

# An Architectural Pattern for Adaptable Middleware Infrastructure

Jason Mitchell

[JMitchell@newroadsoftware.com](mailto:JMitchell@newroadsoftware.com)

Arturo J Sánchez-Ruíz

University of North Florida, CIS Department

4567 St Johns Bluff Rd S

Jacksonville, FL3224

[Arturo@acm.org](mailto:Arturo@acm.org)

**Abstract** - *Middleware technologies change so rapidly that designers must adapt existing software architectures to incorporate new emerging ones. This paper proposes an architectural pattern and guidelines to abstract the communication barrier whereby allowing the developer to concentrate on the application logic.*

*We demonstrate our approach and the feasibility of easily upgrading the middleware infrastructure in the context of a sample application and three case studies using three different middlewares on the .NET framework.*

**Keywords:** Middleware, Architectural Pattern, Infrastructure, Distributed Communication.

## 1 Introduction

Let us suppose that we have chosen a middleware and have written a distributed system. This means that we have an application logic that interfaces with the middleware Application Programming Interface (API). This also means that we have probably defined a messaging infrastructure whereby we have defined the messages being passed between certain components of the application. Most of the time this is done by means of some sort of interface definition language (IDL) so the messaging infrastructure knows how to marshal/unmarshal the complex types across the network, which calls for mappings between our application-specific complex types and the types defined as our messages. An application so designed is inherently prone to be tightly dependent on the middleware in question!

What does this mean for our application developers? They would potentially need to modify large segments of the logic that uses the middleware API when evolution imposes the use of a new middleware. This also means that they would need to write a new set of classes to map to the new set of interface definition types. This not only means more development but the applications themselves need to be recompiled, retested, and redeployed.

This is far too much overhead for something that could have been avoided from the beginning. This paper shows an approach to avoid these pitfalls, which could lead to a better utilization of resources such as time and money.

## 2 Related Work

There is a plethora of middleware architectures, frameworks, and protocols. They try to tackle different problems and complexities. Each additional feature of a middleware has a cost associated with it; most of the time it's a performance hit or a new learning curve to tackle for the team.

New policy-driven middleware approaches like QuO handle many scenarios such as dynamic security requirements, ad hoc networking of devices, and context-aware computing [15].

Resource management becomes a key factor in the middleware arena. Resource awareness and dynamic reallocation of resources are important responsibilities of a resource management system. A way of adapting to the network via reflection techniques is a key approach one framework has attempted to accomplish [02].

Many other examples of middleware architectures make use of some of the aspects discussed already. Some examples are Artic Bean developed at the University of Tromsø [01], a composable reflective framework at the University of California, Irvin [17], an open network platform protocol developed at Ericson [09], and an Advanced Communication Toolkit (ACT) developed at Rutgers University [05].

There are architectural designs that apply to the system as a whole that address interaction between the application domains; two of which are Federated Architectures [20] and Service Oriented Architectures by Microsoft. They don't however directly address the ability to abstract the middleware from the application logic that utilizes it.

Most of these implementations are either built or based on commercial object-oriented middleware technologies such as OMG's CORBA, Sun's RMI, Microsoft's COM+, and IBM's MSQ. All of these commercial implementations offer great advantages when building a distributed system, and work well for certain scenarios.

It is even easy to choose which one will work best for the current implementation of the application given its domain. The unavoidable problem that arises is change: the domain, the complexity, the environment, or the application will change, and this may mean that the middleware infrastructure needs to be changed to adapt to the new requirements.

### 3 An Architectural Pattern

#### 3.1 The Problem

When developing a component-based architecture a variety of middlewares may potentially be used to handle different scenarios in the context of a distributed enterprise system. The task for the architect is therefore to design the system in such a way that adapting to change is accomplished with minimal effort.

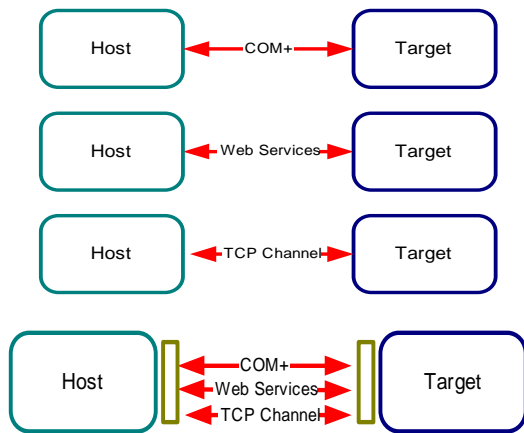


Figure 1: Using layers to expose the middleware. Logically, host and target can be a client or a server.

This requires a layer of abstraction between the applications and the middleware. This layer is shown in the bottom portion of figure 1. This abstraction layer must be non-intrusive. Otherwise, it might compromise the flexibility of the concrete middleware or introduce unnecessary dependencies in the application.

In order to study the influence of the middleware in the whole design, we have identified two focus areas. The first is the Application Programming Interface (API) and the second is the messaging infrastructure. We shall now explain these in more detail.

#### 3.2 The Application Programming Interface (API)

Each application must be written to interact with the API of the middleware. This means that we must reference external libraries so that some business logic to interoperate with the middleware. Switching to a new middleware therefore entails changing that code to now interface with the new middleware's API. Now that we have changed some of our business logic the whole system needs to be retested and the interaction code needs to be redeployed.

An example of this scenario that is often found is when the application logic passes reference to the remote server, which could potentially (tightly) couple the system to the middleware.

The solution to this is to have the host applications bind to an interface. Then, implement code to bind the logic between the interface and the middleware. This not only allows new bindings to be introduced but also saves us from having to recompile, and even re-test the application logic that uses it.

#### 3.3 The Messaging Perspective

The other place where designers might not foresee the need for future changes is in the messaging between the components. Two applications communicate with each other through the transfer of complex data types.

Current implementations of middleware offer some sort of interface definition language to define the complex types so that they can be marshaled and un-marshaled to be sent across the network

This creates a problem when designers couple the application to specific data representation schemes. With each message being passed between applications we must define the types and instruct the middleware how to send types across the network. A substantial amount of work is required to map large data objects in any interface definition language. If repeated several times to accommodate different middleware, the headache of re-implementation surfaces quickly [03].

Our pattern promotes extensibility by delegating on applications the responsibility of parsing and interpreting the messages being passed across. The only thing the middleware knows about is that a character string is being passed across. This exchanging of strings is where the flexibility and decoupling of data and messaging definition come into play. Our idea is akin to the "open binding" approach of Fitzpatrick et alia [04].

The current primary choice for this is the extensible markup language (XML), which offers a generic loosely

coupled integration environment. The messaging infrastructure is overall more extensible and adaptable and lays the messaging infrastructure foundation for a heterogeneous and diverse market of middleware communications [11].

With the introduction of an adapter-like pattern abstracting the API and an extensible messaging infrastructure the groundwork is laid for our architectural pattern, which when instantiated appropriately leads to highly flexible and adaptive distributed systems. This will become more and more important as many new middlewares will be introduced in the next years to come.

### 3.4 Adaptable Middleware Pattern

We have summarized below the architectural pattern into a recipe like format. This format is similar to the one used by Stephen Stelting [14], and includes a recipe for instantiating it.

#### 1) Pattern Properties

Type: Behavioral

Level: Component/Architectural

#### 2) Purpose

To introduce an abstraction layer that decouples the application from the middleware being used and to have the application interpreting the data independently of the transport.

#### 3) Introduction

Let's assume we have a distributed system. We would then probably decide to go with some form of middleware among components. We might then later decide to switch middlewares. We want to limit the changes necessary to switch between them. We would also like to limit any other efforts such as testing, compiling, and deploying already completed systems.

#### 4) Applicability

This pattern is very useful when distributed systems are using some sort of middleware. It is also applicable when the two communicating applications are built under different platforms.

#### 5) Description

This pattern is broken up into two parts. It involves separating the application logic from the Application Programming Interface (API) and separating the data interpretation from the middleware.

To separate the API we define an interface between the target and host application. On both sides we build the business logic to bind to these interfaces, illustrated by the class diagram in figure 2. This implies that once the logic is built and tested as long as the interfaces don't change this existing logic also doesn't need to be changed either.

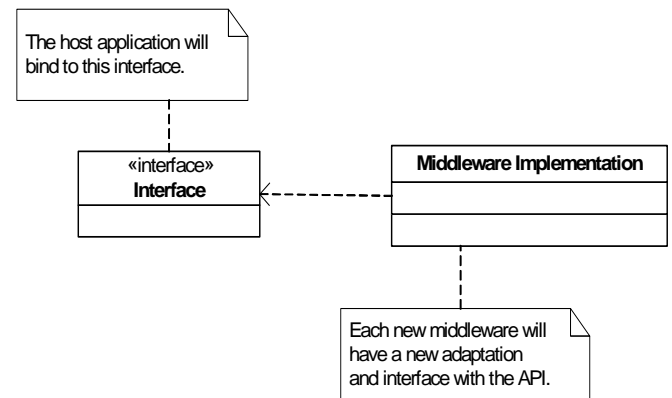


Figure 2: Class Diagram for our pattern.

Now to separate the data from the communication medium we define a way to have our applications actually interpret the data independently of the transport. Providing meta-data information within our data messages does this. We will only allow one type of message to be passed across the middleware and that is a character string. This interface will adapt to any middleware of choice.

#### 6) Implementation

Each service will have an interface defined, and each version of middleware will implement the interface, providing the middleware integration now decoupled from the application.

Secondly, the messaging infrastructure will be defined by passing a character string as the in parameter and returning a character string as the output. This way we can pass XML messages and the interpretation of the messages will be done within the application, independently of the middleware being used, which implies that there is no need to redo any mapping or defining the types with respect to the new middleware (see figure 3).

#### 7) Benefits and Drawbacks

This will significantly reduce the overhead of switching to a new middleware infrastructure. This will also provide a way in which the application logic doesn't have to be retested and deployed. Only the code integrated the new

API will have to be written and tested. Thirdly, this messaging infrastructure provides for a more extensible framework.

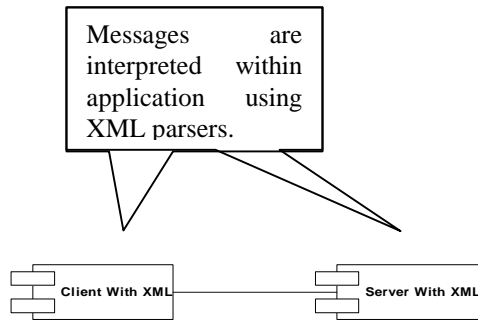


Figure 3: Messaging Infrastructure.

The only drawback might be a loss of flexibility with respect to the services that specific middlewares might provide.

## 4 A Case Study

The application that spawns our case studies is a sports statistic software system used to report real time statistics of athletic events. We will use the terms “web server” and “application server” to distinguish between the two servers.

Depending on factors such as network layout, performance requirements, and flexibility we would choose among many different middlewares to best fulfill the requirements of the system. For our case study we will use three different types of middleware all provided by the .NET framework. We chose the .NET framework because of the inherent XML tools it provides. We could have just as easily used any other platform.

For our design we will now instantiate our architectural pattern (c.f. Section III, D) and develop a solution that will enable the swapping of a new middleware easily.

The “Implementation” section in our pattern description is about allowing our applications to interpret the messaging infrastructure. We also mention that the best way to do this is by using XML as the format for passing such messages.

The .NET framework offers some tools when it comes to serializing classes into character streams using XML. For our case studies we have decided to use these tools.

We just have to create our complex types that we would like to use and then auto-generate the XML marshaling of

the complex type to a character string, which implies that we will not have to describe our types to the middleware; we only express one type of message going across the middleware, namely a simple character string.

Therefore, the messaging infrastructure will be the same for any middleware we decide to use. We will not mention the messaging in the three different case studies because they are all the same. Our applications will do the interpretation (parsing) of the messages independent of the middleware. This de-coupling allows the middlewares to change and we will never have to describe to the middleware how to marshal our messaging infrastructure.

The .NET framework has a tool that will create the methods to marshal any complex data type into a character string and read from a character string back into our object. The generated class has all of our typed parameters being passed across for ease of use within our other code.

This allows our application logic to handle any changes in our communication messages without ever depending on changes to the middleware, which means that if we change the middleware we don’t have to map our objects or define our complex types to the specific middleware.

The first step of our pattern describes a way to de-couple the middleware API from the host application. To do this we will create an interface that the host will bind to. Then we will implement the interface with a class that will act as a bridge to the target that will service the request.

The implementation of this object will be dynamically loaded. As long as the interface doesn’t change, the host application logic would not have to change or be re-compiled or be re-tested.

### 4.1 Case 1: COM+

Using Microsoft’s distributed communication protocol COM+, code snippets 1 and 2 in Appendix A shows an example in C# of how to obtain a reference to the remote server and invoke the service layer to retrieve a team object.

This class would implement the Item Service interface and act as a proxy to the remote server. There are references and configuration setup that would be coupled with this class. This implementation of the client side API to COM+ retains all syntax referring to COM+. The (XML) messaging is returned to the host and its code knows nothing about the interactions with COM+.

Now on the server side, the class that accepts the COM+ request would then forward the request onto the service layer. The service layer would retrieve the respective

team and return the XML payload string to this method to be passed back over COM+.

Just like with the client's side, all API references are kept within this abstraction layer so that they are not coupled with the applications that are using them. Furthermore, these classes would be kept in a separately linked library so that none of the application logic using this abstraction layer would have to be re-compiled after the initial release.

#### 4.2 Case 2: .NET Remoting

Let's suppose that COM+ did not suffice as a middleware between the applications. Now we have to change all of the code that references the COM+ API and change it so it will then reference .NET remoting syntax. As long as we dynamically load this class we won't have to compile, test, or re-deploy and host application code.

The minimal configuration changes and the new implementation of this class is all that was needed to swap out one middleware infrastructure to another one from the server end. Once again, we didn't have to make any changes to the business logic class or anything it uses. This saves us from having to test it. A sample is provided in Appendix A code snippets 3 and 4.

#### 4.3 Case 3: Web Services

As a third example we will now communicate with the remote server using web services. Under the .NET framework we would need to change some configuration information, such as add a reference to the web service, and compile the web service proxy. Each toolkit used to create a web client or server would be different. Once the proxy is built you just refer to it like any other object. The .NET framework has done a lot to make the integration with web services very seamless. There are many more protocols such as CORBA that make it more difficult to integrate with.

The server side portion is not so straightforward. Not only one needs to extend a web class, but also to mark each method as one published by this web service. A sample is provided in Appendix A, code snippet 5.

Once a request is accepted this service will then pass the request onto the middleware independent application layer that will handle the request.

#### 4.4 Summary

As shown in all three cases, the decoupling of the middleware infrastructure from our business applications can be achieved by applied our architectural pattern, which calls for the separation of the application from API-specific functionality, and the introduction of an extensible

messaging framework. We demonstrated this strategy with three different middlewares but this could just as easily been done with any middleware on the market.

Although the case studies have been implemented using the .NET technology, applying our pattern under J2EE is equally seamless because the middleware used at the core of J2EE is remote method invocation (RMI). RMI abstraction could also be achieved with our pattern similarly to the way in which we did it with .NET remoting.

## 5 Conclusions and Future Work

The case studies presented in our paper demonstrate that swapping among three different middlewares can be accomplished by a small amount of configuration changes and only a few systematic modifications to the source code. The real key is that none of the actual business logic on the client and server side needed to be recompiled or tested. Only the code that depended on the specific middleware infrastructure had to be altered.

Since changes associated with the middleware are inevitable for some application domains, developers should prepare in advance to face them. In this paper we have presented an architectural pattern that enables the interchanging of middlewares with minimal effort and overhead for the development team

As with every form of software design pattern, there are trade offs. A possible trade off with this approach could be some performance loss with the introduction of message parsing done at a textual level and the additional overhead that comes with another layer of indirection. Future work includes a performance analysis of this approach as opposed to a less flexible approach, and the automatic generation of some of the components from simple specifications.

## Appendix

### 1) COM +Client

```
public class COMTeamClient : ITeamService
{
    public string GetTeamById(string teamrequest)
    {
        try
        {
            TeamMgr mgr = new TeamMgr();

            return mgr.GetTeamById(teamrequest);
        }
    }
}
```

```

    }
    catch(Exception ex)
    {
        throw ex;
    }
}
}

```

## 2) COM + Server

```

[Transaction(TransactionOption.Required)]
[Guid("822A6BC5-1C84-4052-838E-FA47E6EDADC3")]
public class TeamComponentService :
    ServicedComponent, ITeamService
{
    public string GetTeamById(string teamId)
    {
        return new
            TeamService().GetTeamById(teamId);
    }
}

```

## 3) .NET Remoting Client

```

public class RemotingTeamClient : ITeamService
{
    public string GetTeamById(string teamrequest)
    {
        try
        {
            string url =
                "http://localhost/TeamService/Team.rem";

            TeamMgr mgr = (ITeamService)
                Activator.GetType(typeof (ITeamService),url);

            return mgr.GetTeamById(teamrequest);
        }
        catch(Exception ex)
        {

```

```

            throw ex;
        }
    }
}

```

## 4) .NET Remoting Server

```

public class TeamRemotingService :
    MarshalByRefObject, ITeamService
{
    public TeamRemotingService()
    {
    }

    public Team GetTeamById(int teamId)
    {
        return new
            TeamService().GetTeamById(teamId);
    }
}

```

## 5) Web Service Server

```

public class TeamService : WebService,
    ITeamService
{
    [WebMethod]
    public Team GetTeamById(int teamId)
    {
        return new
            TeamService().GetTeamById(teamId);
    }
}

```

## References

- [1] Anderson, Anders. "Artic Beans: Configurable and Reconfigurable Enterprise Component Architectures". IEEE Distributed Systems Online, 2001, Vol 2, Number 7. See <http://dsonline.computer.org/0107/features/and0107.htm>

- [2] Duran, Hector. "A Resource Management Framework for Adaptive Middleware". IEEE, March 2000 pp 206-209.
- [3] Emmerich, Wolfgang. Schwarz, Walter. Finkelstein, Anthony. "Markup Meets Middleware". ACM. Proceedings of the Seventh IEEE Workshop on Future Trends of Distributed Computing Systems. December 20 1999, Tunisia, South Africa, p 261.
- [4] Fitzpatrick, Tom. Blair, G. Coulson, G. Davies, N. Robin, P. "Supporting Adaptive Multimedia Applications through Open Bindings". Proceedings of the International Conference on Configurable Distributed Systems, March 04-06, 1998. Annapolis, Maryland.
- [5] Franco, Cristian. Marsic, Ivan. "An Advanced Communication Toolkit for Implementing the Broker Pattern".. Proceedings of the 19th IEEE International Conference on Distributed Computing Systems. May 31 – June 4 1999. Austin, Texas, p 458.
- [6] Geihs, Kurt. "Middleware Challenges Ahead". IEEE-Computer, Jan-June 2001, Volume 34, pp 24-30.
- [7] [07] Thompson, Charles. "A Scout's Guide to Three-Tier Architecture". Database Programming and Design, August 1997.
- [8] [08] Gold-Berstein, Beth. "Race to the Middle", Database Programming & Design: Volume 11, February 1998, pp 28-33.
- [9] [09] Jozic, Danijel. Osmanlic, T. Huljenic, D. Sinkovic, V. "Open Network Platform for Multiprotocol Communication". Proceedings of the 25th Annual IEEE Conference on Local Computer Networks. November 09-10, 2000, Tampa, Florida.
- [10] Mullender, Maarten. "Some Architectural Patterns for the Enterprise", Webcast, Microsoft 2002. See <http://www.microsoft.com/usa/webcasts/ondemand/960.asp>
- [11] Nusser, Gerd. Schimkat, Dieter. "Rapid Application Development of Middleware Components by Using XML", Proceedings of the 12th International Workshop on Rapid System Prototyping. June 25-27, 2001, Monterey, California.
- [12] OMG. "The Common Object Request Broker: Architecture and Specification Revision 2.2". 492 Old Connecticut Path, Framingham, MA 01701, USA, February 1998.
- [13] Schaeffer, Jonathan. Sztipanovits, Janos. Karsai, Gabor. Moore, Michael. Ledeczi, Akos. Long, Earl. The Enterprise Model for Developing Distributed Applications. Proceedings of the IEEE Conference and Workshop on Engineering of Computer-Based Systems. March 07-12, 1999, Nashville, Tennessee, pp 225.
- [14] Stelling, Stephen. Maassen, Olav. Applied Java Patterns. Published by Sun Microsystems Press A Prentice Hall Title. 2002.
- [15] Tripathi, Anand. "Challenges Designing Next-Generation Middleware Systems". Communications of the ACM, June 2002, Volume 25, No. 6, pp 39-42.
- [16] Venkatasubramanian, Nalini. Deshpande, Mayur. Mohapatra, Shivjit. Sebastian, Gutierrez-Nolasco, Wickramasuriya, Jehan. "Design and Implementation of a Composable Reflective Middleware Framework". Proceedings of the 21st International Conference on Distributed Computing Systems, April 16-19, 2001, Mesa, Arizona, pp 644.
- [17] Venkatasubramanian, Nalini. "Safe Composability of Middleware Services". Communications of the ACM, June 2002, Vol 45, No. 6, pp 49-52.
- [18] .NET Remoting, Tutorial, Microsoft 2002. See <http://msdn.microsoft.com/library/default.asp?url=/library/en-us/dndotnet/html/hawkremoting.asp>
- [19] .NET Remoting, Tutorial, 2002. See <http://www.dotnetremoting.cc/>
- [20] Zhao, Liping, Fernandez, George, Wijegunaratne, Inji. "Design Patterns for a Federated Architectural Form". Journal of Object Technology, Vol 1, No. 2, July-August 2002