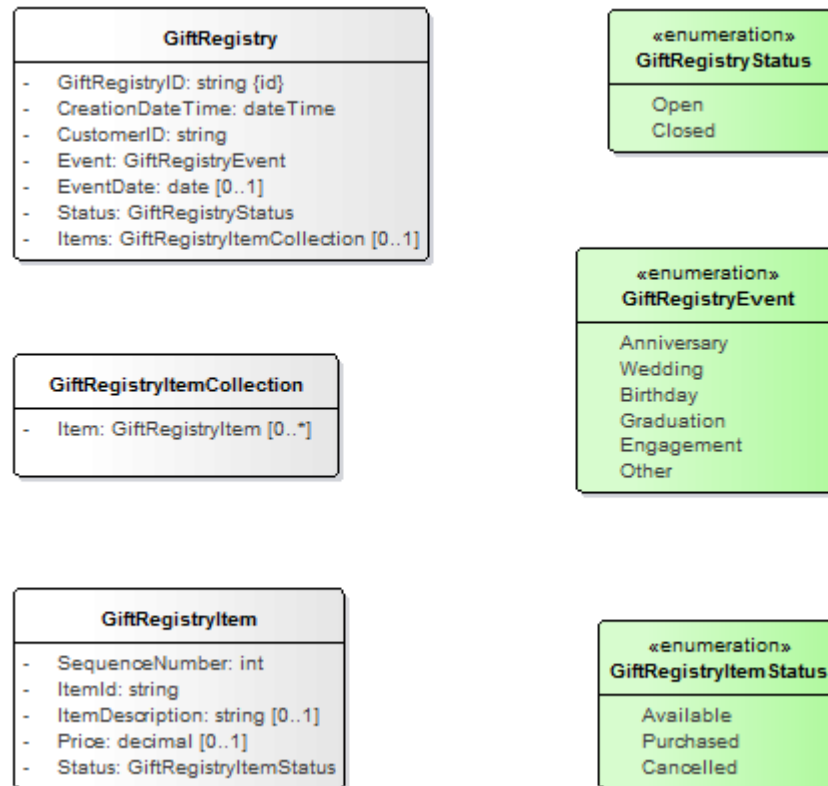
Figure 3.4 UML Diagram of Simple Gift Registry Service Data

The diagram above was created using Enterprise Architect tool [35]. There are three data structures that represent the core of the service information model: GiftRegistry, GiftRegistryItem, and GiftRegistryItemCollection.

3.4.2 Creation of XML and JSON Schemas

Enterprise Architect supports advanced XML schema generation using special stereotyped classes that allow creating UML that can represent the arbitrary XML schema. It also has a special Schema Composer tool that can create XML and JSON schemas from a class model. Generated schemas might need some slight manual adjustments.

The schemas provide foundation for formal description of services.

Generated XML schema:

```xml
<?xml version="1.0" encoding="utf-8"?>
<xs:schema targetNamespace="http://www.nrf-arts.org/IXRetail/namespace/"
  elementFormDefault="qualified"
  xmlns="http://www.nrf-arts.org/IXRetail/namespace/"
  xmlns:xs="http://www.w3.org/2001/XMLSchema"
>
 <xs:complexType name="GiftRegistry">
  <xs:sequence>
   <xs:element name="GiftRegistryID" type="xs:string" minOccurs="1" maxOccurs="1"/>
   <xs:element name="CreationDateTime" type="xs:dateTime" minOccurs="1" maxOccurs="1"/>
   <xs:element name="CustomerID" type="xs:string" minOccurs="1" maxOccurs="1"/>
   <xs:element name="Event" type="GiftRegistryEvent" minOccurs="1" maxOccurs="1"/>
   <xs:element name="EventDate" type="xs:date" minOccurs="0" maxOccurs="1"/>
   <xs:element name="Status" type="GiftRegistryStatus" minOccurs="1" maxOccurs="1"/>
   <xs:element name="Items" type="GiftRegistryItemCollection" minOccurs="0" maxOccurs="1"/>
  </xs:sequence>
 </xs:complexType>
```

```xml
<xs:simpleType name="GiftRegistryEvent">
 <xs:restriction base="xs:string">
  <xs:enumeration value="Anniversary"/>
  <xs:enumeration value="Wedding"/>
  <xs:enumeration value="Birthday"/>
  <xs:enumeration value="Graduation"/>
  <xs:enumeration value="Engagement"/>
  <xs:enumeration value="Other"/>
 </xs:restriction>
</xs:simpleType>
<xs:complexType name="GiftRegistryItem">
 <xs:sequence>
  <xs:element name="SequenceNumber" type="xs:integer" minOccurs="1" maxOccurs="1"/>
  <xs:element name="ItemId" type="xs:string" minOccurs="1" maxOccurs="1"/>
  <xs:element name="ItemDescription" type="xs:string" minOccurs="0" maxOccurs="1"/>
  <xs:element name="Price" type="xs:decimal" minOccurs="0" maxOccurs="1"/>
  <xs:element name="Status" type="GiftRegistryItemStatus" minOccurs="1" maxOccurs="1"/>
 </xs:sequence>
</xs:complexType>
<xs:complexType name="GiftRegistryItemCollection">
 <xs:sequence>
  <xs:element name="Item" type="GiftRegistryItem" minOccurs="0" maxOccurs="unbounded"/>
 </xs:sequence>
</xs:complexType>
<xs:simpleType name="GiftRegistryStatus">
 <xs:restriction base="xs:string">
  <xs:enumeration value="Open"/>
  <xs:enumeration value="Closed"/>
 </xs:restriction>
</xs:simpleType>
<xs:simpleType name="GiftRegistryItemStatus">
 <xs:restriction base="xs:string">
  <xs:enumeration value="Available"/>
  <xs:enumeration value="Purchased"/>
  <xs:enumeration value="Cancelled"/>
 </xs:restriction>
</xs:simpleType>

</xs:schema>
```

Generated JSON schema:

```json
{
 "$schema": "http://json-schema.org/draft-04/schema#",
 "id": "http://www.nrf-arts.org/IXRetail/namespace/#",
 "type": "object",
 "properties":
 {
  "GiftRegistry" :
  {
   "type": "object",
   "properties":
   {
    "GiftRegistryID": {"type": "string"},
    "CreationDateTime": {"type": "string"},
    "CustomerID": {"type": "string"},
    "Event": {"$ref": "#definitions/GiftRegistryEvent"},
    "EventDate": {"type": "string"},
    "Items": {"$ref" : "#definitions/GiftRegistryItemCollection"},
    "Status": {"$ref" : "#definitions/GiftRegistryStatus"}
   },
   "required": ["GiftRegistryID", "CreationDateTime", "CustomerID", "Event", "Status"]
  }
 },
 "definitions":
 {
  "GiftRegistryEvent":
  {
   "type": "string",
   "enum":
   [
    "Anniversary",
    "Birthday",
```

```
          "Engagement",
          "Graduation",
          "Other",
          "Wedding"
        ]
      },
      "GiftRegistryItem":
      {
        "type": "object",
        "properties":
        {
          "SequenceNumber": {"type": "integer"},
          "ItemId": {"type": "string"},
          "ItemDescription": {"type": "string"},
          "Price": {"type": "number"},
          "Status":
            {
              "$ref" : "#definitions/GiftRegistryItemStatus"
            }
        },
        "required": ["SequenceNumber", "ItemId", "Status"]
      },
      "GiftRegistryItemCollection":
      {
        "type": "array",
        "items": {"$ref": "#definitions/GiftRegistryItem"}
      },
      "GiftRegistryItemStatus":
      {
        "type": "string",
        "enum":
        [
          "Available",
          "Cancelled",
          "Purchased"
        ]
      },
      "GiftRegistryStatus":
      {
        "type": "string",
        "enum":
        [
          "Closed",
          "Open"
        ]
      }
    }
}
```

These schemas generated from the UML diagram represent formal definition of the data structures that are used to exchange the information between the registry service and its consumers. They provide the foundation for the formal definition of the service interface.

3.4.3 Defining Service Capabilities

The next step in defining the service interface is the formal description of the service capabilities. SOAP and RESTful services use different approaches for the formal service description.

SOAP-based service interfaces are formally described using WSDL. There are currently several techniques that can be used to describe the RESTful APIs. ARTS work teams can, for example, use RAML since it supports both JSON and XML payloads.

3.4.3.1 Defining SOAP Service Interface Using WSDL

When designing a SOAP service interface, service capabilities are expressed in terms of operations.

Best Practices for Services Implementation Using ARTS Standards

Service Capability	Service Operation	Input	Output
Create a new gift registry	GiftRegistryCreate	GiftRegistry Data	None.
Add an item to the gift registry	GiftRegistryItemAdd	GiftRegistryItem Data	None.
Delete an item from the gift registry	GiftRegistryItemDelete	Integer Item Sequence Number	None.
Update an item in the gift registry	GiftRegistryItemUpdate	GiftRegistryItem Data	None.
Get the gift registry data	GiftRegistryGet	String Gift Registry ID	GiftRegistry Data
Get a certain item from the gift registry	GiftRegistryItemGet	String Gift Registry ID and Integer Item Sequence Number	GiftRegistryItem Data
Get list of items with a certain status from the gift registry	GiftRegistryItemsGet	String Gift Registry ID and String Item Status	Collection of GiftRegistryItem Data
Close the gift registry	GiftRegistryClose	String Gift Registry ID	None.

A WSDL document consists of a set of definitions that describe a SOAP service.

```xml
<?xml version="1.0" encoding="utf-8"?>
<!-- WSDL definition structure -->
<definitions
  name="ServiceName"
  targetNamespace="http://example.org/ServiceName/"
  xmlns="http://schemas.xmlsoap.org/wsdl/"
>
  <!-- abstract definitions -->
  <types> ...
  <message> ...
  <portType> ...

  <!-- concrete definitions -->
  <binding> ...
  <service> ...

</definition>
```

To create a WSDL-based service definition, ARTS work teams only have to deal with the first three elements (types, message, and portType) that constitute the programmatic service interface. The last two elements (binding and service) describe the concrete implementation details such as service address, communication protocol, etc.

The definition of portType inside a WSDL document is conceptually close to the definition of an interface within a programming language. It includes a description of all operations and input

and output messages for each operation. For the Gift Registry service the portType element would look like:

```xml
<?xml version="1.0" encoding="utf-8"?>
<wsdl:definitions name="GiftRegistryService"
 targetNamespace=http://www.nrf-arts.org/IXRetail/namespace/GiftRegistry/
 xmlns:tns="http://www.nrf-arts.org/IXRetail/namespace/GiftRegistry/"
 xmlns:arts="http://www.nrf-arts.org/IXRetail/namespace/"
 xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/"
 xmlns:xsd="http://www.w3.org/2001/XMLSchema"
>
  <!-- abstract definitions -->
  <types> ...
  <message> ...

  <!-- GRI (Gift Registry Interface) -->
  <wsdl:portType name="GRI">
   <wsdl:operation name="GiftRegistryCreate">
    <wsdl:input message="tns:GRI_GiftRegistryCreate_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryCreate_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryItemAdd">
    <wsdl:input message="tns:GRI_GiftRegistryItemAdd_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryItemAdd_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryItemDelete">
    <wsdl:input message="tns:GRI_GiftRegistryItemDelete_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryItemDelete_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryItemUpdate">
    <wsdl:input message="tns:GRI_GiftRegistryItemUpdate_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryItemUpdate_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryGet">
    <wsdl:input message="tns:GRI_GiftRegistryGet_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryGet_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryItemGet">
    <wsdl:input message="tns:GRI_GiftRegistryItemGet_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryItemGet_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryItemsGet">
    <wsdl:input message="tns:GRI_GiftRegistryItemsGet_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryItemsGet_OutputMessage"/>
   </wsdl:operation>
   <wsdl:operation name="GiftRegistryClose">
    <wsdl:input message="tns:GRI_GiftRegistryClose_InputMessage"/>
    <wsdl:output message="tns:GRI_GiftRegistryClose_OutputMessage"/>
   </wsdl:operation>
  </wsdl:portType>

  <!-- concrete definitions -->
  <binding> ...
  <service> ...
</definition>
```

Basically, portType element contains all operations and names of input and output messages. Structure of the messages is defined inside WSDL message elements.

```xml
<wsdl:message name="GRI_GiftRegistryCreate_InputMessage">
 <wsdl:part name="parameters" element="tns:GiftRegistryCreate"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryCreate_OutputMessage">
 <wsdl:part name="parameters" element="tns:GiftRegistryCreateResponse"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemAdd_InputMessage">
 <wsdl:part name="parameters" element="tns:GiftRegistryItemAdd"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemAdd_OutputMessage">
 <wsdl:part name="parameters" element="tns:GiftRegistryItemAddResponse"/>
</wsdl:message>
```

```xml
<wsdl:message name="GRI_GiftRegistryItemDelete_InputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemDelete"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemDelete_OutputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemDeleteResponse"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemUpdate_InputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemUpdate"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemUpdate_OutputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemUpdateResponse"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryGet_InputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryGet"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryGet_OutputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryGetResponse"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemGet_InputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemGet"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemGet_OutputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemGetResponse"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemsGet_InputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemsGet"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryItemsGet_OutputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryItemsGetResponse"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryClose_InputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryClose"/>
</wsdl:message>
<wsdl:message name="GRI_GiftRegistryClose_OutputMessage">
  <wsdl:part name="parameters" element="tns:GiftRegistryCloseResponse"/>

</wsdl:message>
```

Every message element contains the element attribute that is defined under WSDL types. The WSDL types element encloses data type definitions typically presented as XML schema. The types element can contain one or more schema elements.

For the Gift Registry service example, types element would contain two XML schemas. The first XML schema, which was created from the UML model, defines data structures as a set of complex types (see 3.4.2). The second XML schema defines elements that represent input and output messages used by the Gift Registry service operations.

```xml
<xs:schema elementFormDefault="qualified"
 targetNamespace=http://www.nrf-arts.org/IXRetail/namespace/GiftRegistry/
 xmlns:xs=http://www.w3.org/2001/XMLSchema
 xmlns:arts="http://www.nrf-arts.org/IXRetail/namespace/"
>

  <xs:element name="GiftRegistryCreate">
   <xs:complexType>
    <xs:sequence>
     <xs:element minOccurs="0" name="giftRegistry" nillable="true" type="arts:GiftRegistry"/>
    </xs:sequence>
   </xs:complexType>
  </xs:element>
  <xs:element name="GiftRegistryCreateResponse">
   <xs:complexType>
    <xs:sequence/>
   </xs:complexType>
  </xs:element>
  <xs:element name="GiftRegistryItemAdd">
   <xs:complexType>
    <xs:sequence>
     <xs:element minOccurs="0" name="item" nillable="true" type="arts:GiftRegistryItem"/>
    </xs:sequence>
```

```xml
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemAddResponse">
      <xs:complexType>
        <xs:sequence/>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemDelete">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="itemSequenceNumber" type="xs:int"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemDeleteResponse">
      <xs:complexType>
        <xs:sequence/>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemUpdate">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="item" nillable="true" type="arts:GiftRegistryItem"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemUpdateResponse">
      <xs:complexType>
        <xs:sequence/>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryGet">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="giftRegistryId" nillable="true" type="xs:string"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryGetResponse">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="GiftRegistryGetResult" nillable="true" type="arts:GiftRegistry"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemGet">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="giftRegistryId" nillable="true" type="xs:string"/>
          <xs:element minOccurs="0" name="itemSequenceNumber" type="xs:int"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemGetResponse">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="GiftRegistryItemGetResult" nillable="true"
            type="arts:GiftRegistryItem"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemsGet">
      <xs:complexType>
        <xs:sequence>
          <xs:element minOccurs="0" name="giftRegistryId" nillable="true" type="xs:string"/>
          <xs:element minOccurs="0" name="status" nillable="true" type="xs:string"/>
        </xs:sequence>
      </xs:complexType>
    </xs:element>
    <xs:element name="GiftRegistryItemsGetResponse">
      <xs:complexType>
        <xs:sequence>
```

```
    <xs:element minOccurs="0" name="GiftRegistryItemsGetResult" nillable="true"
     type="arts:GiftRegistryItemCollection"/>
  </xs:sequence>
 </xs:complexType>
</xs:element>
<xs:element name="GiftRegistryClose">
 <xs:complexType>
  <xs:sequence>
    <xs:element minOccurs="0" name="giftRegistryId" nillable="true" type="xs:string"/>
  </xs:sequence>
 </xs:complexType>
</xs:element>
<xs:element name="GiftRegistryCloseResponse">
 <xs:complexType>
  <xs:sequence/>
 </xs:complexType>
</xs:element>
</xs:schema>
```

The XML schema above completely describes the Gift Registry Service SOAP service interface and references XML schema that was generated from the UML diagram.

3.4.3.2 Defining REST API Using RAML

When designing a RESTful API, service capabilities are expressed in terms of resources and standard HTTP verbs.

Service Capability	HTTP Verb	Resource	Input	Output
Create a new gift registry	POST	/giftregistry	GiftRegistry Data	
Add an item to the gift registry	POST	/giftregistry/{ID}/item	GiftRegistryItem Data	
Delete an item from the gift registry	DELETE	/giftregistry/{ID}	Integer Item Sequence Number	
Update an item in the gift registry	PUT	/giftregistry/{ID}/item	GiftRegistryItem Data	
Get the gift registry data	GET	/giftregistry/{ID}	String Gift Registry ID	GiftRegistry Data
Get a certain item from the gift registry	GET	/giftregistry/{ID}/item/ {sequanceNumber}	String Gift Registry ID and Integer Item Sequence Number	GiftRegistryItem Data

Get list of items with a certain status from the gift registry	GET	/giftregistry/{ID}/item/ status={statusValue}	String Gift Registry ID and String Item Status	Collection of GiftRegistryItem Data
Close the gift registry	DELETE	/giftregistry/{ID}	String Gift Registry ID	

RAML is a vendor-neutral open specification [44] for description of RESTful APIs.

A RAML API description can be created using any text editor. However, MuleSoft has developed a free browser-based editor for RAML that significantly simplifies the authoring of RAML documents [45].
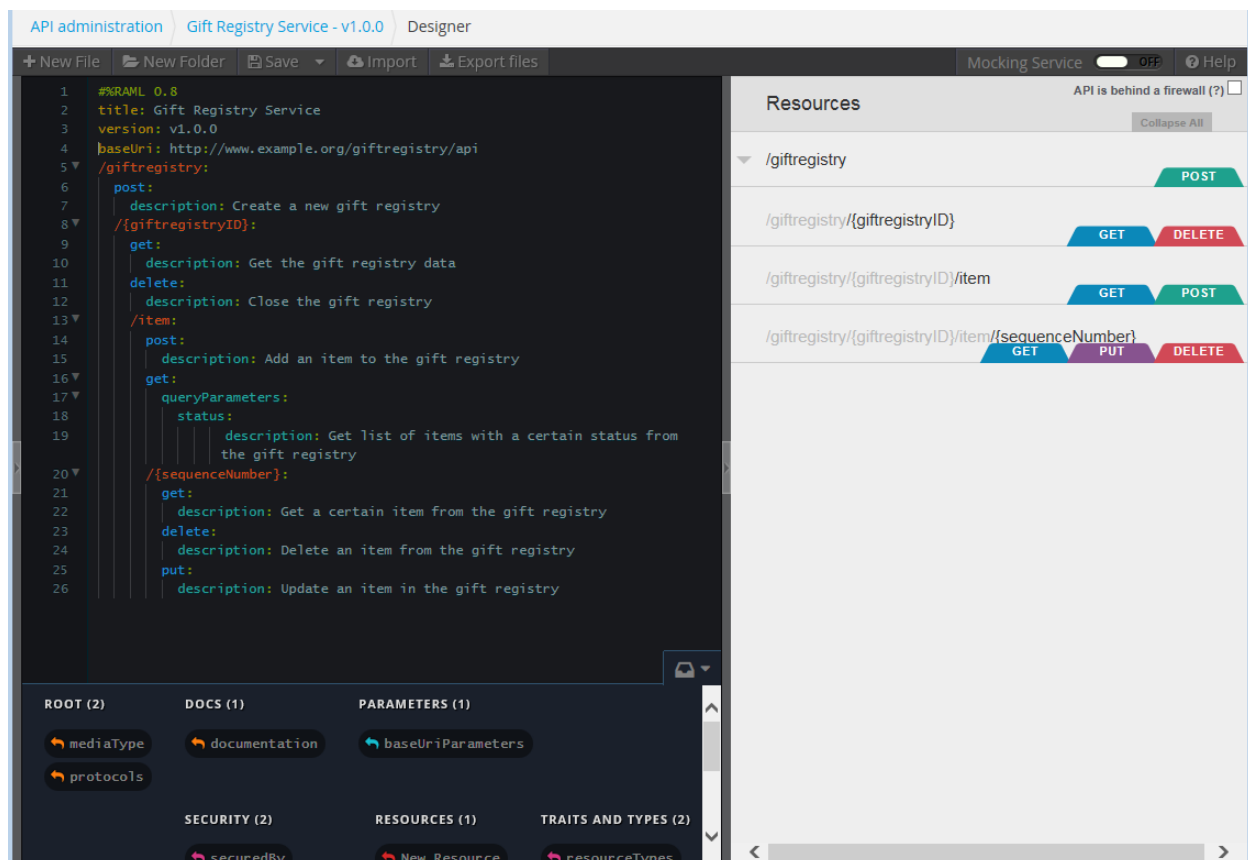
Figure 3.5 REST API Designer

The RAML document below specifies how resources and HTTP verbs are mapped to service capabilities.

#%RAML 0.8

title: **Gift Registry Service**

version: **v1.0.0**

baseUri: **http://www.example.org/giftregistry/api**

/giftregistry:

  post:

```
        description: Create a new gift registry
   /{giftregistryID}:
     get:
       description: Get the gift registry data
     delete:
       description: Close the gift registry
     /item:
       post:
         description: Add an item to the gift registry
       get:
         queryParameters:
           status:
             description: Get list of items with a certain status from the gift registry
       /{sequenceNumber}:
         get:
           description: Get a certain item from the gift registry
         delete:
           description: Delete an item from the gift registry
         put:
           description: Update an item in the gift registry
```

The RAML document above defines the structure of URIs to access the Gift Registry REST API. The structure of a request and/or response body has to be further specified by the schema property under the appropriate media type. XML and JSON schemas can be declared inline or in an external file.

For large API descriptions it is preferable to place schemas into files since it makes the RAML documents more readable and simplifies the reuse of the same data structures. For example, the JSON schema generated from the UML can be split into three separate files: giftregistry.json, giftregistryitem.json, and giftregistryitemlist.json.

After every resource is placed into a separate JSON schema file they can be easily referenced from the RAML document to complete the API definition.

```
#%RAML 0.8
title: Gift Registry Service
version: v1.0.0
baseUri: http://www.example.org/giftregistry/api
/giftregistry:
  post:
    description: Create a new gift registry
    body:
      application/json:
        schema: !include giftregistry.json
  /{giftregistryID}:
    get:
      description: Get the gift registry data
```

```
      responses:
        200:
          body:
            application/json:
              schema: !include giftregistry.json
  delete:
    description: Close the gift registry
  /item:
    post:
      description: Add an item to the gift registry
      body:
        application/json:
          schema: !include giftregistryitem.json
    get:
      queryParameters:
        status:
          type: string
          description: Get list of items with a certain status from the gift registry
      responses:
        200:
          body:
            application/json:
              schema: !include giftregistryitemlist.json
    /{sequenceNumber}:
      get:
        description: Get a certain item from the gift registry
        responses:
          200:
            body:
              application/json:
                schema: !include giftregistryitem.json
      delete:
        description: Delete an item from the gift registry
      put:
        description: Update an item in the gift registry
        body:
          application/json:
            schema: !include giftregistryitem.json
```

Since RAML uses the quite expressive YAML serialization language [43] that was designed to be human-friendly the RESTful API definition is fairly easy to understand.

Another useful feature of RAML is its ability to provide examples as part of the interface description. The examples can use both JSON and XML.

```
body:
```

```
application/json:
  schema: !include giftregistryitemlist.json
  example: |
    "GiftRegestryItems": [
      {
        "ItemId": "550e8400-e29b-41d4-a716-446655440000",
        "SequenceNumber": 2,
        "ItemDescription": "Mens Denim Jeans 36X32",
        "Price": 49.99,
        "Status": "Available"
      },
      {
        "ItemId": "630a6400-229a-41d4-a716-749847710000",
        "SequenceNumber": 4,
        "ItemDescription": "Jogger Sweatpants",
        "Price": 29.99,
        "Status": "Available"
      }
    ]
```

The ability to provide descriptions and examples as part of RAML API definition can be very useful.

3.4.3.3 Example Conclusions

Both the SOAP and RESTful Gift Registry service APIs are based on the same information model that was created using abstract UML. This approach provides semantic consistency between different styles of service interfaces and data representation formats.

4. SERVICE IMPLEMENTATION

This chapter discusses some important aspects of implementing services. It also presents several essential design patterns that can improve the quality of service implementation especially in cloud environment.

4.1 Granularity Considerations

Granularity is an important consideration in designing a service. Selecting the right level of granularity is a balancing act that depends on many factors. It is impossible to provide an answer that would fit all the situations; however, a few guidelines can be very helpful.

The following aspects of granularity are essential to consider when designing a service: service granularity, service capability or operations granularity, and data granularity.

Service granularity is defined by the functional context of the entire service. Services that are intended to cover a larger functional context are considered to have a coarse granularity. On the contrary, fine-grained services expose a narrow specialized functionality.

Capability granularity is defined by a functional scope of a single service capability or operation. It shows how much work is performed by the capability. A service can contain both fine- and coarse-grained capabilities.

Data granularity represents the amount of data exchanged by a service capability during a single invocation. It is usually somewhat related to the capability granularity, because coarse-grained capabilities tend to exchange coarse-grained data. However, it is possible for a fine-grained capability to retrieve a large chunk of data or vice versa.

4.1.1 Service Granularity

A service can be viewed as a package of capabilities related to a particular functional context. It is a unit of design, development, testing, deployment, and maintenance.

A coarse-grained design might complicate a service's maintenance - a change to any part of a large service contract will require a new version, which could impact service consumers who might not have been affected if the service was more granular. In addition, finer-grained services offer more flexible deployment options. It may be difficult to justify the deployment and the potential overhead of a large, expensive component just to be able to utilize a small subset of its functionality. Locating a single desired capability within a coarse-grained service that covers a large functional context might not be an easy task either. This may have a negative impact on service reusability and may result in creation of redundant capabilities.

On the other hand, deploying and managing a large number of fine-grained services can be a daunting job. It might also have performance implications, since crossing service boundaries creates additional overhead. In addition, certain issues, such as transactions and security, are notoriously more difficult to coordinate across multiple services than within a single service.

	Fine Granularity	Coarse Granularity
Pros	Flexible deployment, increased reusability, superior ability to predict and maintain service-level agreements. Ability to scale different	Better performance, smaller number of services.

	parts of the system independently. Ability to divide the development and maintenance in more granular and easy to manage pieces.	
Cons	Performance risks, complex maintenance of a large number of services.	Complex maintenance of a large service, reduced reusability, higher chance of partial operational redundancy. Coarse services are much more difficult to scale.

Looking from the perspective of service classification (i.e., utility, entity, task, process):

- Granularity of utility services is usually defined by grouping infrastructural capabilities with a common purpose—for example, logging service.

- The functional context of entity services is scoped by the entities that they manage. For example, granularity of the customer service is defined by the customer entity. In this case customer service would include all the capabilities that are necessary to maintain customer entity.

- Task service typically contains a group of capabilities related to the same business task—for example, a tax calculation service. Task services tend to be fairly granular.

- Process services, on the other hand, usually deal with a larger functional context defined by the encapsulated business process—for example, customer order-processing service would include all the capabilities necessary to manage a customer order.

4.1.2 Capability Granularity

Capability or operations granularity deals with the amount of logic that should be performed by a service capability during a single invocation.

A major consideration in defining the right level of capability or operations granularity is performance. Splitting a single capability into a set of finer-grained capabilities could result in a chattier interface, which might negatively impact performance. However, if a fine-grained capability represents a distinct, reusable, useful piece of business logic, such decomposition might be very useful. Sometimes it makes sense to expose both a coarse-grained capability and a set of corresponding finer-grained capabilities, although doing so does create some redundancy.

A second important consideration is that every capability should perform a complete unit of work. This important requirement helps avoid maintaining a transient state between invocations of capabilities. On the other hand, it is important to avoid creating unnecessary dependency between autonomous pieces of business logic simply to avoid state considerations. Ideally, service capabilities should represent well-defined, self-contained business actions. This business suitability criterion is an important consideration in defining the right level of the granularity for capabilities.

	Fine Granularity	**Coarse Granularity**
Pros	Increased reusability.	Less chatty service interactions.
Cons	Performance risks, complex maintenance of large numbers of service capabilities.	Reduced reusability. Only small subset of business logic is exposed via service interface.

Looking at services classification,

- Capabilities inside utility services should have a high level of reusability. As a result, utility services tend to have somewhat more granular capabilities than do business services.

- Entity services typically have entity-level CRUD capabilities, along with some data validation. Because entity services often perform data modifications within a data store, it is important that every capability completes a single unit of work that does not leave any data in an inconsistent state.

- Capability granularity for task services is usually selected according to the business suitability criterion. A well-defined business task provides a good outline of the capability's functional scope. Unit of work considerations should also be taken into account for task services.

- Business suitability criterion is also extremely important for defining capabilities for process services. The functional scope of business process capabilities is defined in such a way that these capabilities would facilitate the transition of the business process from one consistent state to another.

4.1.3 Data Granularity

Data granularity defines the amount of data exchanged during a single invocation of a service capability. As in the case of capability granularity, performance is a crucial consideration in determining the chunkiness of the data. When capabilities are defined and scoped, the data granularity becomes implicitly specified to support the capabilities. For this reason, data granularity should be an essential consideration when identifying the granularity of capabilities. There has been a tendency for services to exchange fairly large, document-style messages. This is in contrast to more traditional finer-grained remote procedure call RPC–style communications. Passing data in smaller chunks requires more round trips, which might negatively impact the performance of the service.

	Fine Granularity	**Coarse Granularity**
Pros	More flexibility.	Better performance due to reduced number of round trips.
Cons	Performance risks due to potential chattiness.	Reduced flexibility.

- Because capabilities of utility services tend to be more granular and often follow the RPC approach, utility services tend to exchange more granular data.

- Data granularity for entity services is mostly defined by the amount of data the entity contains. Entity services that manage complex entities usually have coarser data granularity. If a collection of coarse-grained entities must be returned to a service consumer, the amount of data can become too large. One technique for increasing the granularity (i.e., making it more fine grained) of entity services that encapsulate large hierarchical entities is called "lazy loading." In this approach detailed information is returned only for the root-level element, while child collections are represented by a narrow subset of mostly reference data. A service consumer can request more detailed information about any member of a child collection using an additional call. Obviously, this approach potentially requires more round trips, but sometimes it is the only practical option available.

- Task services should have data granularity that is entirely dictated by the specific needs of the business operation, avoiding the tendency to send more data than required for the business task at hand. Limiting data to what is needed avoids unintended data coupling, which can sometimes result from the service implementation using data simply because they are available, even though the service capability indicated no explicit need for it.

- Process services tend to have a fairly coarse data granularity, because a business process service is often used to build complex data structures. These structures have to be repeatedly communicated to a service consumer to represent the state of the process, which results in coarse data exchange.

It is important to note that in recent years there has been a trend to use finer service granularity. The design approach based on the concept of microservices became very popular among distributed system architects.

4.2 Microservices

Microservices definitely go beyond just simple granularity considerations. This architectural approach has been a subject of many discussions in the recent years. The idea is to architect a complex system as a set of highly-cohesive granular services that can evolve independently over time. Even though the article by James Lewis and Martin Fowler [15] does not give a precise definition of microservices they describe their common characteristics.

- *Componentization via Services.* Services are used as independently replaceable and upgradable units of software.

- *Organized around Business Capabilities.* Services implement a well-defined business capability.

- *Products not Projects.* Development teams own Products (microservices) they created rather than just participate in their development Projects and then move on.

- *Smart endpoints and dumb pipes.* Avoid using complex middleware technologies like ESB.

- *Decentralized Governance.* There is no single standard technology or platform to be used by all team.

- *Decentralized Data Management.* Every service can have its own fit-for-purpose data store.

- *Infrastructure Automation.* Continuous delivery becomes an essential part of how the software is deployed.

- *Design for failure.* Instead of trying to build system that will not fail assume that failure will happen and build the system that can recover from failure.

- *Evolutionary Design.* Decomposition of the software system into more granular services enables evolutionary approach since services can evolve more independently.

Microservices is an approach of decomposing a large software system into smaller, more manageable pieces. Although decomposition of a large system into sub-systems has been around for a long time, the microservices approach takes into account some of the best practices that have been developed by the leaders of the industry designing complex SOA software. Also the distributed deployment model that is enabled by this approach can potentially deliver much better scalability. Since microservices can be scaled independently the system can be implemented to adapt to different patterns of workload resulting in much greater flexibility and efficiency.

It is important to note that while every single microservice is much less complex the complexity of the overall system increases with number of microservices that need to be deployed and maintained. The microservices approach is about managing the complexity of a fairly large system through decomposition, which introduces a new set of issues associated with a distributed system (discovery, remote calls failures, data synchronization, and etc.). To justify such approach the software system has to have enough inherent complexity that can be addressed by decomposition. Of course the complexity of the deploying and maintaining large distributed system can be alleviated using containers and technologies offered by underlying modern platforms like load balancing, auto-scaling, etc.

Microservices can be thought of as a fine-grained style of SOA, where every service is centered on a single capability of the business domain and can be built and deployed independently of other services. Well known cloud technology expert Adrian Cockcroft describes microservices as "loosely coupled service oriented architecture with bounded context". The concept of bounded context was introduced by Eric Evans in his famous "blue" book on Domain-Driven Design (DDD) [39]. It describes a consistent subset of a business domain suitable for independent development. Therefore a microservice within a bounded context is self-contained for the purpose of software development.

The key principles behind microservices - focus on a specific business capability, well-defined contracts, loose coupling, and well-designed and stable APIs are the same as the original principles behind SOA. Microservices also added important DevOps and continuous delivery considerations that imply important organizational and cultural changes. Microservices approach promotes a way of designing complex software systems that enforces the service-orientation. No longer can slapping a SOAP interface on top of monolithic system be seen as a service-oriented solution.

The ability to evolve services independently is a key to the evolution of very complex distributed system especially in the cloud environment. It requires careful management of dependencies and strict backward compatibility policies.

With the advent of the cloud microservices approach is gaining more and more popularity. It requires some important organizational changes that remove friction from the development process. Since small software development teams operate in high trust, low process environment they can deliver new features to the marketplace very fast. This approach promotes the culture of "freedom and responsibility" [46]. Using small development teams to implement microservices approach is consistent with Conway's Law that states that the structure of a system reflects the communication structures of the organization that designed the system [47]. There is a well-known concept of "two-pizza team" introduced by Amazon's founder Jeff Bezos. The two-pizza team rule states that if two pizzas are not sufficient to feed a team, then the team is too big in size.

It is important to note that microservices approach is not suitable for every project. It introduces certain degree of complexity due to its distributed nature. This inherent complexity is difficult to justify for smaller systems. However, as the size of a system grows and its complexity increases, the advantages of the microservices approach become more obvious.

Recommendation 4.1 Use Microservices to Design, Develop and Deploy Complex Systems

Recommendation	Build large and complex software system as a set of highly cohesive and loosely coupled services that can be updated and even replaced in autonomous manner independently from the rest of the system.
Rationale	Microservices approach enables design of complex software systems using simple granular services that can be evolved independently. Every microservice represents a specific granular autonomous piece of business logic. Since every microservice can evolve independently this approach gives software development teams unprecedented level of control and flexibility and enables them to deliver innovative solutions at rapid rate.

Since proper management of microservices requires certain infrastructure and automation to be in place they are only useful for fairly large systems. Some industry experts argue that starting with monolith first and then evolving system to microservices is much more pragmatic approach that historically has been more successful [48]. But even monolith application should be carefully designed with modularity in mind paying special attentions to boundaries and interfaces. In fact, monolith is a great way to refine them since refactoring of monolithic application is so much easier than changing service APIs or moving pieces of logic from one service to another. With proper design it will take much less effort to refactor an internal component to be an independent service. Microservices can be thought of as the next stage in a lifecycle of a complex software system.

Microservices architecture promotes a more agile approach to software development, which is crucial for modern retail enterprise. It enables companies to react more quickly to rapidly changing technology landscape and consumer behavior. Also retailers and software vendors can test new innovative approaches and, if necessary, adjust the direction without large upfront investments.

4.3 Service Versioning

Versioning is one of the most critical considerations when implementing a service and publishing its API. There are two major types of versions: version of the service API and versions of the internal components that are part of service implementation.

The primary focus of this paper is the service interface versioning. Versions of implementation components are very important for service maintenance operations but they should have no impact on service consumers and can be considered internal implementation details.

Recommendation 4.2 Version Every Release of Service API

Recommendation	Always version every release of the service API.
Rationale	Service API versioning is essential to determine compatibility issues. It is even more critical for public API. Once API is published consumers will start depending on it. To be able to evolve the service and implement new and/or improved capabilities changes to the API might be necessary. Unfortunately, not all the changes can be implemented in backwards compatible manner. Therefore it might be necessary to run the new and the old version of the API concurrently until all the service consumers migrate to the new version. API versioning provides mechanism to clearly distinguish between different API versions.

4.3.1 Versioning Scheme

According to the versioning guidelines in ARTS Operations Guide [49], ARTS uses three-sequence version identifier (major.minor.fix). The first sequence (major) must be incremented if the new release breaks backwards compatibility. The second sequence (minor) is incremented if the new release contains some additions to the current functionality but they do not break backward compatibility. Finally, the third sequence (fix) is incremented when corrections are made to fix identified issues and errors.

Major	**Minor**	**Fix**
Breaks backward compatibility. Requires a new Charter. Requires modification to all segments of the documentation.	Adds a new use case or device or subject area. Requires modification to conformance.	Staff Fixes Problem. No Expanded Scope or intent. No impact to Conformance. Corrections due to oversight.

ARTS versioning scheme is very similar to major principals of Semantic Versioning [50] that can be summarized as follows:

Given a version number MAJOR.MINOR.PATCH, increment the:

1. MAJOR version when you make incompatible API changes,
2. MINOR version when you add functionality in a backwards-compatible manner, and
3. PATCH version when you make backwards-compatible bug fixes.

Additional labels for pre-release and build metadata are available as extensions to the MAJOR.MINOR.PATCH format.

Thus Semantic Versioning approach can be considered as an extension of ARTS versioning.

It is important to note that backwards compatibility depends on the technology that is used to implement service consumers. For example, certain changes to the resource structure can lead to deserialization failure on the client side. On the other hand if the structure from the wire is consumed as raw XML or JSON no automatic mapping to framework objects is necessary. Technologies like XPath can be used to retrieve the value of a certain XML element or attribute even if there were quite significant changes to the structure of the document. Therefore, different client implementations have different levels of robustness.

4.3.2 Versioning Techniques

Since according to the Recommendation 3.3 defining standard data structures for nouns (business entities, business concepts, etc.) is the foundation of the service API definition therefore versioning of these data structures represents the basis of API versioning.

4.3.2.1 Schema Versioning

The approach taken by ARTS workgroup to version XML schema documents was to embed special "fixed" version attributes inside the schema.

For example, ARTS XML POSLog V6 schema has the following version attributes defined as part the complex type for the root POSLog element:

```
<xs:attribute fixed="6" form="unqualified" name="MajorVersion" use="required"/>
<xs:attribute fixed="0" form="unqualified" name="MinorVersion"/>

<xs:attribute fixed="0" form="unqualified" name="FixVersion"/>
```

It is interesting that version attributes are defined on the TransactionBase complex type, which is not the type of the root element. The idea is that Batch element can contain transactions with different versions, which adds additional flexibility.

It is interesting that POSLog XML schema also defines version attributes on the TransactionBase complex type, which is not the type of the root element. The idea is that the Batch element can contain multiple transactions with different versions, which adds additional flexibility:

```
<xs:attribute fixed="6" form="unqualified" name="MajorVersion" type="xs:integer">
  <xs:annotation>
    <xs:documentation>POSLog allows many different transactions to be merged into one
      Batch. This is the version for this transaction. If it is the same as the one
      in POSLog it can be left out.</xs:documentation>
  </xs:annotation>
</xs:attribute>
<xs:attribute fixed="0" form="unqualified" name="MinorVersion" type="xs:integer"
  use="optional"/>
<xs:attribute fixed="0" form="unqualified" name="FixVersion" type="xs:integer"

  use="optional"/>
```

Another approach to implement schema versioning would be to define special complex type Version that could be embedded inside different complex types that require versioning.

Similar approach can be used to define versions inside JSON schemas. Even though JSON schema does not have "fixed" attribute semantics, an enumeration with a single value could be used to specify fixed version numbers.

```
{
  "$schema": "http://json-schema.org/draft-04/schema#",
  "definitions":
  {
    "Version":
```

```
{
  "title": "Version Object",
  "type": "object",
  "properties":
  {
    "MajorVersion": {"type": "number", "enum":[6]},
    "MajorVersion": {"type": "number", "enum":[0]},
    "FixVersion": {"type": "number", "enum":[0]}
  },
  },
  "required": ["MajorVersion"]
}

}
```

Every set of XML or JSON schemas that provide foundation for service API should be versioned.

Recommendation 4.3 Version Every Set of XML or JSON schemas

Recommendation	Set version attributes for at least every root element in the set of XML or JSON schemas that provide foundation for the service interface.
Rationale	Versioning of the root elements helps validation and also allows using appropriate business logic for every version of the element.

One of the potential drawbacks of the schema versioning approach is that application has to pre-parse document to know what schema version it uses. Some frameworks deserialize payload into programming language constructs like objects but to use the correct object type for deserialization it might be necessary to know the version number that is buried inside the payload. It does not mean that schema versioning is not necessary it just means that only schema versioning might not be enough.

There are two major versioning approaches that do not require deserialization of the payload. Version numbers can be embedded in either inside URIs or inside headers.

4.3.2.2 URI Versioning

If URI based versioning is used then a client application depending on the API version it supports would have to use different URIs to access service capabilities. There several ways to embed the version number into URI.

Hostname	v6.rti.example.org
Path Prefix	example.org/rti/V6/.../
Path Suffix	example.org/rti/transaction/2015-01-17-1001-Reg1-1011/V6
Query Parameter	example.org/rti/transaction/2015-01-17-1001-Reg1-1011?ver=6

Since only the major version defines backwards compatibility it is not necessary to include minor versions inside the URI. In fact, such approach could potentially result in unnecessary breaking changes.

URI design is an important aspect of RESTful API and as such represents commitment to the service consumers. It is possible to evolve an API adding new resources and/or new data

elements to existing resources. Such backwards compatible modifications should not result in a change of URI.

From a REST purist point of view the resource URI should not change just because the representation of a resource changed breaking RESTful API consumers. Such a purist approach definitely imposes some limitations on how the service can evolve over time.

A more pragmatic approach is to embed the major version number as part of the URI. Such an approach allows side by side deployment of multiple service versions and guarantees behavioral consistency to service consumers. When the interface to a service changes in a non-backwards-compatible way, it can be viewed as an entirely new service.

Recommendation 4.4 Embed Major Version Number into Service URI as Path Prefix

Recommendation	If using URI based versioning the version number should be embedded into service URI as a path prefix: example.org/rti/V6/.../
Rationale	Embedding major version number into the service URI allows dealing with non-backwards-compatible changes and provides guarantees to service consumers that the service will behave consistently. Using path prefix is the most common and explicit approach.

It is also possible to implement a mixed approach, in which the URI without a version number will be mapped to the latest implementation. For example, if the latest version of RTI interface were version 6 then the following URIs would return the same implementation of the transaction resource:

example.org/rti/transactions/2015-01-17-1001-Reg1-1011

example.org/rti/V6/transactions/2015-01-17-1001-Reg1-1011

With this approach the implementers of the RTI service client have a choice to either always develop and run against the latest version of the API and to use the same URI or to use a versioned URI that would have to be changed when moving to the latest implementation.

4.3.2.3 Header Versioning

It is possible to use extensibility of the underlying protocol to pass the version information. For example, for SOAP-based services it is fairly easy to define a custom SOAP header that would contain version information. For RESTful APIs, custom HTTP header could be a good option.

It is also possible to version representation of resources using HTTP content negotiation. In this case the desired version of the resource is specified by client using HTTP Accept header. For example

GET /transactions/ HTTP/1.1

Host: www.example.org

Accept: application/POSLog.V6+xml

Putting a version number in the header can be optional and if omitted the service would assume the latest version.

Typically API version will be put into x-API_name-Version header.

GET /transactions/ HTTP/1.1

Host: www.example.org

x-RTI-Version: 6.0

The major drawback of the header-based versioning is that it might be not as intuitive as URI-based versioning and it puts additional burden on service consumers the have to deal with the additional complexity of headers.

4.3.2.4 Choosing between URI and Header Versioning

Even though both approaches can be implemented with reasonable defaults, the URI-based versioning is more explicit. Therefore if the intent is to clearly communicate version information to the service consumers then URI-based versioning would be a preferred approach. It is also a simpler option.

On the other hand, if a service is designed to use the same URI and the implementers want to separate versioning from the API then using headers is probably a better option. In this case service consumers still have an option to request a specific version of the API but in a more subtle manner.

4.4 Service Discovery

In SOA, services have to have a mechanism to find each other. Service discovery is a key component of most large distributed systems. Since services need to communicate with each other they need to know the information about endpoints. Such information can be placed inside a caller's configuration file but this simplistic approach becomes problematic as the number of services grows.

For example, traditional retail store networks connect all network capable devices into one local subnet. Devices and servers get static IP addresses. Number ranges within the subnet are kept identical in all stores.

This simple approach reaches its limit, as the numbers of network capable systems are increasing, omnichannel solutions require exposure of services to a broader range of devices, UPOS peripherals become network aware and are also exposed as services and a variety of mobile devices hit the retail floor. Besides, more and more services are deployed in the cloud where provisioning and deallocation of the resources is even more dynamic. Updating statically configured services would affect too many devices so service discovery becomes crucial for higher flexibility and lower maintenance effort.

In this new dynamic world of intercommunicating services it is very difficult to maintain up-to-date correct information about all service endpoints. Hence a discovery mechanism is necessary to remove the dependency on brittle static configuration.

Service discovery is an essential aspect of SOA because it helps to avoids early binding of service consumers to particular service instances. Removing such coupling provides much greater flexibility for reconfiguration of the overall system and promotes service reuse.

4.4.1 Discovery Methods

In addition to static configuration approach in which service URIs are specified inside service consumer configuration file, there are other more dynamic techniques to discover service endpoints.

- Services can be designed to announce their appearance and respond to multicast discovery requests (for example, using UDP protocol).

- Services can leverage existing discovery mechanisms, like DNS.

- Service consumers can use centralized registry service that maintains information about available services.

4.4.1.1 Ad-Hoc Discovery

With ad-hoc discovery services and their consumers have to implement certain discovery support. Typically every discoverable service should announce its availability by broadcasting a special announcement message on start up. Service consumers have to listen to such announcements and to process them accordingly. On the other hand, service consumers should be able to broadcast special probing messages trying to discover a required service.

Ad-hoc discovery does not require any special centralized registry but the services and their consumers should be able to broadcast and reply to certain discovery messages. Typically this discovery type is implemented using UDP (User Datagram Protocol). UDP is a connectionless protocol and there is no direct connection required between the sender and the receiver. Therefore special UDP endpoints that support ad-hoc discovery have to be exposed by the services and the service consumers.

WS-Discovery standard [51] supports dynamic discovery and therefore can be used to implement ad-hoc discovery solutions. This protocol can only be used for the discovery of SOAP services since it specifically relies on WS specifications.

The major drawback of the ad-hoc approach is its complexity. There are also infrastructural limitations on how far the broadcast messages can reach.

4.4.1.2 DNS Discovery

Service discovery can be implemented using standard mechanisms of Domain Name System (DNS). DNS is a distributed data store for name and address information for computer hosts, services or other resources on a network. DNS data store resides on a hierarchy of special servers.

The primary function of DNS is to translate meaningful domain names (for example, www.nrf.com) into IP addresses.

This approach already provides a level of indirection that allows changing the physical location of a service (IP address and port number) while still keeping the same URI. Therefore, a service URI represents logical location and intent rather than physical location of the service instance.

DNS-based service discovery standard (DNS-SD) [52] allows clients to discover a list of services by type in a particular domain using standard DNS queries. DNS-SD uses ubiquitous, time-tested, powerful, and reliable internet technology.

A service instance is described using DNS SRV and DNS TXT record types. The following SRV record

_auth._tcp.example.org.   IN   SRV 0 0 113 security.example.org.

specifies that an authentication service (_auth) is available at port 113 on the host security.example.org. TXT records provide additional metadata about the service instance as key-value pairs. Definition of the metadata keys is defined by service type specification.

DNS service discovery uses the logical service naming syntax from SRV records but adds one level of indirection. First, the client queries PTR records that returns the list of available instances of a given service type. Thus, <Service>.<Domain> PTR-query returns zero or more PTR records that contain service instance names <Instance>.<Service>.<Domain>. Then service instance names are mapped to SRV/TXT records.

DNS-SD leverages existing reliable internet technology for discovery of services. To improve performance DNS systems extensively use caching techniques. Therefore, if service location information changes fairly often DNS-base service discovery can potentially return obsolete cached data. Thus it is important to provide proper time-to-live (TTL) configuration that balances performance and ability to detect changes.

4.4.1.3 Service Registry

The service registry approach is not new. One of the first implementations UDDI [53] was written in 2000. UDDI has not been widely adopted by the industry and major vendors withdrew their support for the standard. Still there are many implementations of registry based discovery.

Service discovery in retail has additional considerations. Upcoming UPOS v2 treats devices as services. Mobile POS systems often do not have all the peripherals attached directly to them. Registry-based discovery services can have additional metadata that, for example, would help to locate the closest receipt printer. Registry can also contain information about health of the service and failover options.

Another important consideration, especially in a retail store environment, is security. Service registry should be able to prevent rogue services from registration.

ARTS is working on the device services registry that will be released after UPOS v2.

4.5 Service Implementation Patterns

This section provides brief descriptions of some useful service implementation patterns. There is no intent to cover the patterns in detail. Rather, the goal is to show some practical implementation approaches that should be considered when designing services for a modern retail enterprise.

4.5.1 Idempotence

The terming idempotence comes from mathematics. It is used to describe mathematical operations that can be applied multiple times without changing the result beyond the initial application. For example, applying absolute value function multiple times results in the same value as applying it only once.

In the context of services the term idempotence means that invoking a service capability multiple times would not result in any unintended side effects. This property is very important since it means that idempotent operations can be safely retried.

In a distributed system a request to a remote service may result in timeout. In this case the caller does not know if the request was processed successfully, failed or is still being processed. It is possible that the request never reached the service or it takes too long to process. It is also possible that request was process successfully but the response could not reach the caller.

The key point is that if the service capability is known to be idempotent then caller can safely retry the request.

Some service operations are inherently idempotent. For example, almost all read operations are idempotent since they typically do not modify service data. Among four major HTTP methods GET, PUT, DELETE, and POST, the first three are idempotent.

Even if a service capability is not inherently idempotent it can still be implemented to have this important characteristic. For example, placing a new order using a POST request is not necessarily idempotent. But if a client application assigns a unique order ID to every order then the service can check if an order with such ID has already been placed.

Recommendation 4.5 Design Service Operations to Be Idempotent

Recommendation	Design inter-service communications to be idempotent.
Rationale	Since communications between services have to cross service boundaries it is possible that a request would return no response resulting in timeout. Idempotence makes it safe to retry the request until it succeeds.

Idempotence implies that at-least-once delivery would work exactly the same as only-once delivery.

To guarantee the delivery of messages it was a common pattern in enterprise systems to use queues and distributed transactions. In this approach queue manager and database management system would enlist in a distributed transaction to guarantee that a message that is received from the queue is successfully persisted in the database. Such approach creates tight coupling and is not suitable for queueing services in the cloud since cloud queueing services cannot be participate in a distributed transaction. However if the service capability that stores message in the database is idempotent it is still possible to provide guaranteed delivery of messages in the cloud environment.

Typical interaction with a cloud queueing service consists of the following steps.

1. Receive a message from the queue. The message becomes invisible for a configurable visibility timeout period so that other queue processors would not try to receive this message while it is being processed.

2. Process the message. This should take much less time than visibility timeout.

3. Delete the message from the queue.

4. If message has not been explicitly deleted from the queue during the visibility timeout then a failure of the processor is assumed and the message becomes visible again. This will effectively result in reprocessing of the message.

In this scenario it is possible that message has been processed successfully but the processor crashed before deleting it from the queue. However, if message processing is idempotent its reappearance on the queue will not cause any problem.

It might be useful to note that even idempotent operations still have to properly deal with concurrency. For example, even though making PUT request multiple times would always result in the same outcome it might override changes made by other service consumers. Obviously, the larger interval between the reties is, the greater is the chance of such collision.

4.5.2 Throttling

To provide consistent performance and availability services should be able to control the consumption of their resources by each client. It is especially important for multitenant services deployed in the cloud that have to meet certain SLAs (Service Level Agreements).

For many retail services the load can vary significantly based on many factors such as time of day, day of the week, holidays, weather, etc. There may also be sudden and unexpected bursts in activity.

If the processing requirements of the service consumers exceed the capacity of the service resources, it will suffer from poor performance. It is even possible that the service may fail completely, which could be unacceptable.

Typically cloud services provide a certain degree of elasticity but sometimes it takes certain time to provision new resources. So, it is quite possible that many service consumers could experience performance degradation just because one of the clients has a rapid increase of activity.

Sometimes it might be possible to use load leveling approaches, for example using queues, but it does not work well for interactive communications.

Simple throttling approach rejects service invocations from a certain client that has already accessed service API more than a particular predefined value. It implies that service keeps track of each client's activity.

More sophisticated throttling solutions can meter consumption of service resources and reject the calls that exceed a certain threshold. Again, it means that service should be able to meter the usage of its resource for each client (or tenant). Metering of resources is important aspect of cloud computing (see 1.4.1).

4.5.3 Retry

A service that communicates with other services, especially in the cloud, should be designed to deal with transient faults that are not uncommon in that environment. Such faults can occur because of a momentary loss of network connectivity, throttling (4.5.2), a temporary unavailability of the target service, etc.

Transient faults are typically self-correcting, and if the request that resulted in a fault is repeated after some delay it is likely to succeed. For example, an entity service that is processing a large number of concurrent requests for certain data may implement a throttling strategy that temporarily rejects any further requests until its workload has eased. An application attempting to get the data may get the error that indicated that it was throttled due to high load, but if it tries again after a reasonable delay it is likely to succeed.

Recommendation 4.6 Implement Retry Logic for Transient Failures

Recommendation	Implement retry logic to handle transient errors that happen during communications with a remote service.
Rationale	Since communications between services have to cross service boundaries it is possible that a request would return no response resulting in timeout. Idempotence makes it safe to retry the request until it succeeds.

The retry approach should be implemented only if the failure is transient in nature. Many failures, like denied access or violation of business constraints, are unlikely to disappear no matter how many times operations are retried. Thus proper classification of potential failures is a key to implementing proper retry logic.

Type of Failure	Retry Logic
Non-transient errors, such as access violation, database constraint violations, business exceptions, etc.	No retries. An appropriate exception should be reported.
Transient rare, one-off errors, such as corrupted network packet.	Retry can be done immediately after the failure since it is unlikely to occur again.
Transient common errors typically associated with some kind of resource contention.	Retry after an appropriate delay.
Timeout error.	If operation that resulted in timeout is idempotent (4.5.1) then retry is good approach, otherwise an appropriate exception should be reported.
Unknown error.	In this situation the best course of action mostly depends on the nature of the operation.

It is important to note that aggressive retry logic can aggravate the situation. It is not a good idea to keep re-submitting requests to a service that might be experiencing problems with handling the current workload. One approach to deal with this issue is to increase the retry interval after every transient failure. Exponential backoff [54] is a commonly used algorithm that doubles the retry interval until it reaches a certain threshold.

Another approach to avoid making a lot of repeated calls is to implement the Circuit Breaker pattern [55]. In this approach a special circuit breaker component after a number of failed attempts prevents further communications with the service (breaks the circuit). This component can probe the service to determine if the problem has been resolved. If the service appears to function properly the client is allowed sending new requests.

These techniques to reduce the number of retries are especially important in the cloud environment where every service request may result in incurring additional costs. Since clients often communicate with services to obtain some data, putting that data into a cache can improve performance and reduce the number of requests to the services that provide the data.

4.5.4 Gateway

Gateway is a common approach to consume service APIs. It represents a single access point and functions as a proxy for one or more services.

A service gateway is especially useful in the context of microservices (4.2) since clients can access multiple granular services through a single endpoint. The client is presented with a

single façade that hides the complexity of a horde of services working together to provide necessary capabilities.

Using a service gateway insulates clients from the complexity of a distributed system built as set of granular services. Clients don't have to deal with determining the locations of service instances.

A service gateway can also improve the performance since a single request can carry the payload that can be used to communicate with multiple services in a single round-trip.

In addition to simplifying access to service capabilities the gateway can perform a number of useful functions such as authentication and authorization, validation, routing, discovery, transformations, logging, etc.

Recommendation 4.7 Implement Service Gateway

Recommendation	Implement a service gateway that presents consumers with a single endpoint and abstracts the complexity of the implementation based on multiple granular services.
Rationale	It is much easier for service consumers to communicate with the gateway than multiple services that might be required to support the client application. Also, service gateway can address useful cross-cutting concerns like security, discovery, logging, etc.

Interestingly, Software-Defined Architecture (SDA) described in [17] uses special SDA gateway that separates services from consumers by virtualizing internal services APIs. In this approach, SDA gateway exposes API that is much easier to consume than application agnostic APIs exposed by the underlying services.

5. SERVICE SECURITY

Modern services (APIs) should be designed so that they could be deployed in different environments. Many retail enterprises are transitioning some of their system into the cloud, which brings some new security considerations. Also, the adoption of mobile devices as a platform for retail operations that need to access those APIs adds even more security requirements. Concerns about security and relative complexity of the issues force many businesses to proceed with great caution in moving some of their services to the cloud. Security challenges have emerged as some of the most significant obstacles to faster and more widespread adoption of cloud computing.

With the advent of mobile and cloud computing in modern retail enterprises, the traditional enterprise security based on some kind of directory services began to experience serious difficulties. The problem was that the conventional approach assumed that all the resource and users are managed by centralized enterprise security systems tightly controlled by IT departments; with cloud and mobile IT has lost that control.

To cross trust domain boundaries SOAP-based technologies (WS-Security, WS-Federation, WS-Trust, etc.) were often used but they led to quite heavy solutions that relied on SAML (Security Assertion Markup Language). Because of their size, these technologies were inadequate for mobile devices that originally came from the consumer space and were poorly equipped to deal with heavy XML processing.

To communicate in cloud and mobile world, it is absolutely crucial to make sure that APIs are secure. If implemented correctly APIs can provide a way for retailers to enable new innovative business processes, to expand into entirely new markets, and to acquire new customers. On the other hand, if services endpoints are not properly secured then APIs can open the enterprise up to a huge array of potential attacks. Any security breaches can cause major disruption to retailer's operations and become a public relations nightmare. It is especially true for public APIs that often become a target for hackers.

It would be a mistake to assume that the same methods and techniques that were used to secure the traditional browser-based web applications can be used to protect service endpoints. Even though it is true that APIs share many of the same threats that plague the web, they are fundamentally very different and have entirely new security risks that must be addressed.

A lot of security aspects are covered in ARTS Security Technical Report [56]. This chapter is focused on the security aspects that are more specific to the implementation of services such as transport security, authentication, and authorization.

5.1 Transport Security

Transport security is responsible for confidentiality and integrity of data transferred over a computer network.

RESTful APIs are implemented on top of HTTP protocol that by itself does not deal with transport security issues. Therefore, it is the common practice to use HTTP on top of a secure transport layer resulting in what is also known as HTTPS.

RFC 5246 [57] defines the Transport Layer Security (TLS) protocol that is based on SSL protocol specification published by Netscape. The TLS protocol uses X.509 certificates to authenticate communicating parties and to negotiate a symmetric session key. This shared

key is used to encrypt the data on the wire. The protocol provides data integrity and confidentiality by making sure the data on the network between the peers of the TLS session has not been tampered with and it has not been exposed to a third party. Thus, one of the main goals of using transport security is to protect against man-in-the-middle attack (MITM).

In addition to integrity and confidentiality TLS protocol can be used to authenticate the communicating parties. In practice, however, TLS is most commonly used to authenticate the server in order to guarantee that data is exchanged with a legitimate party. Clients more often use some other authentication mechanism over an already established secure TLS session.

Another approach to guarantee integrity and confidentiality is to use message level security. It means that it secures messages rather than whole communication pipe. For example, SOAP-based services can use WS-Security that is part of WS-I Basic Security Profile [58]. WS-I (Web Services Interoperability) [37] is an OASIS Member Section focused on promoting best practices for interoperability of SOAP-based web services.

WS-Security cannot be recommended as a general approach since it is only applicable to web services that use SOAP. Also, WS-Security adds significant performance overhead. As it was pointed out in an article dedicated to the subject of WS-Security performance [59] WS-Security added an order of magnitude overhead when it was compared to just encryption and signing of 100KB array of data. Therefore, even SOAP services should use WS-Security diligently only when a specific feature like end-to-end security is necessary or if a transport level security is not available.

Currently the most practical approach to achieve integrity and confidentiality of the transferred data is to use the TLS protocol.

Recommendation 5.1 Use the TLS Protocol

Recommendation	Use TLS to secure HTTP communications with services.
Rationale	The TLS protocol makes it much less likely that communications between a service and its consumers will be exposed to and/or manipulated by a malicious third party. It is especially important when privileged information such as security credentials or payment data is exchanged between the parties. The minor TLS performance overhead is a small price to pay for the provided security of the data in transfer.

As mentioned above, the TLS protocol uses X.509 certificates. The security of the certificate signature depends on the strength of the hashing function. The problem is that a lot of certificates today use the SHA-1 hashing algorithm, which does not provide enough security. The collision resistance of SHA-1 algorithm became a major concern among security experts. For this reason, the industry is transitioning to more secure SHA-2 ciphers.

Recommendation 5.2 Use Strong Certificate Signature Algorithm

Recommendation	Use strong certificate signature algorithms, like SHA-2.
Rationale	The SHA-1 hashing algorithm is potentially insecure. Certificates that use SHA-1 are planned to be phased out by the end of 2016.

5.2 Identity and Access Management

There are a variety of authentication and authorization mechanisms for different types of services and scenarios. It is important to note that management of keys and security credentials is crucial to successful implementation of a solution, especially in the cloud. If credentials or keys are compromised even the strongest security mechanisms provide no protection.

To better understand the issues with securing Web APIs it might be useful to consider typical scenarios of how services can be accessed in today's retail enterprise.
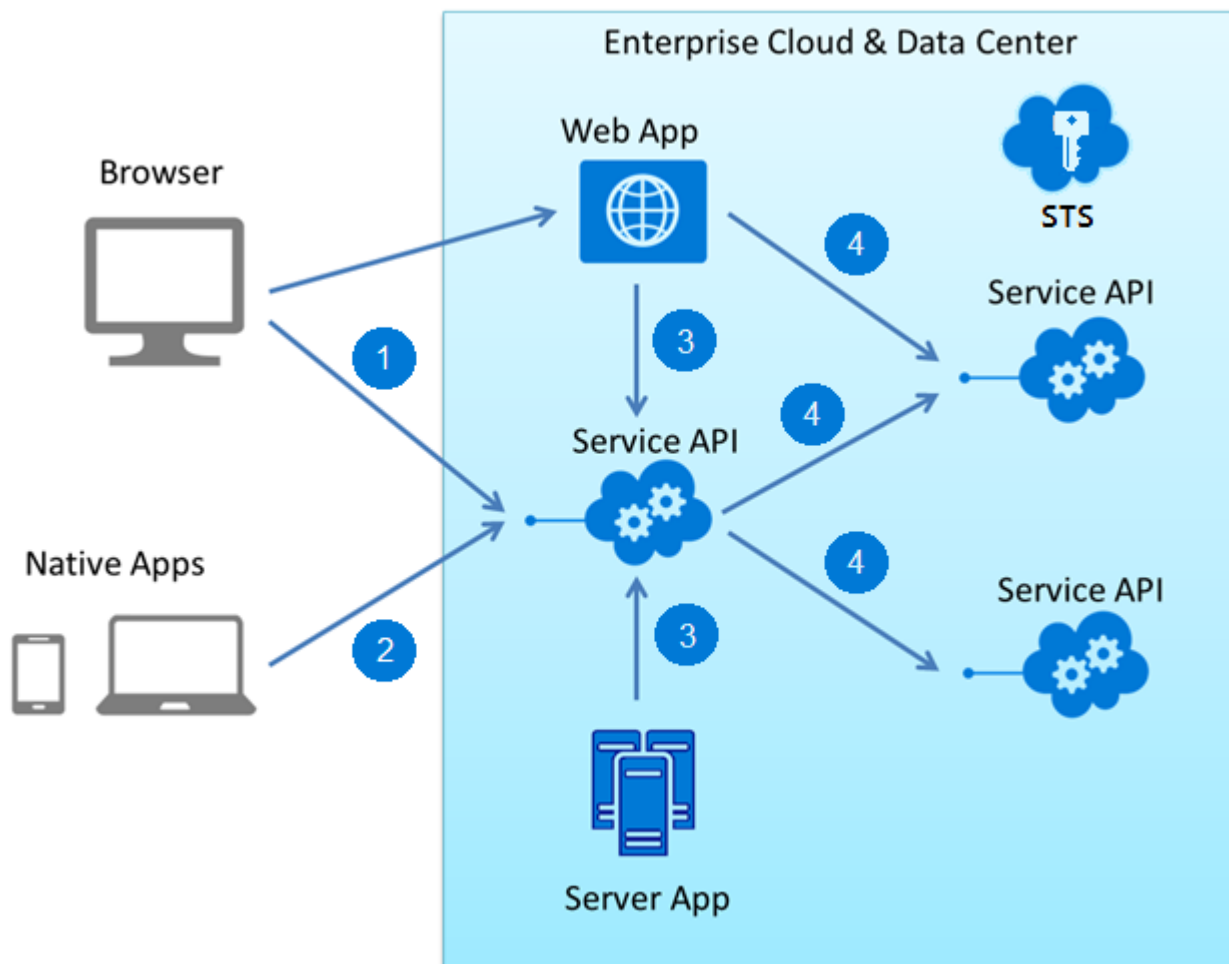
Figure 5.1 Service Access Scenarios

1) *Browser to frontend service call.* Services can be accessed by applications that run in a browser. These are typically AJAX/AJAJ (Asynchronous JavaScript and XML/JSON) calls performed from JavaScript that is executed in a browser.

2) *Native application to frontend service call.* Services can be accessed by a variety of native applications that run on different devices (mobile, laptop, desktop, etc.).

3) *Server application to frontend service call.* Service can be accessed either by web applications that execute on a server side or other server applications. The difference here is that API invocations are performed on the server side in a more secure and controlled environment.

4) *Service to backend service call.* Of course, services can also invoke other background services that might not even be exposed to the outside clients.

The four scenarios shown above can be handled differently from the security point of view and there are two complementary approaches: trusted subsystem and delegated access authorization.

5.2.1 Trusted Subsystem

Trusted subsystem is a common approach used in web development [60]. In this case, user communicates with a web application that provides frontend services (UI markup rendering, processing HTTP requests, etc.) and this web application talks to services that run in the back. The backend services trust the web application therefore it represents a trusted subsystem. In this scenario, even if the user is authenticated the services typically do not care which particular user communicates with the web application. In trusted subsystem a service trusts the direct caller.

5.2.2 Delegated Access Authorization and STS

With delegated access authorization the services need to know the user so that they can control which resources the calling application can access in the context of a particular call.

A special software based identity provider called Security Token Service (STS) [61] can help to address security concerns for these different use cases in a consistent manner.

If a service is exposed as an open API and can be accessed by different types of clients it is critical that service capabilities are provided in a manner that safeguards the security and privacy of customers and a retail enterprise.

Only authorized client application should be able to successfully access the service API. Although simple login based authentication is fairly easy to implement and might work fine for a simple system, it creates some of the following serious issues for enterprise solutions.

Supplying credentials with every API call increases the risk that they can be compromised and if it happens the attacker gets complete control over all the user's resources. If the same credentials are used to access multiple services, which is fairly common, then compromising a single service puts the whole system at risk. This widens the surface of attack and increases potential damage. It also significantly complicates the revocation process if a security breach is detected.

To be able to supply the credentials for every call they should be kept in memory, and since re-entering passwords, especially on mobile devices, can negatively impact user experience client applications often opt for storing user credentials on the device. Such approach potentially creates additional vulnerabilities.

Also, the service has to validate passwords on every request, which can incur significant computational costs because of the special techniques used to protect against dictionary attacks [62].

Another problem with using passwords comes up if users want to allow third-party software to access their resources via API. If the API requires user's credentials to be passed with a request then the only way to make it work would be to share the password with third-party software. Again it increases the chance for exposure.

Security tokens provide a more granular and time-constrained mechanism to grant access to certain resources exposed by API. Users can use their credentials to obtain a security token that can be given to an application to get access to secured resources.

There are two major types of security tokens that are used for controlling access to APIs: an access token and an identity token. In RFC 4949 [63] they are referred to as correspondingly capability token and authentication token and defined in the following way:

- *Capability (access) Token* – A token (usually an unforgeable data object) that gives the bearer or holder the right to access a system resource. Possession of the token is accepted by a system as proof that the holder has been authorized to access the resource indicated by the token.

- *Authentication (identity) Token* – A data object used to verify an identity in an authentication process.

Since dealing with security issues is not a trivial task it is not a prudent approach to burden every service with handling complex activities like authentication and delegated access authorization. Therefore, it is a good practice to use a common Security Token Service that is focused on keeping track of the users and issuing security tokens.

Recommendation 5.3 Use common STS to issue scoped and time-limited security tokens.

Recommendation	Use centralized STS that deal with security issues like authentication and delegated access authorization.
Rationale	Security software is notorious for being difficult to implement correctly and requires special skills. Software developers that develop business APIs often are not security experts. Besides it does not make sense to implement security handling inside every service again and again. STS should issue tokens that have limited scope and expiration date and time.

5.2.3 Tokens and Security Protocols

There are different types of tokens and different security protocols that can be used to facilitate the API access scenarios described above.

One of the most popular protocols used today when communicating with Web APIs is OAuth2 [64]. OAuth2 is a protocol that allows clients to obtain an access token from the token service and then to use them when accessing the API. The key here is that the user authenticates with the token service and the business domain service only has to handle security tokens and does not have to deal with complicated user authentication and identity management issues.

There is one very important nuance here. OAuth2 is not an authentication protocol. At the time when a client application accesses an API there is no guarantee that the user who authenticated with the token service is even present. The whole idea behind OAuth2 is that an access token can be given to a client application and then that application can use the token much later until it expires.

There are two major protocols that have been used to deal with authentication and they both are based on SAML security tokens. Both protocols can handle federated user identities that

can be shared among multiple identity management systems, which makes authentications techniques like Single Sign-On (SSO) possible. Both protocols are maintained by OASIS. The first, more widely accepted protocol is SAML 2.0 [65]. It has been mostly used by Java community. The second protocol is WS-Federation [66], which is a part of WS Security framework that applies to SOAP services.

But according to many industry experts SAML is headed down the legacy path. At the 2012 Cloud Identity Summit, KuppingerCole's Distinguished Analyst Craig Burton pronounced that "SAML is dead". That statement stirred up quite a bit of controversy at the time. Burton called SAML "Windows XP of Identity" and qualified his remarks further: "SAML is dead does not mean SAML is bad. SAML is dead does not mean SAML isn't useful. SAML is dead means SAML is not the future."

Since OAuth2 is the most popular protocol to deal with securing access to APIs it makes a lot of sense that OpenID Connect [67] authentication protocol that is based on OAuth2 is quickly gaining momentum and becoming the protocol of choice for identity management.

5.2.4 OAuth 2.0

According to OAuth website [64], "OAuth 2.0 is the next evolution of the OAuth protocol which was originally created in late 2006. OAuth 2.0 focuses on client developer simplicity while providing specific authorization flows for web applications, desktop applications, mobile phones, and living room devices".

Basically, OAuth2 is a standard for delegated access authorization over HTTP protocol. It is defined by RFC 6749 [68].

With OAuth2 an application gets access rights to an API using an access token. If the application is not trusted, the user does not have to provide it with login credentials. Instead, the user first communicates with the authorization server (STS) and a special security token is passed to the application. The security token has limited scope and grants access to a subset of data for a limited time interval, which is a much more secure approach than directly using user credentials.

OAuth2 uses so-called bearer tokens that are simpler to use but should always be communicated over a secure channel with some kind of transport security. RFC 6750 [69] defines bearer token as "a security token with the property that any party in possession of the token (a "bearer") can use the token in any way that any other party in possession of it can. Using a bearer token does not require a bearer to prove possession of cryptographic key material (proof-of-possession)".

5.2.4.1 OAuth 2.0 Roles

To accommodate different scenarios OAuth2 specification [68] defines four distinct roles.

- *Resource Owner* – An entity capable of granting access to a protected resource. When the resource owner is a person, it is referred to as an end-user.

- *Resource Server* – The server hosting the protected resources, capable of accepting and responding to protected resource requests using access tokens.

- *Client* – An application making protected resource requests on behalf of the resource owner and with its authorization. The term "client" does not imply any particular implementation characteristics (e.g., whether the application executes on a server, a desktop, or other devices).

- *Authorization Server* – The server issuing access tokens to the client after successfully authenticating the resource owner and obtaining authorization.

It is important to note that OAuth2 makes clear separation between end-user (resource owner) and an application (client) that needs to access the resources. This separation allows placing them into their own security boundaries. That means that client application that is used to access protected resource can be treated as untrusted system and therefore resource owner credentials should not be exposed to it.

For example, a retailer might offer customers access to their digital receipts via a website. A customer that frequently shops at retailer's stores signed up for third party budgeting services. If the budgeting application can consume standard ARTS digital receipts [70] and the retailer exposes an API that serves standardized digital receipt documents, then the budgeting app can import the digital receipts from the retailer. This example clearly shows the difference between the resource owner (customer) and the client (budgeting app) that accesses the resource server (retailer's digital receipts API) on behalf of the customer.
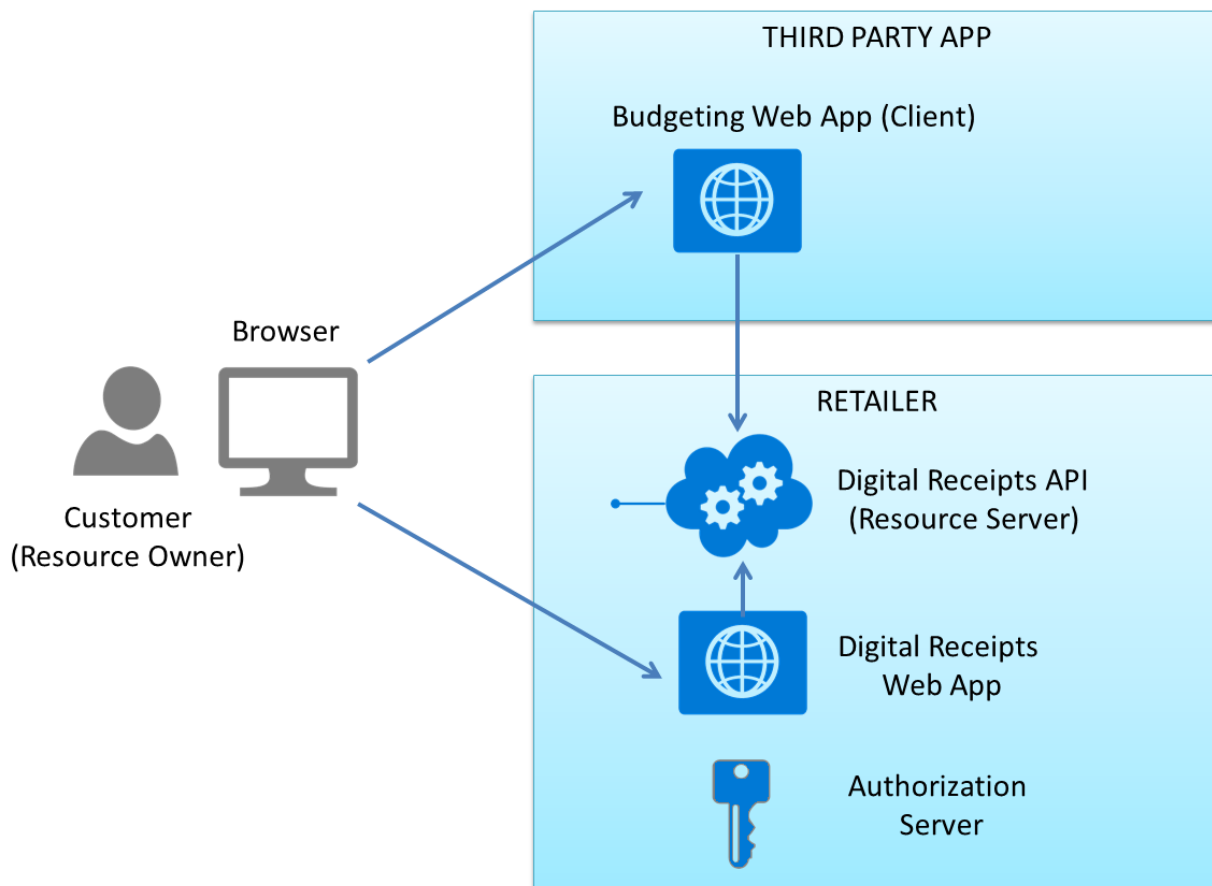
Figure 5.2 Resource Owner vs Client

If the retailer and the budgeting app support OAuth2 protocol then the customer can use retailer's authorization service to obtain an access token that can be passed to the budgeting app. OAuth2 protocol specifies the orchestration of this process and the only thing the customer has to do is to login into retailer's service and approve the consent form.

The customer as a resource owner of the digital receipts (protected resource) can of course access them directly using retailer's website. But in the context of OAuth2 it is more important

that the customer can delegate the access rights to a third party application. Of course, the Digital Receipts API should be able to validate security tokens but it is not involved in the authentication process and can process requests from retailer's website as well as from any third party application in a similar fashion.

To be able to communicate with OAuth2 authorization server a client must go through a registration process, that is not defined by OAuth protocol. As part of the registration process the client provides the authorization server with some important data like redirect URIs, authorization scopes, etc. The authorization server issues the client a special identifier (client_id) and for some client a special password (client_secret). These client credentials are used for authentication of clients to the authorization server. Client applications that can store client_secret securely are called confidential clients. Some clients do not have capability to securely store confidential information. Such clients are called public.

It is a good practice to supply a special state parameter with requests to authorization server. If the state parameter was provided than the response must also include state with the same value. This helps to mitigate against cross-site request forgery.

5.2.4.2 OAuth 2.0 Endpoints

The OAuth2 specification defines three special endpoints that are used by different flows to accomplish the access authorization process. These endpoints represent RESTful services that behave according the OAuth2 specification. Not every flow utilizes all three endpoints.

The authorization server exposes two endpoints.

- *Authorization endpoint* – used by the client to obtain authorization from the resource owner via user-agent redirection. The user-agent communicates with this endpoint to obtain an authorization grant. The authorization server must authenticate the user before such grant can be issued. The OAuth2 does not specify the authentication mechanism. After the authentication is completed the resource owner can be explicitly asked to confirm the delegation of the access rights for the protected resource.

- *Token endpoint* – used by the client to exchange an authorization grant for an access token, typically with client authentication. The client uses HTTP basic authentication to gain access to the token endpoint. The idea is the token endpoint can be only accessed by the clients that went through the registration process and therefore are known to the authorization server.

The client exposes one endpoint.

- *Redirection endpoint* – used by the authorization server to return responses containing authorization credentials to the client via the resource owner user-agent. This endpoint is typically configured in the authorization server during the client registration process. The redirection endpoint is not called directly. It is accessed via HTTP-redirect command (HTTP status code 302) from the authorization server.

5.2.4.3 OAuth 2.0 Tokens

The OAuth2 uses so-called bearer tokens that are simpler to use but should always be communicated over a secure channel with some kind of transport security. RFC 6750 [69] defines bearer token as "a security token with the property that any party in possession of the token (a "bearer") can use the token in any way that any other party in possession of it can.

Using a bearer token does not require a bearer to prove possession of cryptographic key material (proof-of-possession)".

Token information is stored in a database inside authorization server. There are two major types of OAuth2 tokens.

Access Token

Access tokens are used by the client application to access protected resources via API exposed by the resource server. Since bearer access tokens are very sensitive they are typically valid for relatively short time. Also short lifetime provides an opportunity for token revocation. Typically access tokens are opaque to the client application. There are two types of access tokens: self-contained and reference. A self-contained token includes security information and a reference token represents a cryptographically strong identifier that is used to retrieve the security information. Reference tokens have smaller size and the resource server can be designed to validate that the access rights have not been revoked. However they require an extra step of retrieving the security data.

Refresh Token

Since frequent logins would be very burdensome for the users OAuth2 includes special refresh tokens. When the client application obtains the access token from the authorization server the response may also include a refresh token. The refresh token is used to obtain a new access token when the current access token is about to expire.

5.2.4.4 OAuth 2.0 Flows

The OAuth2 can handle different scenarios and different types of client applications (web, mobile, browser-based JavaScript apps, desktop, etc.). The specification defines four so-called flows that describe requests and responses sent by different OAuth roles to obtain security tokens.

The detailed description of OAuth2 flows goes beyond the scope of this document but high level overview can be useful to understand the capabilities and the applicability of a certain OAuth2 flow to a particular scenario.

Authorization Code Flow

Authorization code flow is typically used when the client is a server-side web application. In this case client_secret can be stored securely. It is the most secure and complex flow. The authorization server authenticates the resource owner using the credentials. Since the resource owner typically uses browser (often referred to as user-agent) the authorization server certificate is validated. The authorization server authenticates the client using client_id and client_secret and the client authenticates the authorization server using the certificate and URI. The resource owner's credentials are never shared with the client. Also, the access token is passed directly to the client bypassing the user-agent.
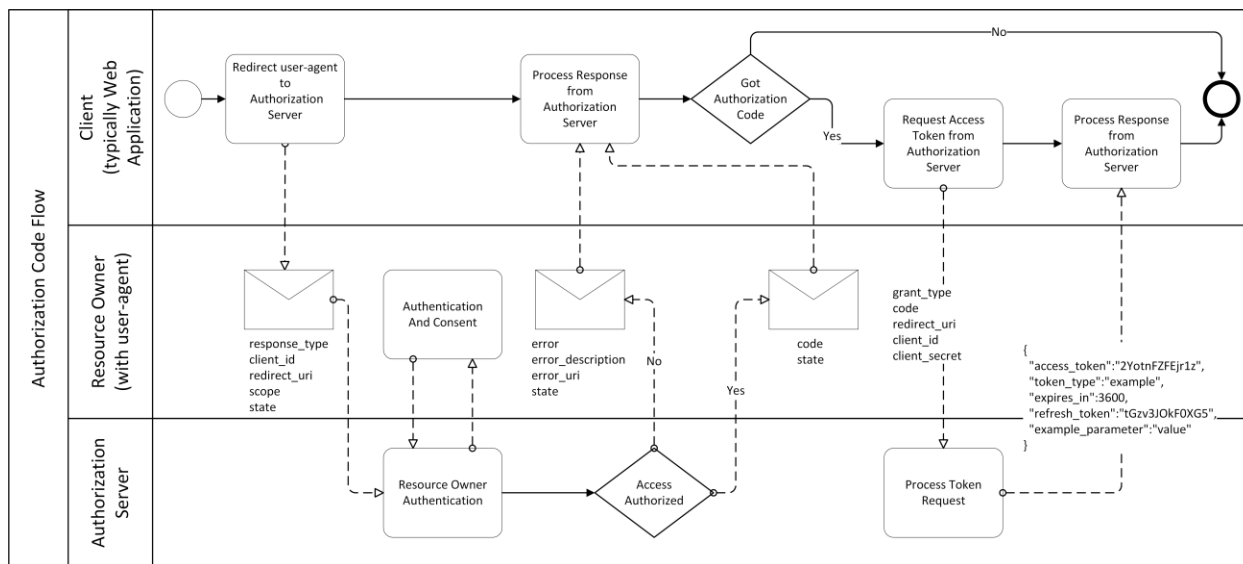
Figure 5.3 Authorization Code Flow

The client application initiates the authorization code flow by redirecting the resource owner's user-agent to the authorization server. The response_type must be set to "code".
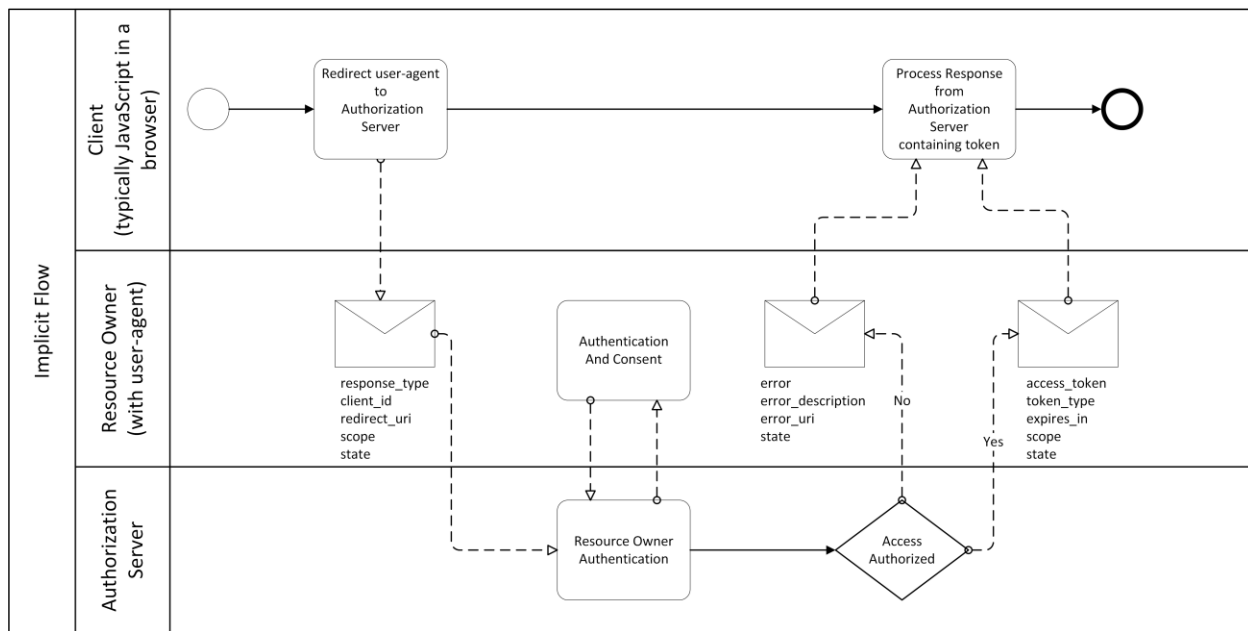
The authorization server authenticates the resource owner. OAuth2 does not specify how the user is authenticated. It gives the authentication server the flexibility to choose the most appropriate mechanism ranging from simple form-based user name and password to multi-factor authentication (MFA). Typically after successful authentication authorization server will ask for user's consent to give the client access to certain protected resources. Some implementations allow users to specify the time interval during which the access token is valid.

If access to the resources is granted the authorization server redirects the user-agent back to the client application using the redirection URI. The authorization code is embedded inside the URI as a query parameter.

It is important to note that authorization code is only valid for a short period of time. The client application authenticates with the authorization server and sends the authorization code as a query parameter. The grand_type must be set to "authorization_code". If the process is successful the client receives the access token and possibly a refresh token.

Implicit Flow

Implicit flow is typically used when clients cannot securely store client_secret and refresh token. Such clients are typically implemented as JavaScript code running in a browser. So, the implicit flow is a simplification of the authorization code flow as the client authentication part does not make sense since client_secret is not available. The client receives the access token as a result of the authorization request. This flow is less secure that the authorization code flow and the access token may be exposed to the resource owner.

Figure 5.4 Implicit Flow

The client application initiates the authorization code flow by redirecting the resource owner's user-agent to the authorization server. The response_type must be set to "token".

The authorization server authenticates the resource owner. The step is the same as for the authorization code flow.

If access to the resources is granted the authorization server redirects the user-agent back to the client application using the redirection URI. The authorization token is embedded inside the URI. In addition to access token the URI contains token_type and may also include additional parameters like scope, state, and expire_in.

Resource Owner Password Credentials Flow

Resource owner password credentials flow is typically used with so-called trusted clients. Trusted clients are usually native applications that are either were developed internally or came from a trusted source. In this scenario the resource owner provides his credentials directly to the client. This flow is the closes to enterprise style authentication used in the past.

There is a significant difference between the resource owner password credentials flow and using passwords to access the resource server. In this flow the password is only used to obtain an access token. Then, right after the token was received, the username and password of the resource owner should be discarded. The client application still can store the access and refresh tokens if secure storage is available.

Since flow has limited applicability since it can only be used with certain types of clients but is fairly simple and requires only a single call to obtain the tokens.
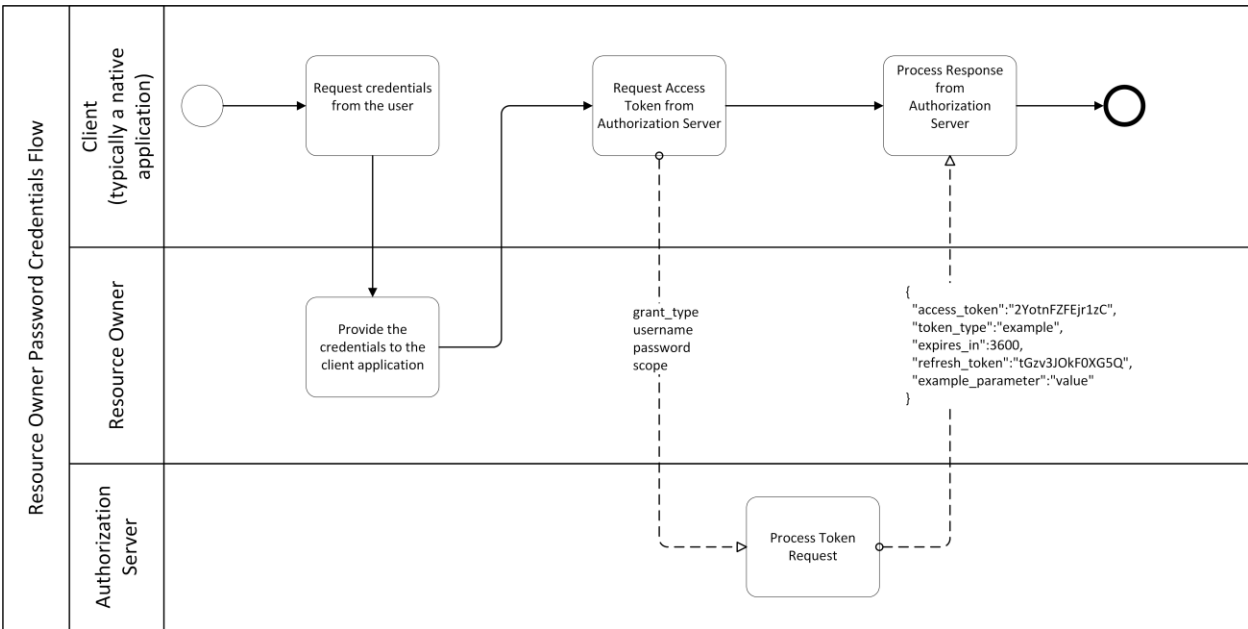
Figure 5.5 Resource Owner Password Credentials Flow

The client application displays some kind of form to the user to enter the username and password.

After the resource owner provides the credentials the client request the access token from the authorization server. The grand_type must be set to "password". The username and password are also supplied as part of the request.

It is important to note that clients that were issued client credentials during the registration process must authenticate with the authorization server.

Since the refresh token can be also returned with this flow, a new access token can be requested without having to prompt the resource owner for the username and password.

Client Credentials Flow

In this scenario no resource owner is present. Conceptually this flow describes a scenario where the client application owns the protected resource. For example, if a web application that runs on the server-side has access to a resource the client credentials flow can be used to get the access token.
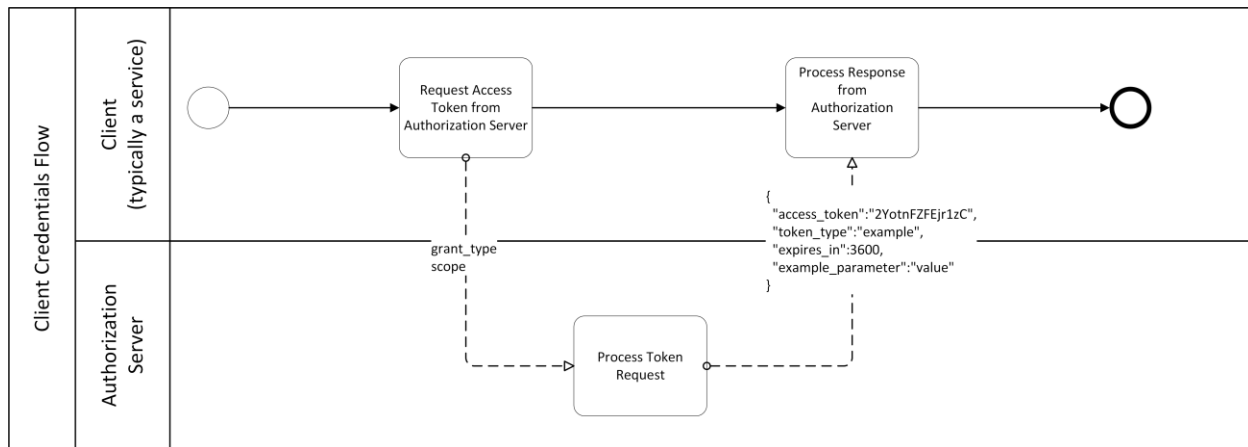
Figure 5.6 Client Credentials Flow

The client application authenticates with the authorization server and requests an access token. The grand_type must be set to "client_credentials".

The authorization server validates the client credentials, and if valid, issues the token.

Refresh tokens should not be used with this flow since it is typically implemented in more secure server-side context and the access token can be obtained without interactions with the user.

Summary

The table below summarizes the typical usage of different OAuth2 flows.

	Authorization Code Flow	**Implicit Flow**	**Resource Owner Credentials Flow**	**Client Credentials Flow**
Typical Client	Server-side web application.	JavaScript application running in a browser. Also untrusted native applications.	Trusted native applications (mobile, desktop)	Server-side applications that given access to resources regardless of the user.
Client Type	Confidential	Public	Public	Confidential
response_type auth. Endpoint	code	token	Does not use authorization endpoint	Does not use authorization endpoint
grant_type token endpoint	authorization_code	Does not use token endpoint	password	client_credentials

Refresh Token	Typically used.	Not used.	Typically used.	Should not be used.

OAuth2 is widely accepted framework to access protected resources on the web.

Recommendation 5.4 Use OAuth 2.0 to secure APIs

Recommendation	Use OAuth 2.0 specification to secure access to service APIs.
Rationale	OAuth2 is the standard framework for securing access to modern APIs. It was specifically designed to work well with RESTful services and lightweight mobile application. OAuth2 is widely accepted in the industry and major software companies have tools supporting the protocol. Many modern cloud, mobile and web applications use OAuth2 under the hood.

5.2.5 JSON Web Token

JSON Web Token (JWT) is a JSON-based open standard security token that is used to represent claims [71] to transfer them between parties. It is specified in RFC 7519 [72].

Even though OAuth2 does not mandate the use of the JWT, it is the commonly used type of token to pass claims information. The OpenID Connect protocol [67] mandates the use of the JWT.

The JWT has two parts: header and a set of claims. The header contains metadata about cryptographic algorithms and properties. The RFC defines the following standard claims that can be used inside a JWT claim set:

- Issuer (iss) – principal that issued the JWT.

- Subject (sub) – principal that is the subject of the JWT.

- Audience (aud) – recipients that the JWT is intended for.

- Expiration time (exp) – expiration time on or after which the JWT must not be accepted for processing.

- Not before (nbf) – time before which the JWT must not be accepted for processing.

- Issued at (iat) – time at which the JWT was issued.

- JWT ID (jti) – unique identifier for the JWT.

The JWT can also contain other claims as part of the claim set.

Here is an example of the JWT:

Header	**{**
	  "alg": "HS256",
	  "typ": "JWT"
	}

Claims	{
	
	  "sub": "1234567890",
	
	  "name": "John Doe",
	
	  "admin": true
	
	}

Signature	HMACSHA256
	
	(
	
	  base64UrlEncode(header) + "." +
	
	  base64UrlEncode(claims),
	
	  secret
	
	)

To get the representation of the JWT on the wire the header and the claim set are Base64url encoded and combined together separated by the period ('.') character. Then the signature of the resulting string is calculated and Base64url encoded as well. The signature is appended at the end separated by the period ('.') character:

**eyJhbGciOiJIUzI1NiIsInR5cCI6IkpXVCJ9.eyJzdWIiOiIxMjM0NTY3ODkwIiwibmFtZSI6Ikpv
aG4gRG9lIiwiYWRtaW4iOnRydWV9.TJVA95OrM7E2cBab30RMHrHDcEfxjoYZgeFONFh7
HgQ**

The result is three Base64url encoded strings separated by dots that can be easily passed in HTML and HTTP environments.

5.2.6 OpenID Connect

As was stated earlier the OAuth2 was designed to deal with delegated access authorization not authentication. However, most implementations of the OAuth2 include some kind of user authentication that is necessary to issue a grant to access protected resources. Because of that there have been several adaptations of the OAuth2 to handle the user authentication. Since the OAuth2 itself does not have enough features to accomplish the authentication securely such implementation resulted in security vulnerabilities and incompatible implementations.

OpenID Connect [67] protocol is designed to address these problems with using OAuth2 for authentication. It provides a standard authentication layer on top of the OAuth2 framework. The OpenID Connect website defines it in the following way:

"OpenID Connect 1.0 is a simple identity layer on top of the OAuth 2.0 protocol. It allows Clients to verify the identity of the End-User based on the authentication performed by an Authorization Server, as well as to obtain basic profile information about the End-User in an interoperable and REST-like manner. OpenID Connect allows clients of all types, including Web-based, mobile, and JavaScript clients, to request and receive information about authenticated sessions and end-users. The specification suite is extensible, allowing participants to use optional features such as encryption of identity data, discovery of OpenID Providers, and session management, when it makes sense for them."

OpenID Connect standardizes the concept of the so-called UserInfo endpoint that can be used to get user profile information. It also introduces a special identity token (id_token) that contains user claims intended for the client. Therefore, the identity token is very different from the access token. The access token is opaque to the client and intended for the resource server. The identity token should be validated by the client application immediately upon its receipt.

As was mentioned earlier OpenID Connect flow is implemented on top of OAuth2 flows. Conceptually it is similar to getting the authorization to access the UserInfo endpoint. So, OpenID Flow uses OAuth2 flows with some important authentication augmentations. To initiate such a flow the scope parameter inside the request to the authorization endpoint must begin with openid. An authorization request that starts with openid scope is called an authentication request. The mandatory openid scope is followed by one or more optional scope values of profile, email, address, phone, and offline_access that have associated sets of claims.

Scope	Claims
profile	name, family_name, given_name, middle_name, nickname, preferred_username, profile, picture, website, gender, birthdate, zoneinfo, locale, updated_at
email	email, email_verified
address	address
phone	phone_number, phone_number_verified
offline_access	Requests refresh token

The user goes through the same authentication process as defined by OAuth2 flow and then the client communicates with the authorization server token endpoint to obtain the security tokens. In this case the client application receives both access_token and id_token. The access token is the standard OAuth2 access token that can be used to access the UserInfo endpoint. The ID token is represented as a JWT that contains claims about the authentication event. It must be validated by the client. Since a JWT contains the audience and expiration claims the client application can verify that the ID token has been issued for its use and is relatively recent and therefore the ID token means successful authentication.

The subject claim provides unique identification of the authenticated user. If more information about the user is necessary then the client can access the UserInfo endpoint using the access token.

5.3 API Vulnerabilities in Retail Store

There are certain threats that should be taken into consideration when store associates use applications on devices to access a backend service.

Just because employees cannot readily see communications between devices in their possession and service APIs, does not mean that these communications cannot be quite easily discovered and even manipulated. Unfortunately it is true even if communications are performed over TLS.

Best Practices for Services Implementation Using ARTS Standards

With proper tools it is possible to get under the veneer of an application and see how it communicates with APIs. Therefore, implementers of backend services should not assume that network traffic is not discoverable because such assumptions may lead to additional vulnerabilities.

One of the attacks that can happen in a retail store environment is so-called men-in-the-middle (MITM) attack. It occurs when the attacker is sniffing the communication channel (Wi-Fi, ISP, hijacked connection, etc.). In this scenario the device, the client application and the backend service are not compromised but vulnerability exists in one of the hops between the device and the backend server hosting the service. The most common way to protect against MITM attack is to use transport layer security. Other defenses include securing sensitive data inside the payload via encryption or tokenization.

The situation can be potentially even more dangerous if the attacker controls the device that runs retailer's application. There are multiple network sniffing tools that can be configured as a proxy for the device.

Even if devices do not allow configuring a proxy they can be breached by using special network auditing tools that can either be configured as a wireless access point or can be physically connected using network cables.

Using these techniques the attackers can easily discover service endpoints. This information can be used to attempt a denial of service (DoS) attack and force an offline situation inside a retail store. The attackers can also learn about size of payload and API flow. They can even manipulate the payload to perform unauthorized activity.

The implementers of services should be mindful about such potentials threads and avoid designing leaky APIs that disclose more information that is necessary for the client application to function. Also, API developers should be careful about error conditions and what kind of information is communicated back to client application to avoid disclosing internal implementation details.

If attackers control devices that are used to communicate with services they can use a fake certificate to bypass the transport layer security. If the fake certificate is installed on a device, then a network sniffing tool can use it to appear as a trusted server to the device effectively implementing MITM attack.

One way to mitigate such attack is to use so-called certificate pinning [73]. The idea is that rather than relying on the certificate chain, the client application is programmed to trust only the specific certificate or only certificates signed by the specific certificate.

6. SERVICES INTEGRATION

This chapter focuses on the services implementation best practices in the context of integration within a retail enterprise. One of the main goals of ARTS is to facilitate the integration among different software sub-systems through the use of standards.

Different aspects of integration are especially important in the context of SOA and cloud computing. It is a fairly common approach that complex cloud systems are built as a large set of fairly granular services the work together to facilitate business processes. Even though basic communication and integration patterns essentially stay the same as in a traditional on premises retail enterprise the underlying architecture and the implementation might be quite different. The book by Gregor Hohpe and Bobby Woolf [74] provides a description of different enterprise integration patterns. Gregor Hohpe also has a web site dedicated to patterns and best practices for enterprise integration [75].

There are four major approaches for integrating applications.

1. Remote Procedure Call is an application integration pattern where one application executes some kind of call against an interface exposed by another application.

2. Asynchronous Messaging is a fire-and-forget style messaging that typically involves some kind of middleware or a broker like service bus or queueing sub-system.

3. Shared Database is a data integration pattern that utilizes a single storage system that is shared among multiple applications.

4. Bulk Data transfer is a data integration approach that deals with transfer of large amounts of data typically using files. This integration pattern is often used for batch data loads.

This chapter discusses integration patterns in the context of the retail enterprise. It is important to note that the examples in this chapter are used for illustrative purposes and there are multiple ways to architecture distributed retail systems.

The underlying assumption is that system components are located either in a retail store or in the cloud.

6.1 Remote Procedure Call

Remote Procedure Call (RPC) is the most common mechanism for application integration. Services typically use this pattern to expose their capabilities over API. Normally with the RPC approach the client application does not have to deal with raw data directly. The enterprise application patterns website [75] describes this pattern in the following way:

Problem	How can I integrate multiple applications so that they work together and can exchange information?
Solution	Develop each application as a large-scale object or component with encapsulated data. Provide an interface to allow other applications to interact with the running application.

Old platform specific implementations of the RPC approach include CORBA, DCOM, .NET Remoting, etc. Then RPC technologies like XML-RPC and SOAP allowed crossing platform

boundaries. RESTful APIs is the most popular RPC approach that is used in modern distributed systems.

The RPC is used for information lookup, data updates, invoking certain tasks, etc. Standards are extremely important for successful integration using the RPC pattern.

The RPC approach typically uses synchronous communications. It also implies temporal coupling between the service consumers and the service provider. The service should at least be online to be able to process requests from the clients.

If the service is deployed in the cloud an offline situation is much more likely to occur. Cloud deployment also introduces additional concerns like latency, security, scalability, etc. It is important to consider implementing retry logic (4.5.3) to deal with transient failures.

6.1.1 RPC inside Retail Store

One RPC example inside a retail store is invocation of Payment service from POS. ARTS Payment Integration White Paper [76] suggests that the payment system should be distinct from the selling system, which implies a remote call.

**Figure 6.1 Isolated Architectural Model**

In this model a payment subsystem may be present on the same computer as the selling system, and the payment system manages all interactions with the payment peripheral devices. This model isolates the payments capability into a smaller system component, thereby reducing the PCI-DSS audit envelope. Using this model, it is possible to ensure that the selling system never stores, processes or transmits cardholder data and therefore isolates the selling system from the PCI-DSS envelope. Also, the payment system can perform additional useful functions like, for example, tokenization. Co-location of the selling system and the payment system on the same computer significantly reduces the chance of the offline situation.

The figure above shows that communications between the selling system and the payment system are implemented using standard EPASOrg [77] Retailer Protocol. It also shows that the payment system communicates with devices. Typically communication with devices is

done using NRF-ARTS Unified POS (UPOS) standard. The selling system also uses UPOS to communicate with other POS devices like printers, scanners, etc.

Of course, as technology evolves the amount of embedded logic in retail peripherals continues to increase and the UPOS interactions are more and more advanced. Significant changes are planned as part of UPOS 2.0 that is actively under development by NRF-ARTS. In this new approach devices can be also considered as services in the store environment. Both the selling system and the payments system would communicate with devices using the RPC approach.
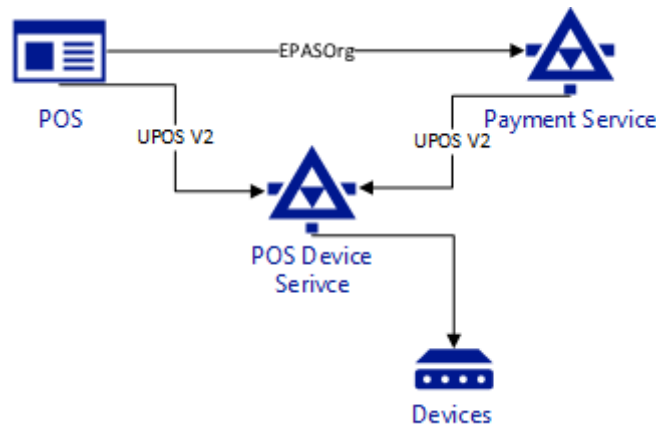
Figure 6.2 RPC inside Retail Store

The example above shows the selling system (POS) that communicates with the isolated payment system using RPC integration approach. Both systems use UPOS to communicate with the device service. It might be useful to note that consuming the device service over HTTP protocol has some limitations since it would not allow using a duplex channel and callbacks. This means that techniques like polling would have to be used to get events from devices. This is one of those examples, where using SOAP protocol still might make sense since it can work with duplex TCP communication channel.

6.1.2 RPC from Retail Store to Cloud

In the next step, to illustrate RPC-based integration from a retail store to the cloud, the POS system is decomposed into two separate components: POS Client Application and POS Service that resides in the cloud and exposes ARTS RTI API [30].

RPC communications from the retail store to the cloud have some interesting considerations. The best way to expose capabilities of edge cloud services is to use simple RESTful APIs. Such APIs are much easier to consume and they also tend to be more stable than APIs with complex data structures. Managing changes to public APIs can be a very challenging exercise.

If the POS Service runs in the cloud and powers the POS Client App then the offline scenario, when the service becomes unavailable, should be part of the consideration. One way to address this problem is to have an instance of POS Service running inside the retail store. That on-premises service would expose exactly the same interface but could have simplified and limited capabilities. When POS Client App determines that it is offline it can start using the on premises failover POS Service.

Since access to cloud services should be secured the identity management becomes very important consideration. Highly available Identity Service can be added to deal with access authorization and authentication matters. High availability for the Identity Service is very important since if the service becomes unavailable, client applications would not be able to get new access tokens or refresh expired tokens that were obtained earlier.
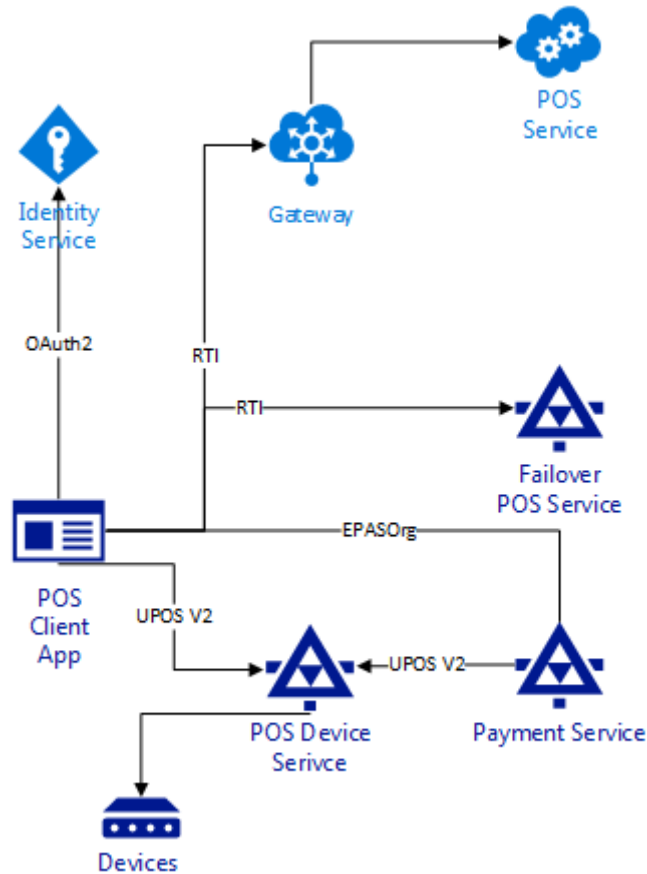
Figure 6.3 RPC from Store to Cloud

Because POS Client App is a trusted application, the OAuth2 user password credentials flow (5.2.4.4) can be used to obtain the security tokens. It means that the user can just enter the login information directly into the client application. Then user credentials are passed to the Identity Service that, after successful validation, returns a security token. Once the client application receives the security token it should immediately discard the user credentials and use the token with every request to the cloud services.

Rather than implementing security token validation logic inside every cloud edge service, it makes sense to create a special Gateway (4.5.4) component that is established as a reverse proxy frontend between the client application and the cloud service. In addition to dealing with security matters, the Gateway can perform other useful functions like monitoring, routing of request, smart load balancing, etc.

Using centralized Identity Service to manage authentication and access authorization makes a lot of sense, but if the Identity Service becomes unavailable it creates a real problem. It is an

unlikely scenario since the Identity service should be designed to have much higher availability than any business service, including POS. If it is absolutely necessary to perform sales in the offline situation, authentication using POS Client App might become the only alternative. To perform such authentication users' credentials have to be located on the client computer. In this case only salted hashes of user passwords should be used. It is a good idea to create an access token signed by the client application so that the same token validation code could be used by the service Gateway.

6.1.3 RPC from Cloud to Cloud

In the next step, to illustrate RPC-based integration between services in the cloud, the POS service is further decomposed into several more granular services that work together to provide necessary POS capabilities.

Typical brick-and-mortar POS system first creates a transient customer order that only exists in computer memory, and then immediately fulfills that order by accepting the payment from the customer. The settlement record of this exchange of merchandise (and/or services) for tender represents a retail transaction.

Therefore, as items are scanned at the POS, the system adds them to a transient customer order that is later settled and a retail transaction is generated. This is an important nuance to understand the decomposition of the POS service.

POS Service becomes a composition controller that coordinates POS activity. It is composed of four other services. The Order Capture service is a process service that keeps track of the whole process of how a transient order is created and modified. The Order Capture service calls the Price Calculation and the Tax Calculation services to calculate prices and taxes for order items. POS Service also calls the CRM service to get customer and loyalty information.

The Tax Calculation service can use ARTS XML Transaction Tax Technical Specification [78] for its interface. Similarly, the Price Calculation service can use ARTS Pricing Service Interface Technical Specification [79]. As far as the CRM service is concerned, it could leverage XML schemas defined in seven volumes of ARTS XML Customer Technical Specification [80] that define data structures for multiple customer-related services (Customer Maintenance, Loyalty, Targeted Offers, etc.).

The RPC communications in the cloud bring forward a whole new set of considerations. Since services are consumed by other services, the interfaces they expose can be more complex. Still careful design of the interfaces is extremely important.

If services are located in the same data center, the latency can be significantly lower than making calls from a service on premises. That is why it is a much better approach to have a single coarse granularity call from the retail store that is fanned into multiple more granular calls to services in the cloud.
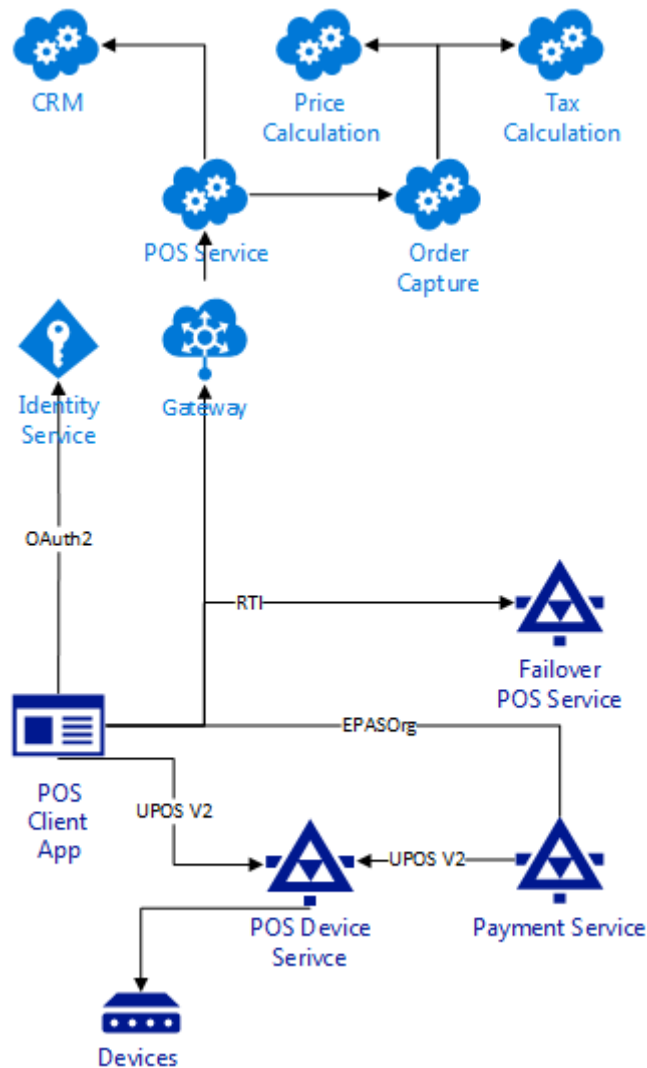
Figure 6.4 RPC from Cloud to Cloud

Inter-service communications in the cloud are inherently more secure and can have multiple layers of protection. Often middle tier services cannot be directly accessed from public internet. A trusted subsystem approach, in which the middle tier services trust the identity of edge services that call them rather than the identity of every caller, is very common in the cloud. The OAuth2 client credentials flow (5.2.4.4) can be used for access authorization between two services. In the example above, since the gateway already performed the necessary authentication and authorization it can use its own credentials to access the POS Service.

6.1.4 RPC from Cloud to Premises

Sometimes it might be necessary for a SaaS system running in the cloud to communicate with invoke a service that is deployed on premises behind a firewall. Making such remote call can be quite tricky but in some situations this might be the only possible solution. For example, such an approach can be used to provide integration with a legacy system that cannot be deployed in the cloud but still offers some valuable capabilities. It also can be used to access data that, because of business policies or regulatory requirements, should be kept on premises.

For example, a retailer could have a policy that all customer credit cards numbers should be tokenized and stored on premises. The payment authorization system uses credit cards to authorize and perform payments but only credit card tokens are passed back to the caller. To achieve this functionality payment authorization communicates with the tokenization service and exchanges credit card numbers for special tokens. In this scenario the eCommerce site that communicates with Order Capture service using ARTS SSOI (Self-Service Order Interface) standard [81] uses the Payment Authorization service in the cloud. The Payment Authorization service communicates with the Tokenization service in a corporate data center to obtain tokens for credit card numbers. The same Tokenization service can be used by payment systems in stores.
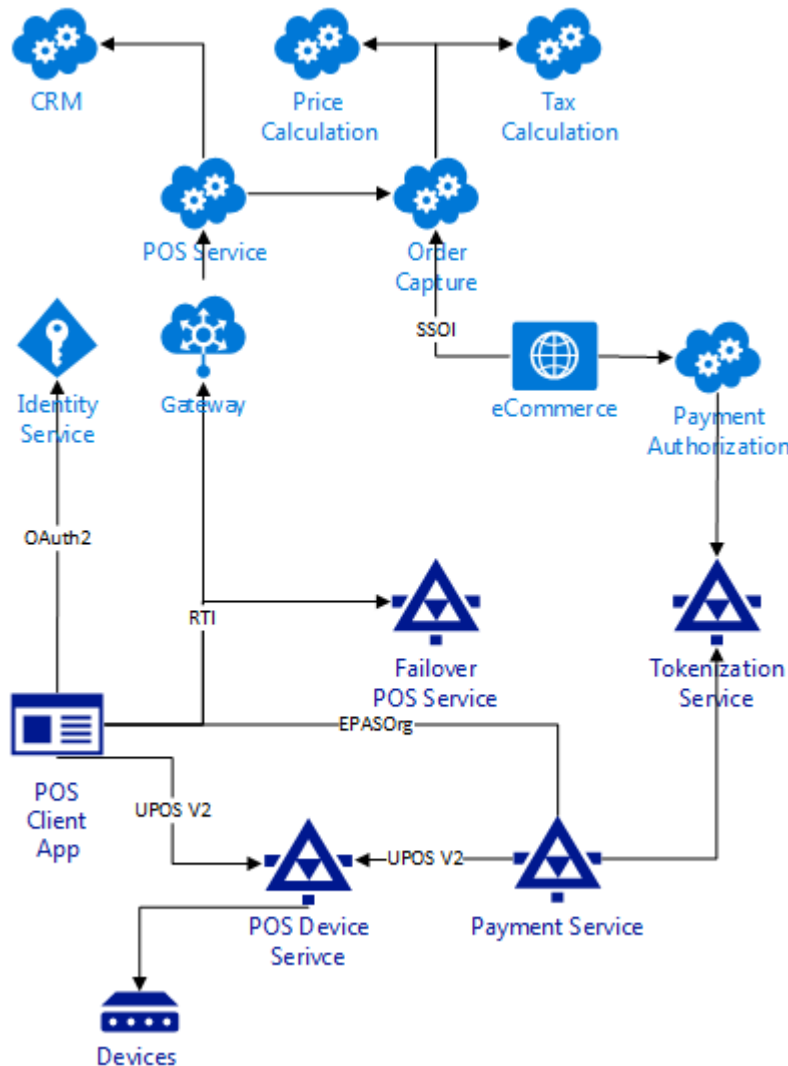
Figure 6.5 RPC from Cloud to Premises

The RPC communications from the cloud to an on-premises system bring forward a whole new set of considerations. Since such communications perform a round trip from cloud to on-premises data center they may experience some latency. Also, the retailer's data center has limited scalability and the load on the service has to be carefully tested since it might become a

performance bottleneck. Since typically a cloud based system can be fairly quickly scaled out it can overwhelm on-premises service that might not be able to keep up with the load.

Allowing access to your on-premises service from the cloud creates important security concerns.

One option to implement such approach is to use a VPN. Most cloud providers offer VPN connectivity between networks in the cloud and a private corporate network. In such scenario a service running in the cloud can securely communicate with a service on premises. Unfortunately, this option might not always be available.

Another approach is to implement an internet-facing proxy inside the DMZ. Cloud services can only have direct access to the proxy which will securely communicate with the service in the corporate data center. This approach introduces additional components that makes the solution more complex and also adds some latency.

The third option is to place a proxy in the cloud and have a special gateway on premises that would establish an outbound bidirectional connection through the firewall to that proxy. The cloud services would communicate with the cloud proxy that would forward the request to the on-premises gateway. This approach is interesting but it is not trivial to implement.

6.2 Asynchronous Messaging

Asynchronous messaging is a common integration mechanism for building loosely coupled distributed systems. Typically this pattern is used with some kind of intermediary (broker, message bus, queueing, etc.). Due to its asynchronous nature this pattern provides loose coupling between the communicating subsystems. It is often used to implement publish-subscribe message exchange pattern (MEP) and event-driven designs. The enterprise application patterns website [75] describes this pattern in the following way:

Problem	How can I integrate multiple applications so that they work together and can exchange information?
Solution	Use Messaging to transfer packets of data frequently, immediately, reliably, and asynchronously, using customizable formats.

When a client sends an asynchronous message to a service it typically uses some kind of messaging infrastructure that is highly available. Therefore, the probability of error when using asynchronous messaging is much less than that of RPC. Also, since messages can be consumed by multiple instances of the service this integration pattern works very well with horizontal scalability.

The messaging intermediary can perform a number of useful functions like persisting messages in durable storage, logging, filtering, content-based routing, transformations, etc.

Even though asynchronous messaging systems can provide a guaranty that messages will be delivered using some kind of store and forward approach but they typically cannot provide any latency guaranties. As a rule asynchronous communications are slower than RPC and application might have to deal with eventual consistency.

Idempotence (4.5.1) is often necessary so that the message could be safely resubmitted. This is important for guaranteed delivery and simplifies the programming model.

6.2.1 Messaging inside Retail Store

Asynchronous messaging can be effectively used for communications between different subsystems inside a retail store. For example, certain events that happen at a register can be fired so that interested subsystems in the store could then consume them. ARTS has a special standard Notification Event Architecture in Retail (NEAR) [82] that describes the envelope for such messages.

Another common example of in-store messaging is forwarding retail transactions, typically represented as POSLog [83], from registers to the store server. A special transaction processing service that runs in the back-office can receive POSLog messages and perform some useful functions like updating inventory or saving data for operational reporting inside a database that uses ARTS Operational Data Model [36].

Figure 6.6 In-Store Messaging

Asynchronous messaging is a good option to move transactions inside the store since it can provide guaranteed delivery of messages to the in-store server. Even if the transaction

processing service is temporarily unavailable messages can be stored and forwarded later when the service is back online. That is a key feature of asynchronous messaging.

6.2.2 Messaging from Retail Store to Cloud

Since asynchronous messaging reduces coupling between interacting services, it is a good option for store-to-cloud communications.

For example, it could be used for sending POSLog messages containing retail transactions to the sales audit service located in the cloud. Also, POSLog messages that contain customer orders placed at the store register can be sent to the centralized Order Management System (OMS).

Figure 6.7 Messaging from Retail Store to Cloud

If a service in the cloud can be offline, durability of messages becomes important. Data security is another critical consideration since messages are sent from a corporate network to a public cloud. It might be necessary to use message level security in this situation.

It is also very important to consider the capabilities of the messaging system such as message broker location, security options, guaranteed delivery, etc.

6.2.3 Messaging from Cloud to Cloud

Since asynchronous messaging provides loose coupling between senders and receivers it is one of the preferred methods of communications in the cloud. This pattern is especially well suited for backend services that do not participate in interactions with the user.

Similar to the in-store example, eCommerce site could use asynchronous messages to send POSLog customer orders to the OMS. There are plenty of other messages types that services use to exchange information in the cloud. Often asynchronous messages are used to implement Event-Driven Architecture (EDA) [84]. In this approach a service publishes an event about relevant data changes and other services consume events they need for their functionality. For example, the merchandising system could publish an event that a new item was added and the inventory service could consume the event and perform necessary item initialization logic. ARTS has the Item Maintenance XML schema that can be used to add an item [85]. The event is shown as a New Item message on the figure below.

Figure 6.8 Messaging from Cloud to Cloud

Security of the messages is still an important concern even for cloud to cloud messaging. Also, since communications happen completely in the cloud the ability to monitor the flow of data becomes crucial.

6.2.4 Messaging from Cloud to Premises

It is also possible to use asynchronous messaging for communications from the cloud to a retail store or a corporate data center.

This scenario is often used for near real-time synchronization of data from cloud services to on-premises systems. Typically this type of synchronization is used for non-bulk data since there are performance and scalability restrictions. Also, many cloud providers impose charges for any outbound traffic from the cloud.

Another concern with communicating to on-premises systems from the cloud is scalability. Internal resources are more difficult to scale than cloud resources. Therefore it is important to make sure that services running on premises can keep up with the load from the cloud.

One example of using messaging from the cloud to a retail store could be OMS sending reservation directive message to the in-store inventory service.
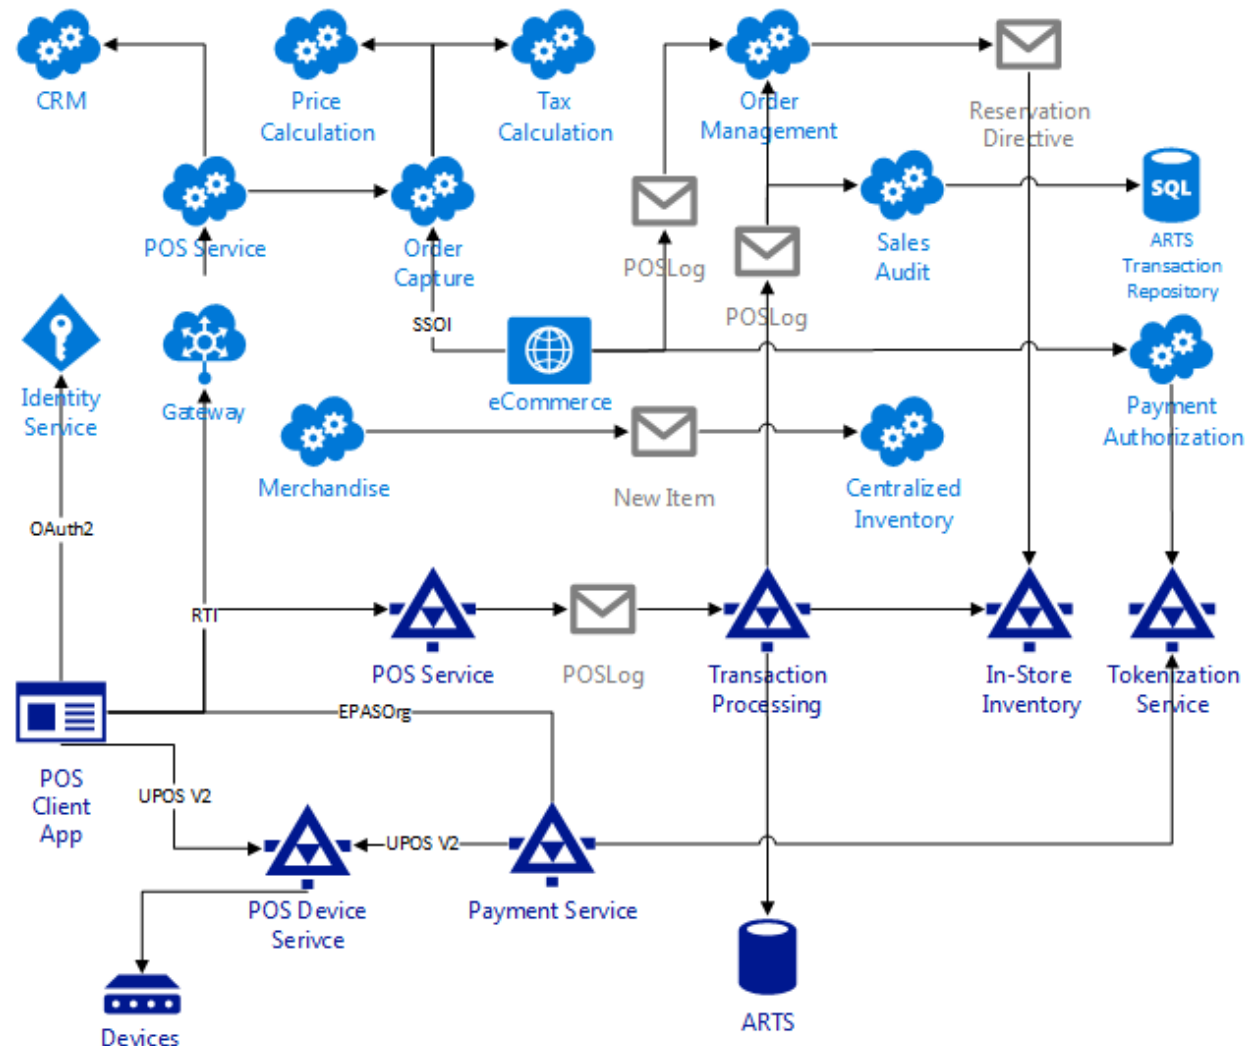

Figure 6.9 Messaging from Cloud to Premises

Every time corporate resources are accessed from a public cloud security becomes a major concern. Such endpoints should be properly protected. One approach is to use a VPN and establish a secure connection between the on-premises network and the virtual network segment in the cloud. It is also possible to use an internet-facing proxy. If a message broker is part of the design then it can be placed in the cloud and a special on-premises agent could poll it periodically.

Some messaging from the cloud can take advantage of mobile phones communications infrastructure. For example, an in-cloud notification system can broadcast messages to mobile phones used by associates.

6.3 Common Data Store (Shared Database)

Common data store is still a fairly common integration mechanism even though it couples multiple communicating services to the underlying database schema. For this reason, it is often used to exchange information between somewhat "related" services when changes of the database structures can be coordinated. This integration pattern typically provides a very high degree of data consistency. As soon as one service commits data modifications they become available for consumption by another service. The enterprise application patterns website [75] describes this pattern in the following way:

Problem	How can I integrate multiple applications so that they work together and can exchange information?
Solution	Integrate applications by having them store their data in a single Shared Database.

One of the biggest challenges with using the shared database integration pattern is the dependency on the common data format. ARTS standards provide tremendous help to solution architects in establishing canonical data structures. This is true for both relational data models as well as the structure of XML/JSON documents that can be stored in a shared NoSQL database.

When a lot of services use the same database simultaneously it can become a performance bottleneck. Databases typically use locking to guarantee consistency of the data. So, when multiple applications attempt to access the information they might need to wait until the locks are released.

Database scalability is another concern especially for relational databases. Scaling data servers up by using a more powerful machine is not an optimal option for the cloud. Some data engines support sharding [86] but such approaches usually result in more complex solutions. Another option is to have multiple database servers that contain the same data but then maintaining data consistency may become an issue. CAP theorem [87] states that it is impossible for a distributed database system to simultaneously provide all three of the following guarantees:

- *Consistency* (all nodes see the same data at the same time)
- *Availability* (every request received a response: success or failure)
- *Partition Tolerance* (the system can operate despite of network partitioning)

Since cloud systems are designed to operate when network might get partitioned distributed databases have to make a trade-off between availability and consistency. Many NoSQL database sacrifice consistency so that they can be highly available and survive network failures. Such databases might not be strictly consistent at any particular moment but they will eventually become consistent when all the changes are propagated to all the nodes.

6.3.1 Shared Database in Retail Store

It is a fairly common scenario when databases inside a retail store are shared among multiple applications to exchange the information. For example, in-store transaction processing service can store retail transactions data inside a store database that can be used to perform operational reporting.
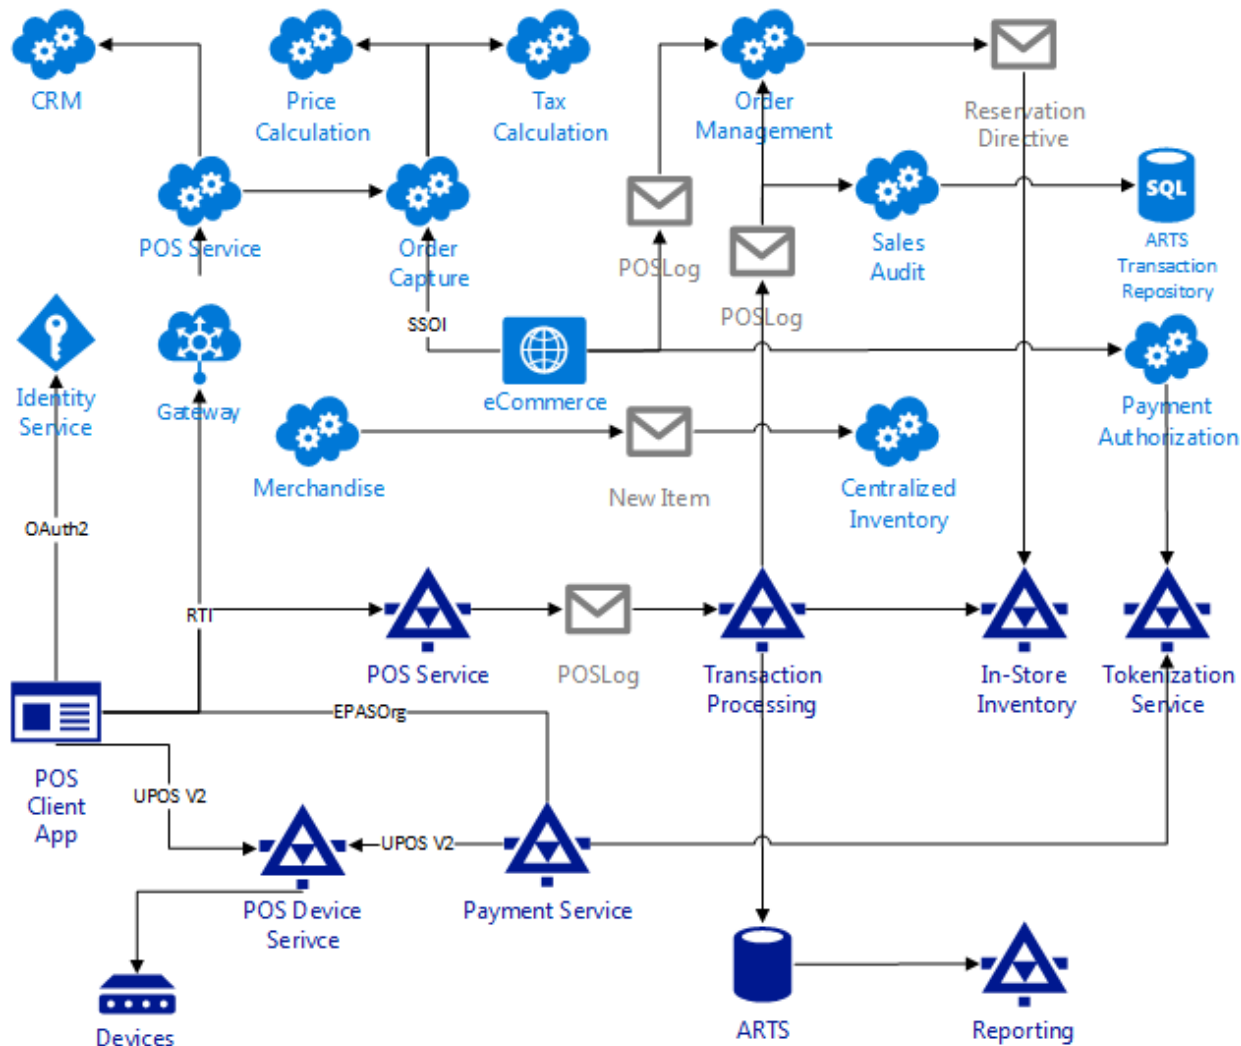
Figure 6.10 Shared Database inside Retail Store

Using ARTS operational data model simplifies integration based on common data store. Another example of this approach inside store could be updates to the PLU table that are made by the price maintenance application and picked up by the POS when items are scanned.

It is not a trivial task to use a shared database on-premises to integrate with services in the cloud. Exposing a database connection to public cloud is a significant security risk. There are some solutions that allow establishing a secure database connection from the cloud to an on-premises database but it is not a common approach. Using a VPN to extend an on-premises network could be another option. It might be useful to note that if the data has to be stored on premises to expose it to the cloud it is a much more common approach to wrap the database in some kind of a service.
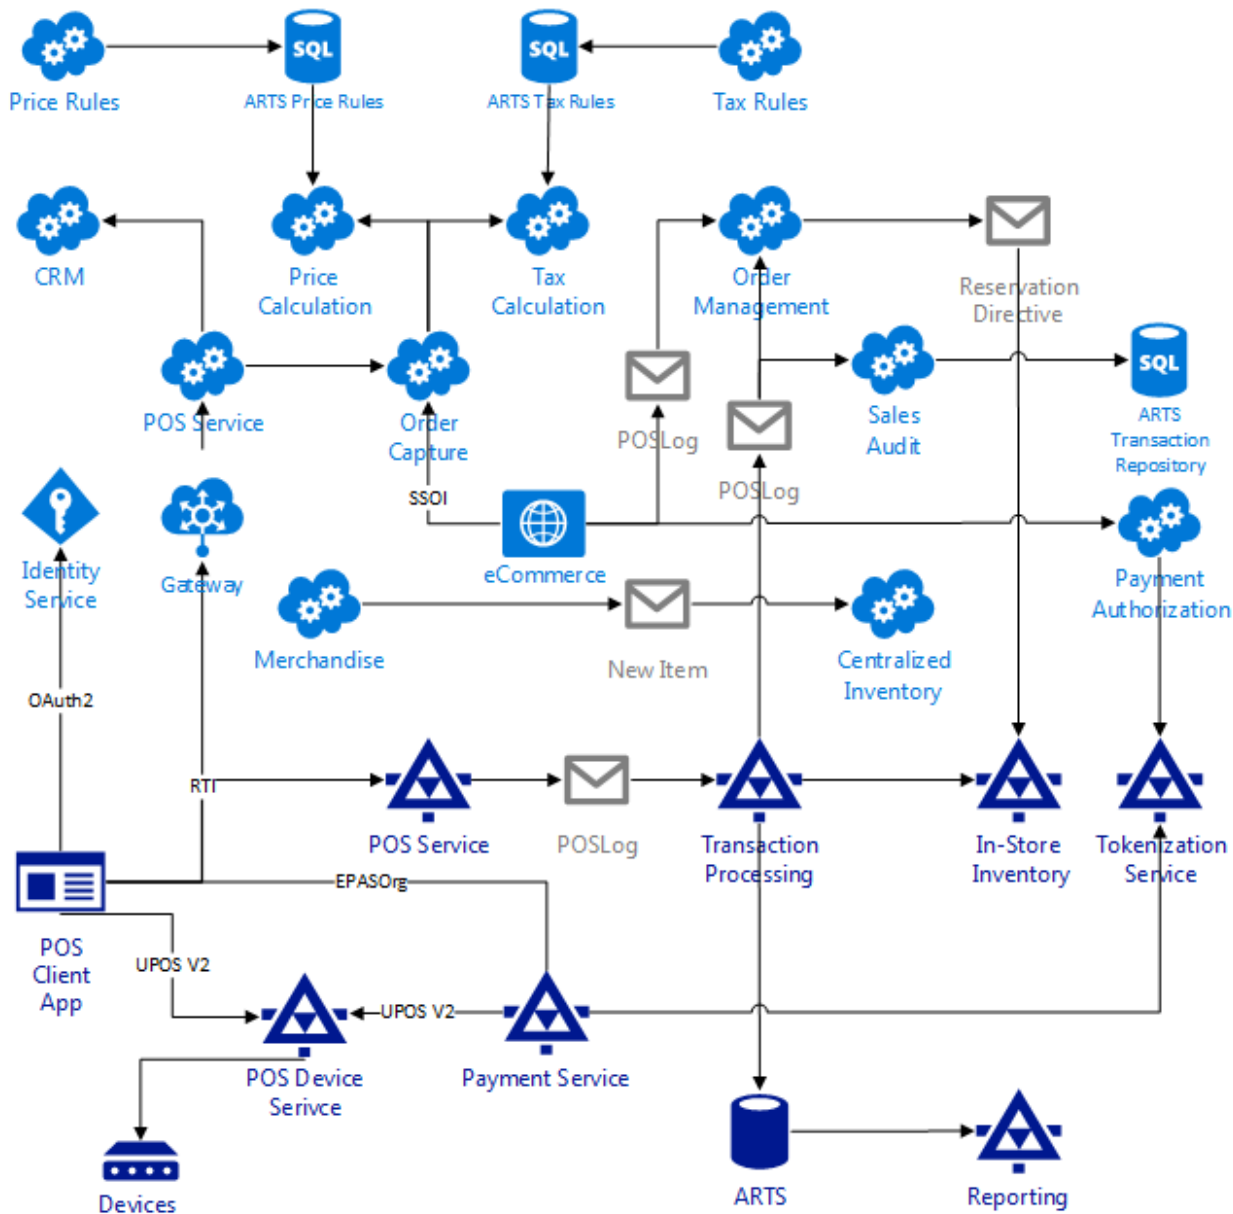
6.3.2 Shared Database in Cloud

Using a shared database in the cloud could be a very sensible option. If the database size increases rapidly and it is difficult to estimate and provision storage resources on premises then a cloud database could be a good choice. Also, if the services that use a shared database are deployed in the cloud then the database should be deployed in the cloud as well. It reduces the latency, minimizes costs, and improves reliability and security of the solution.

For example, the tax calculation service can load data from a shared tax rules database that is maintained by the tax rules service. Similarly, the price calculation service can use price derivation rules loaded from a database that it is shared with the price rules service. Since ARTS data model defines entities that support both tax rules and price derivation rules, it can be used to establish standard relational structures.

Figure 6.11 Shared Database in Cloud

A database in the cloud can be accessed not only from the cloud but also directly from premises. This creates and interesting integration option between cloud and on-premises services. Such an approach has to be able to tolerate latency and even potential network partitioning. Many cloud database services expose HTTP based RESTful APIs to access the data. In such a scenario security becomes an important consideration. The database either exposes a public data access endpoint or all database communications are performed over the VPN. Providing a direct access to a database using the public HTTP endpoint can result in security vulnerabilities. If there is a need to expose data for public consumption, it is a better practice to access it indirectly via entity services.

6.4 Bulk Data Synchronization (File Transfer)

File transfer is a fairly old data synchronization pattern that has been extensively used in retail. It was very common to perform nightly downloads of PLU files. This pattern can be used when latency is acceptable and there is no need for real-time data consistency between the source and the destination. This pattern is typically used for significant volumes of data when approaches like asynchronous messaging become impractical. The enterprise application patterns website [75] describes this pattern in the following way:

Problem	How can I integrate multiple applications so that they work together and can exchange information?
Solution	Have each application produce files containing information that other applications need to consume. Integrators take the responsibility of transforming files into different formats. Produce the files at regular intervals according to the nature of the business.

Since business analytics works with large volumes of data that are collected over extended periods of time the file transfer pattern is often used to move data from operational data stores to a data warehouse. The data files participate in ETL (Extract, Transform, and Load) process [88].

One of the considerations with using this pattern is the significant amount of time it might take to process the input file. It is especially true if the data load includes a lot of cleansing and business logic.

This type of integration takes data from one system, puts it into a file and then loads it into another system, effectively creating copies of business data in the files. Therefore security and privacy concerns become very important if sensitive data is shared in this manner.

When moving large volumes of data between systems in the cloud and on premises, bandwidth constraints can become a limiting factor. Also, data transfer from a data center in the cloud to on premises systems or systems running in another cloud data center typically incur charges.

6.4.1 File Transfer inside Retail Store

File transfer inside retail store can be used to move updates from the in-store server to registers. This would typically include PLU data, employee file, etc. Sometimes transaction archives can be moved from registers to some centralized location.

Since transferring large volumes of data can have negative impact on normal business operations it is a good idea to perform register updates when the business activity is low.

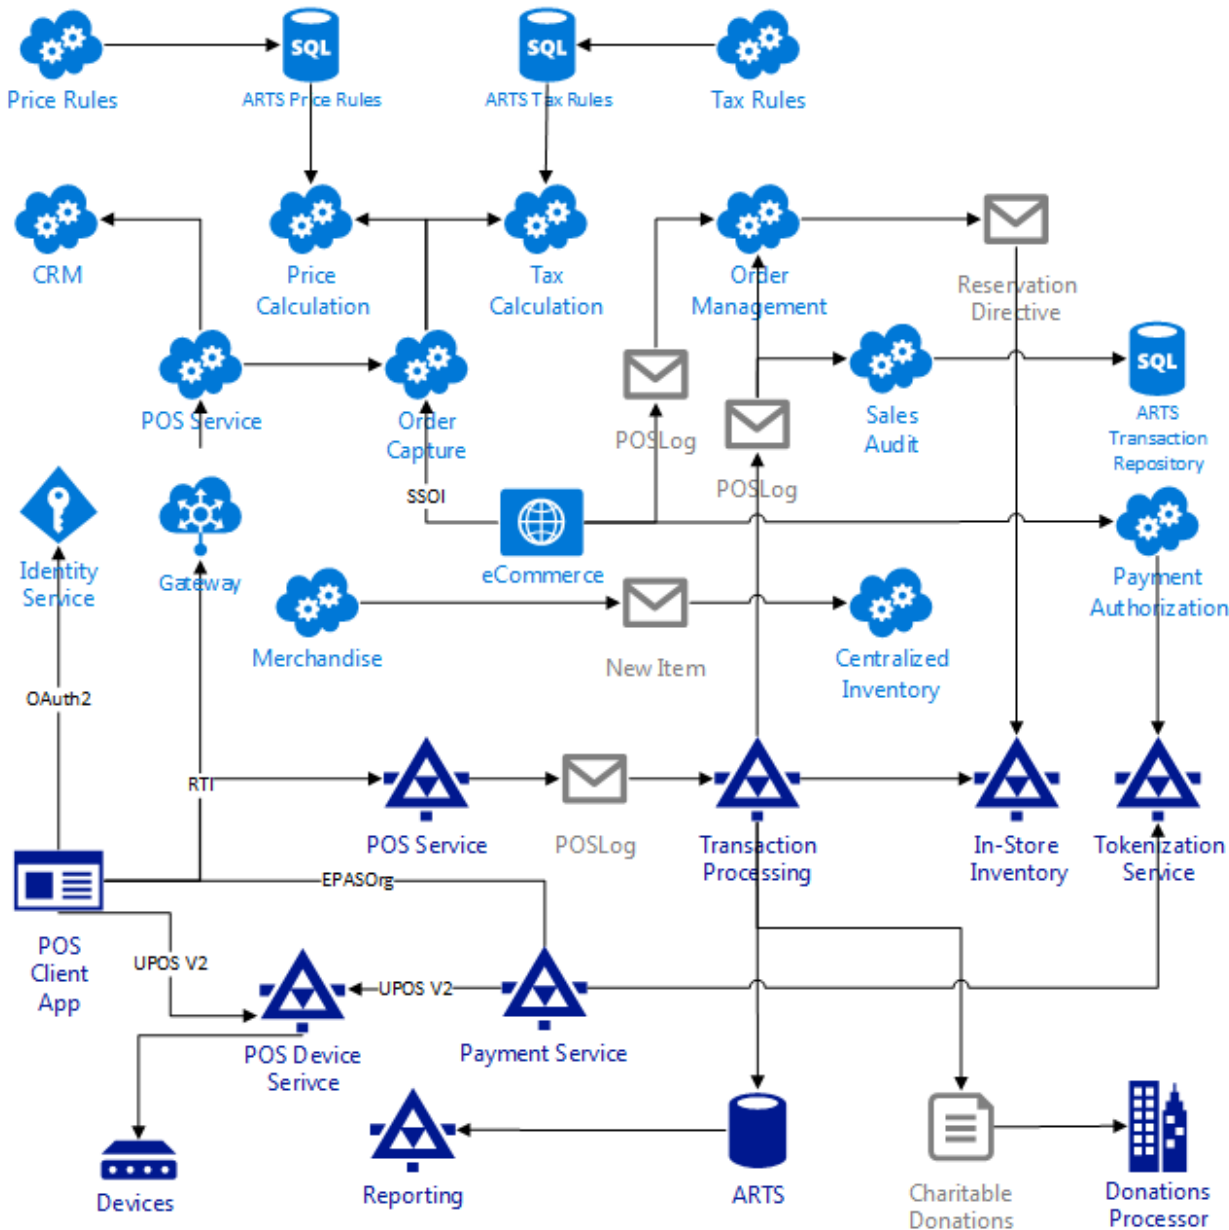### 6.4.2 File Transfer from Retail Store to Cloud

Most of the large volume data generated inside a retail store comes from sales activity at the registers. Modern retail systems often use asynchronous messaging to communicate this data to a centralized location in a near real-time fashion. It is done to have a more accurate picture of inventory counts and facilitate some other business processes that might need access to sales transaction data.

Still it is not uncommon that sales data is transferred from stores using files that contain all the records of sales for the business day. Sometimes it can be done as a part of the recovery

process. ARTS POSLog standard [83] supports a concept of batch to facilitate such functionality.

Another example of moving a batch data transfer out of a retail store is sending charity contributions information to a donation processor, which could be sent using asynchronous messaging. However, in this scenario typically latency is acceptable and periodical file transfer can be a reasonable approach instead. ARTS Change4Charity standard [89] supports sending multiple records of charitable contributions as a single XML file.

Figure 6.12 File Transfer from Store

In addition to latency, file transfer out of retail store has to properly address security and privacy concerns.

6.4.3 File Transfer from Cloud to Cloud

File transfer is typically used for bulk data exchange in the cloud. It is a good idea to implement file based integration in an asynchronous manner. One approach is to combine file transfer with asynchronous messaging. For example, the sales audit service can generate a file containing validated sales data for a business date and put it in some kind of storage system. Then it can publish an asynchronous message that this file is available at a certain URI and any services that might need this information could get a copy of the file.

Another example of the file transfer integration in the cloud is sending sales data collected from all the stores to a data warehouse. It might be useful to note that ARTS developed a reference data warehouse model [90] that is a dimensional extension to the ARTS operational data model. The two models are tightly integrated through a set of ETL queries and views. POSLog batch can also be used to move sales data between an operational database and a data warehouse.

Figure 6.13 File Transfer in the Cloud

One of considerations in this scenario is throughput. Another important concern is security and privacy especially if the files that are used to transfer the data are placed into some kind of staging area for extended periods of time.

6.4.4 File Transfer from Cloud to Premises

It is a very common approach for retail stores to receive data updates from enterprise systems in a file format. Also retailers might have investment in some sophisticated on-premises systems like, for example, sales analytics. Thus, if the sales data repository is located in the cloud it should also periodically feed the on-premises analytics system.

Probably the most common example of file transfer into a retail store is item updates. ARTS has Item Maintenance XML schema [85] that can be used to update a set of items, delete another set of items and add some new items in a single batch XML instance.
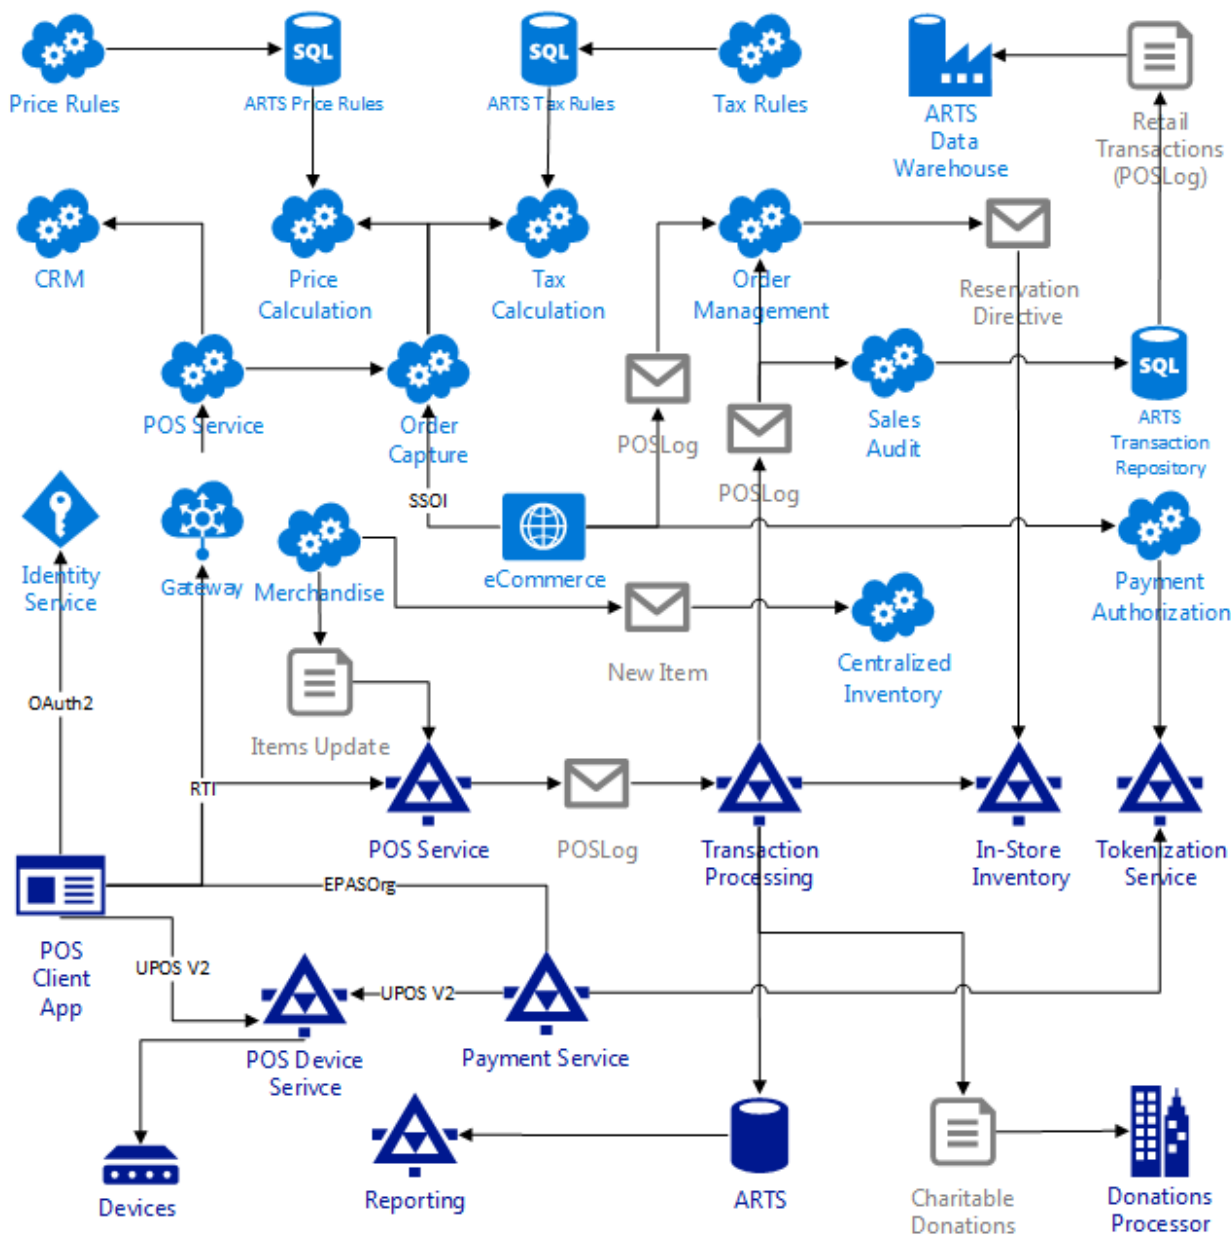
Figure 6.14 File Transfer to Retail Store

When moving data from the cloud to a retail store it is important to consider costs since cloud providers typically impose a charge for the outbound traffic. Also, since file transfer is typically used for large volumes of data bandwidth constraints would be another consideration.

Security is always a concern when data ends up on internal systems. Using a VPN is one option. Another option could be pulling files from some known location in the cloud.

6.5 Conclusion

This chapter used ARTS standards to demonstrated different integration options on premises and in the cloud. Only a small subset of standards was used to create a fairly elaborate picture of a distributed retail enterprise. ARTS technical specifications are built around use cases that provide great insights into how the standards are used in the context of different business processes.

7. ABBREVIATIONS

ACID	Atomicity, Consistency, Isolation, Durability
AMQP	Advanced Message Queuing Protocol
API	Application Programming Interface
ARTS	Association for Retail Technology Standards
BASE	Basically Available, Soft state, Eventual consistency
CQRS	Command Query Responsibility Segregation
CRUD	Create Read Update Delete
CSP	Cloud Service Providers
CSV	Comma-separated values
DDD	Domain-Driven Design
DNS	Domain Name System
DOMS	Distributed Order Management System
DoS	Denial of Service (attack)
EAI	Enterprise Application Integration
EDA	Event-Driven Architecture
EDI	Electronic Data Interchange
ETL	Extract, Transform, Load
HATEOAS	Hypermedia As The Engine Of Application State
HTTP	Hypertext Transfer Protocol
IAM	Identity and Access Management
IPC	Inter-Process Communication
JSON	JavaScript Object Notation
JWT	JSON Web Token
LDAP	Lightweight Directory Access Protocol

MEP	Message Exchange Pattern
MFA	Multi-Factor Authentication
MITM	Man In The Middle (attack)
MOM	Message-Oriented Middleware
NIST	National Institute of Standards and Technology
REST	Representational State Transfer
RPC	Remote Procedure Call
SAML	Security Assertion Markup Language
SDA	Software-Defined Architecture
SLA	Service Level Agreement
SOA	Service-Oriented Architecture
SSL	Secure Socket Layer
SSO	Single Sign-On
STS	Security Token Service
TCP	Transmission Control Protocol
TLS	Transport Layer Security
TTL	Time-To-Live
UDDI	Universal Description Discovery Integration
UDP	User Datagram Protocol
UML	Unified Modeling Language
UPOS	Universal Point Of Service
URI	Uniform Resource Identifier
W3C	World Wide Web Consortium
WADL	Web Application Description Language
WSDL	Web Services Description Language

8. REFERENCES

[1] ARTS, "SOA Best Practices Technical Report," 2008.

[2] ARTS, "SOA Blueprint for Retail," 2008.

[3] ARTS, "Cloud Computing for Retail," 2009.

[4] W. R. Schulte and Y. V. Natis, "Service Oriented Architectures, Part 1," Gartner, 1996.

[5] W3C, "Web Services Glossary," 11 February 2004. [Online]. Available: http://www.w3.org/TR/2004/NOTE-ws-gloss-20040211/#webservice.

[6] N. Malik, "JaBoWS is the Enemy of Enterprise SOA," 17 March 2008. [Online]. Available: http://blogs.msdn.com/b/nickmalik/archive/2008/03/17/jabows-is-the-enemy-of-enterprise-soa.aspx.

[7] R. T. Fielding, "Architectural Styles and the Design of Network-based Software Architectures," University of California, Irvine, 2000.

[8] A. T. Manes, "SOA is Dead; Long Live Services," 5 January 2009. [Online]. Available: http://apsblog.burtongroup.com/2009/01/soa-is-dead-long-live-services.html.

[9] "SOA Manifesto," 23 October 2009. [Online]. Available: http://www.soa-manifesto.org/.

[10 T. Erl, "The Annotated SOA Manifesto," 22 November 2009. [Online]. Available:
]   http://www.soa-manifesto.com/annotated.htm.

[11 The Open Group, "SOA Reference Architecture," The Open Group, 2011.
]

[12 OASIS, "Reference Architecture Foundation for Service Oriented Architecture Version 1.0,"
]   4 December 2012. [Online]. Available: http://docs.oasis-open.org/soa-rm/soa-ra/v1.0/cs01/soa-ra-v1.0-cs01.pdf.

[13 The Open Group, "Service-Oriented Architecture Ontology Version 2.0," The Open Group,
]   2014.

[14 J. Lewis and M. Fowler, "Microservices," 25 March 2014. [Online]. Available:
]   http://martinfowler.com/articles/microservices.html.

[15 S. Jones, "Microservices is SOA, for those who know what SOA is," 18 March 2014.
]   [Online]. Available: http://service-architecture.blogspot.co.uk/2014/03/microservices-is-soa-for-those-who-know.html.

[16 N. Kant and S. Tonse, "Karyon: The nucleus of Composable Web Service," 6 March 2013.
]   [Online]. Available: http://techblog.netflix.com/2013/03/karyon-nucleus-of-composable-web-service.html.

[17 Y. V. Natis, "Software-Defined Architecture: Application Design for Digital Business,"
]   Gartner, 8 May 2014. [Online]. Available:

http://my.gartner.com/portal/server.pt?open=512&objID=202&mode=2&PageID=5553&ref=webinar-rss&resId=2698619.

[18 P. Mell and G. Timothy, "The NIST Definition of Cloud Computing," September 2011.
]   [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf.

[19 D. Box, "A Guide to Developing and Running Connected Systems with Indigo," *MSDN*
]   *Magazine,* January 2004.

[20 The Open Group, "Service Oriented Architecture : What Is SOA?," [Online]. Available:
]   http://www.opengroup.org/soa/source-book/soa/soa.htm.

[21 OASIS, "Reference Model for Service Oriented Architecture 1.0," 12 October 2006.
]   [Online]. Available: http://docs.oasis-open.org/soa-rm/v1.0/soa-rm.pdf.

[22 T. Erl, SOA: Principles of Service Design, Prentice Hall, 2007.
]

[23 G. Prasad, "SOA As Dependency-Oriented Thinking - One Diagram That Explains It All,"
]   20 May 2013. [Online]. Available: http://wisdomofganesh.blogspot.co.uk/2013/05/soa-as-
     dependency-oriented-thinking-one.html.

[24 G. Prasad, "Dependency-Oriented Thinking: Volume 1 - Analysis and Design," December
]   2013. [Online]. Available: http://www.infoq.com/minibooks/dependency-oriented-thinking-
     1.

[25 G. Prasad, "Dependency-Oriented Thinking: Volume 2 - Governance and Management,"
]   December 2013. [Online]. Available: http://www.infoq.com/minibooks/dependency-
     oriented-thinking-2.

[26 W3C, "SOAP Version 1.2 Part 1: Messaging Framework (Second Edition)," 27 April 2007.
]   [Online]. Available: http://www.w3.org/TR/soap12/.

[27 W3C, "Web Services Description Language (WSDL) Version 2.0 Part 0: Primer," June 26
]   2007. [Online]. Available: http://www.w3.org/TR/wsdl20-primer/.

[28 W. Vambenepe, "Amazon proves that REST doesn't matter for Cloud APIs," 5 December
]   2010. [Online]. Available: http://stage.vambenepe.com/archives/1700.

[29 M. Fowler, "Richardson Maturity Model," 18 March 2010. [Online]. Available:
]   http://martinfowler.com/articles/richardsonMaturityModel.html.

[30 ARTS, "ARTS XML Retail Transaction Interface Technical Specification Version 1.0," 2007.
]

[31 OASIS, "AMQP (Advanced Message Queuing Protocol)," [Online]. Available:
]   http://www.amqp.org/.

[32 Wikipedia, "ISO 8601," [Online]. Available: http://en.wikipedia.org/wiki/ISO_8601.
]

[33 Object Management Group, "OMG Unified Modeling Language (OMG UML)
] Superstructure," 6 August 2011. [Online]. Available:
http://www.omg.org/spec/UML/2.4.1/Superstructure/PDF/.

[34 D. Carlson, Modeling XML Applications with UML, Addison-Wesley Professional, 2001.
]

[35 Wikipedia, "Enterprise Architect (software)," [Online]. Available:
] https://en.wikipedia.org/wiki/Enterprise_Architect_(software). [Accessed August 2015].

[36 ARTS, "ARTS Operational Data Model V7.0," 2014.
]

[37 WS-I, "Basic Profile Version 2.0," 9 November 2010. [Online]. Available: http://ws-
] i.org/Profiles/BasicProfile-2.0-2010-11-09.html#Response_Wrappers.

[38 IETF, "Hypertext Transfer Protocol (HTTP/1.1): Semantics and Content RFC 7231," June
] 2014. [Online]. Available: https://tools.ietf.org/html/rfc7231.

[39 E. Evans, Domain-Driven Design: Tackling Complexity in the Heart of Software, Prentice
] Hall, 2003.

[40 API Blueprint, "API Blueprint," 2015. [Online]. Available: https://apiblueprint.org/.
]

[41 RAML Workgroup, "RESTful API Modeling Language," 2015. [Online]. Available:
] http://raml.org/index.html.

[42 Swagger, "Swagger. A Poweful Interface to Your API," 2015. [Online]. Available:
] http://swagger.io/.

[43 O. Ben-Kiki, C. Evans and I. döt Net, "YAML Ain't Markup Language (YAML™) Version
] 1.2," 1 October 2009. [Online]. Available: http://www.yaml.org/spec/1.2/spec.html.

[44 Open Project, "The RESTful API Modeling Language (RAML) Spec," 2015. [Online].
] Available: https://github.com/raml-org/raml-spec.

[45 MuleSoft, "API Designer," 2015. [Online]. Available:
] https://www.mulesoft.com/platform/api/api-designer.

[46 R. Hastings, "Netflix Culture: Freedom & Responsibility," 1 August 2009. [Online]. Available:
] http://www.slideshare.net/reed2001/culture-1798664.

[47 Wikipedia, "Conway's law," 2015. [Online]. Available:
] http://en.wikipedia.org/wiki/Conway%27s_law.

[48 M. Fowler, "MonolithFirst," 3 May 2015. [Online]. Available:
] http://martinfowler.com/bliki/MonolithFirst.html.

[49 ARTS, "ARTS Operations Guide Version 1.0.0," 2013.
]

[50 T. Preston-Werner, "Semantic Versioning 2.0.0," [Online]. Available: http://semver.org/.
]

[51 OASIS, "Web Services Dynamic Discovery (WS-Discovery) Version 1.1," July 2009.
]   [Online].

[52 IETF, "DNS-Based Service Discovery RFC 6763," February 2013. [Online]. Available:
]   http://www.ietf.org/rfc/rfc6763.txt.

[53 OASIS, "OASIS UDDI Specification TC," February 2005. [Online]. Available:
]   https://www.oasis-open.org/committees/uddi-spec/.

[54 Wikipedia, "Exponential backoff," 2015. [Online]. Available:
]   https://en.wikipedia.org/wiki/Exponential_backoff.

[55 Wikipedia, "Circuit breaker design pattern," 2015. [Online]. Available:
]   https://en.wikipedia.org/wiki/Circuit_breaker_design_pattern.

[56 ARTS, "ARTS Security Technical Report V 1.0," 2015.
]

[57 IETF, "The Transport Layer Security (TLS) Protocol Version 1.2 RFC 5246," August 2008.
]   [Online]. Available: http://tools.ietf.org/html/rfc5246.

[58 WS-I, "Basic Security Profile Version 1.0," 30 March 2007. [Online]. Available:
]   http://www.ws-i.org/Profiles/BasicSecurityProfile-1.0.html.

[59 H. Liu, S. Pallickara and G. Fox, "Performance of Web Services Security," [Online].
]   Available: http://grids.ucs.indiana.edu/ptliupages/publications/WSSPerf.pdf.

[60 T. Erl, "SOA Design Patterns. Trusted Subsystem.," [Online]. Available:
]   http://soapatterns.org/design_patterns/trusted_subsystem.

[61 Wikipedia, "Security token service," https://en.wikipedia.org/wiki/Security_token_service,
]   2015.

[62 Wikipedia, "Dictionary attack," July 2015. [Online]. Available:
]   https://en.wikipedia.org/wiki/Dictionary_attack.

[63 IETF, "Internet Security Glossary, Version 2 RFC 4949," [Online]. Available:
]   http://tools.ietf.org/html/rfc4949.

[64 OAuth, "OAuth 2.0," October 2012. [Online]. Available: http://oauth.net/2/.
]

[65 OASIS, "Security Assertion Markup Language (SAML) V2.0 Technical Overview," 25 March
]   2008. [Online]. Available: https://www.oasis-
    open.org/committees/download.php/27819/sstc-saml-tech-overview-2.0-cd-02.pdf.

[66 OASIS, "Web Services Federation Language (WS-Federation) Version 1.2," 22 May 2009.
]   [Online]. Available: http://docs.oasis-open.org/wsfed/federation/v1.2/os/ws-federation-1.2-
    spec-os.pdf.

[67 OpenID Foundation, "OpenID Connect," [Online]. Available: http://openid.net/connect/.
]

[68 IETF, "The OAuth 2.0 Authorization Framework RFC 6749," October 2012. [Online].
]   Available: http://tools.ietf.org/html/rfc6749.

[69 IETF, "The OAuth 2.0 Authorization Framework: Bearer Token Usage RFC 6750," October
]   2012. [Online]. Available: http://tools.ietf.org/html/rfc6750.

[70 ARTS, "ARTS XML Digital Receipt Technical Specification," 2011.
]

[71 Wikipedia, "Claims-based identity," https://en.wikipedia.org/wiki/Claims-based_identity,
]   2015.

[72 IETF, "JSON Web Token (JWT)," May 2015. [Online]. Available:
]   https://tools.ietf.org/html/rfc7519.

[73 OWASP, "Certificate and Public Key Pinning," June 2015. [Online]. Available:
]   https://www.owasp.org/index.php/Certificate_and_Public_Key_Pinning.

[74 G. Hohpe and B. Woolf, Enterprise Integraton Pattern, Addison-Wesley Professional, 2003.
]

[75 G. Hohpe, "Enterprise Integration Pattern.Introduction to Integration Styles.," 2014.
]   [Online].   Available:
     http://www.enterpriseintegrationpatterns.com/patterns/messaging/IntegrationStylesIntro.ht
     ml.

[76 ARTS, "ARTS Payment Integration White Paper," 2013.
]

[77 EPASOrg, "EPASOrg," [Online]. Available: http://www.epasorg.eu/.
]

[78 ARTS, "ARTS XML Transaction Tax Technical Specification," 2008.
]

[79 ARTS, "ARTS Pricing Service Interface Technical Specification," 2014.
]

[80 ARTS, "ARTS XML Customer Technical Specification," 2009.
]

[81 ARTS, "ARTS XML Self-Service Order Interface Technical Specification," 2012.
]

[82 ARTS, "IXRetail Notification Event Architecture for Retail Technical Specification," 2006.
]

[83 ARTS, "ARTS POSLog V6.0.0," 2014.
]

[84 Wikipedia, "Event-driven architecture," 2015. [Online]. Available:
] https://en.wikipedia.org/wiki/Event-driven_architecture.

[85 ARTS, "ARTS XML Item Maintenance Technical Specification Version 1.2.1," 2--7.
]

[86 Wikipedia, "Shard (database architecture)," 2015. [Online]. Available:
] https://en.wikipedia.org/wiki/Shard_(database_architecture).

[87 Wikipedia, "CAP theorem," 2015. [Online]. Available:
] https://en.wikipedia.org/wiki/CAP_theorem.

[88 Wikipedia, "Extract, transform, load," 2015. [Online]. Available:
] https://en.wikipedia.org/wiki/Extract,_transform,_load.

[89 ARTS, "ARTS Change4Charity Technical Specification Version 1.0," 2013.
]

[90 ARTS, "ARTS Data Warehouse Model Narrative Description Version 3.0," 2013.
]

[91 G. Young, "CQRS, Task Based UIs, Event Sourcing agh!," 16 February 2010. [Online].
] Available: http://codebetter.com/gregyoung/2010/02/16/cqrs-task-based-uis-event-sourcing-agh/.