

Dear Reader,

This is a compilation of **drafts** posted by Greg Young to the Google Group. These drafts are an extremely educating read but will probably change by the time the book comes out.

Check out the Google group for the latest updates on these drafts and for submitting any feedback
<http://groups.google.com/group/DDDDReviewDiscussion>

More information on the DDDD/CQRS, article and source code references and updates:
<http://abdullin.com/cqrs/>

Best regards,
Rinat Abdullin

Technology - Lokad.com | Blog - Abdullin.com | +7 917 4613826

DDDD by Greg Young

<http://groups.google.com/group/DDDDReviewDiscussion>

4/16/2010
<http://codebetter.com/blogs/gregyoung/>
Greg Young

Contents

A Stereotypical Architecture.....	3
Application Server.....	3
Client Interaction.....	4
Analysis of the Stereotypical Architecture.....	6
Summary.....	8
Works Cited.....	9
Events as a Storage Mechanism.....	10
What is a Domain Event?.....	10
Other Definitions and Discussion.....	11
Events as a Mechanism for Storage.....	12
There is no Delete.....	16
Performance and Scalability.....	17
Rolling Snapshots.....	18
Impedance Mismatch.....	21
Business Value of the Event Log.....	22
As the events represent every action the system has undertaken any possible model that one would want describing the system can be built from the events.....	23
Works Cited.....	25
Building an Event Storage.....	26
Structure.....	26
Operations.....	27
Rolling Snapshots.....	29
Event Storage as a Queue.....	31
User Do Things.....	35
Commands.....	37
User Interface.....	39
Works Cited.....	42

A Stereotypical Architecture

Before moving into architectures for Domain Driven Design based projects it is important to start off by analyzing what is generally considered to be the standard architecture that many try to apply to projects. We can from that point attempt to improve upon the stereotypical architecture in small rational steps while trying to minimize the cost in terms of productivity for each step towards a better architecture.

Below is shown a diagram of a stereotypical architecture.

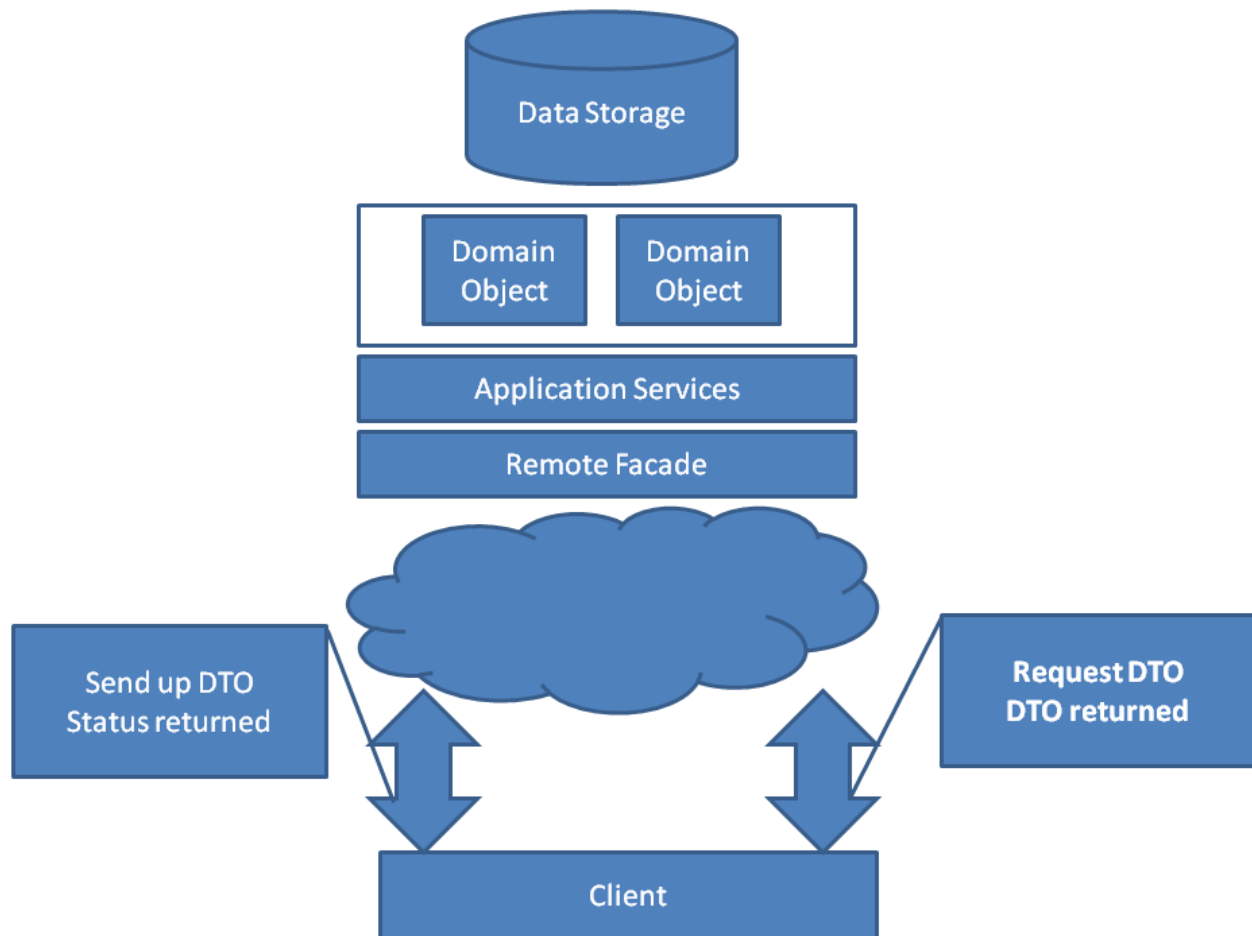


Figure 1 A Stereotypical Architecture

Application Server

The above architecture is centered upon a backing data storage system. This system although typically a RDBMS does not have to be, it could just as easily be a key/value store, and object database, or even plain XML files. The important aspect of the backing store is that it is representing the current state of objects in the domain.

Above the backing data storage lies an Application Server. An area of logic, labeled as the domain in Figure 1 contains the business logic of the system. In this area validation and orchestration logic exists for the processing of requests given to the Application Server.

It is important to note that although Figure 1 is drawn without a data tier one could place a data tier in between the Application Server and the Data Storage. It is also important to note that a "domain" is not necessary to achieve this architecture, one could also use other patterns such as Table Module or Transaction Script. With these only existing as Application Services.

Abstracting the "domain" one will find a facade known as Application Services. Application Services provide a simple interface to the domain and underlying data, they also limit coupling between the consumers of the domain and the domain itself.

On the outside of the Application Server sits some type of Remote Facade. This could be many things such as SOAP, custom TCP/IP, XML over HTTP, TomCat, or even a person manually typing messages that arrive tied to the legs of pigeons. The Remote Facade may or may not be abstracted away from its underlying technology mechanism depending on the situation and tools that are involved.

The overall usage of an Application Server to abstract away the data storage of a system and to provide a centralized location of business logic has become very popular over the years and at the time of this writing is in many circumstances considered to be the default architecture applied to many systems.

Client Interaction

Interacting with the Application Server there is a / are many client(s). The general interaction of the clients can be seen in Figure 2.

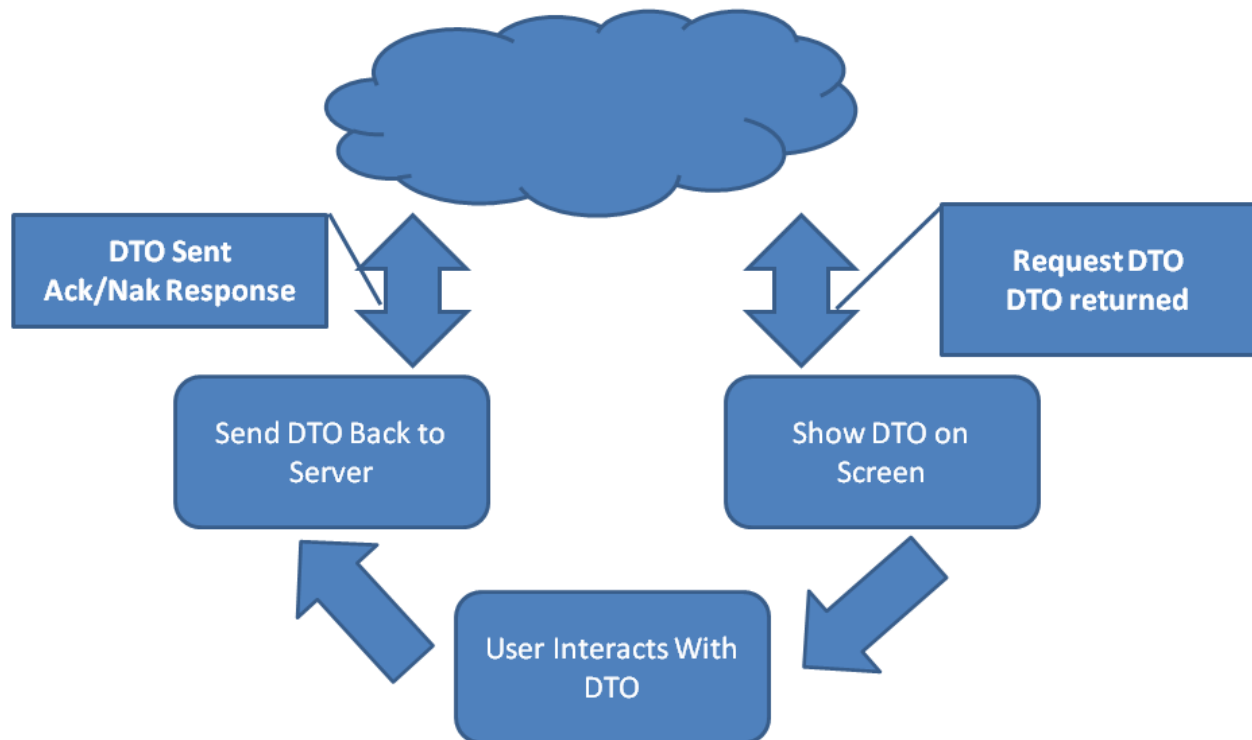


Figure 2 Typical Client Interaction

The basic interaction of the client can be described as a DTO (Data Transfer Object) up/down interaction. Going through the lifecycle of an operation is the easiest way to show the functioning of the API. A user goes to a screen, perhaps to edit a customer. The client sends a request to the remote facade for a DTO representing Customer #id. The Remote facade loads up the domain objects required, and maps the domain objects to a DTO that is then returned to the client. An example of DTO in XML format can be seen in Figure 3 but the basic explanation is that the DTO in this stereotypical architecture contains the current state of the object in questions.

The client will then display the information received from the Remote Facade on the screen allowing the user to interact with it. This is very often done utilizing a view model and/or data binding with the view.

At some point the user will be complete with the editing of the data on the screen and will through some action cause the UI to “Save” the data. Generally this is implemented through a “Save Button” although some User Interfaces will instead just have you leave the current data which forces a save.

```
<Contact id="1234">
  <Name>Greg Young</Name>
  <Address>
    <Street>111 Some St.</Street>
    <City>Vemon</City>
    <State>CT</State>
    <Zip>06066</Zip>
    <Country>USA</Country>
  </Address>
</Contact>
```

Figure 3 Example in XML of a DTO

The processing of a Save on the client will take the data that has been edited by the user on the screen, pack it back into a DTO (usually identical to the DTO it requested from the Remote Façade for displaying to the user). It will then send this DTO back up to the Application Server.

The Application Server receiving the DTO will then start a transaction/session, map the DTO back to the domain objects, allow the domain objects to verify any changes, then save the changes within the domain objects back to the data storage likely through the use of something like an Object Relational Mapper that has the ability to distinguish what has changed in the domain objects and update the data storage accordingly. The Application Server will return to the client either an Acknowledgement (Ack) that the change has been persisted or it will return a series of errors as to why it was unable to persist the changes.

Analysis of the Stereotypical Architecture

The architecture provided above as with any architecture has many properties. Some of these properties are good under certain scenarios and other properties can be extremely bad in others. As architects we should really be trying to align many of these properties to best fit our needs.

Simplicity

When looking at properties it is important to note what the most likely property is for a given architecture becoming popular. In the architecture above the most likely property defining its popularity is that it is simple. One could teach a Junior developer how to interact with a system built using this architecture in a very short period of time. Going along with the simplicity, the architecture is completely generic. One could use this architecture on every project. Along with being able to use it on every project, because many people are doing it, its likely that if a team brings on a new member the new member will be intimately familiar with the general architecture of their system again lowering on ramp up costs.

The combination of these two items allows teams to become extremely adept at applying this architecture and more important it allows them to use it as a default architecture. The thought process

of needing to align non-functional requirements really goes away as they know that this architecture will be “good enough” for 80% of the projects that they run into.

Tooling

Many frameworks exist for the optimization of time required to create systems utilizing the architecture provided above. ORM's are the largest single example as they provide valuable services such as change tracking and transaction management with complex object graphs. Other examples would include the automapping frameworks that map from the domain objects to DTOs on both sides resulting in largely removing the amount of “plumbing code” required to map the DTOs back and forth in the Application Server.

Of course there are however also many not-so-good things associated with the architecture provided above. It being that this document is associated with Domain Driven Design the single most important of the not-so-good properties would be that **it is impossible to apply Domain Driven Design with the architecture provided.**

Domain Driven Design

The application of Domain Driven Design is not possible in the above architecture though many who are “practicing” Domain Driven Design use this architecture. The reasoning for why it is impossible can easily be seen when one looks at how the Ubiquitous Language is represented by the object model.

In the architecture above there are only four verbs (and of course synonyms for those four such as edit instead of update). The four verbs are Create, Read, Update, and Delete or CRUD as they are commonly known in the industry. Because the Remote Façade has a data oriented interface the Application Services must necessarily have the same interface.

This means that there are no other verbs within the domain. When however one talks with domain experts in an effort to refine an Ubiquitous Language, it is extremely rare that one ends up with a language that is focused on these four verbs.

There is a related well-known anti-pattern of domain modeling known as an “Anemic Model”.

“The basic symptom of an Anemic Domain Model is that at first blush it looks like the real thing. There are objects, many named after the nouns in the domain space, and these objects are connected with the rich relationships and structure that true domain models have. The catch comes when you look at the behavior, and you realize that there is hardly any behavior on these objects, making them little more than bags of getters and setters. Indeed these models often come with design rules that say that you are not to put any domain logic into the domain objects. Instead there are a set of service objects which capture all of the domain logic. These services live on top of the domain model and use the domain model for data” (Fowler, 2003)

The model that is being built in this architecture sounds at first to be an anemic domain model. Because the Application Services map data back and forth to DTO's the domain objects have little behavior and

are littered with getters and setters to be used in the mapping process. There is a structure to the domain showing how objects relate with one another but ...

One cannot even create an Anemic Domain Model with this architecture as then all of the business logic would be in services, here the services themselves are really just mapping DTO's to domain objects, there is no actual business logic in them. In this case a large amount of business logic is not existing in the domain at all, nor in the Application Server, it may exist on the client but more likely it exists on either pieces of paper in a manual or in the heads of the people using the system.

Architectures like the one being viewed tend to come with instructions of how to complete complex tasks by editing data in many parts of the system. A stereotypical example of this would be when changing the sex of an employee you must after go edit their health insurance information. This is far worse than the creation of an anemic model, this is the creation of a glorified excel spreadsheet.

Scaling

When one looks at the architecture provided above in the context of scaling one will quickly notice that there is a large bottle neck. The bottleneck in terms of scaling is the data storage. When using a RDBMS as 90%+ currently use this becomes even more of a problem most RDBMS are at this point not horizontally scalable and vertically scaling becomes prohibitively expensive very quickly. It is however also extremely important to remember that most systems do not **need** to scale and as such scalability is really not a grave issue in all cases.

Summary

The DTO up/down architecture employed on many projects is capable of being used for many applications and can offer many benefits in terms of simplicity for teams to work with. It cannot however be used with a Domain Driven Design based project, to attempt so will bring failure to your efforts at applying Domain Driven Design.

This architecture does however make a good baseline and the rest of this document will be focused on improving this architecture in incremental steps while attempting to limit or remove cost while adding business value at each additional step.

Works Cited

Fowler, M. (2003, 11 25). *MF Bliki: AnemicDomainModel*. Retrieved from Bliki:
<http://martinfowler.com/bliki/anemicdomainmodel>

Events as a Storage Mechanism

Most systems in production today rely on the storing of current state in order to process transactions. In fact it is rare to meet a developer who has worked on a system that maintains current state in any other way. It has not always been like this.

Before the general acceptance of the RDBMS as the center of the architecture many systems did not store current state. This was especially true in high performance, mission critical, and/or highly secure systems. In fact if we look at the inner workings of a RDBMS we will find that most RDBMSs themselves not actually work by managing current state!

The goal of this section is to introduce the concept of event sourcing, to show the benefits, to show how one can build a simple event storage system utilizing a Relational Database for underlying data management.

What is a Domain Event?

An event is something that has happened in the past.

All events should be represented as verbs in the past tense such as `CustomerRelocated`, `CargoShipped`, or `InventoryLossageRecorded`. For those who speak French, it should be in *Passé Composé*, they are things that have **completed** in the past. There are interesting examples in the English language where one may be tempted to use nouns as opposed to verbs in the past tense, an example of this would be “Earthquake” or “Capsize”, as a congressman recently worried about Guam, but avoid the temptation to use names like this and stick with the usage of verbs in the past tense when creating Domain Events.

It is absolutely imperative that events always be verbs in the past tense as they are part of the Ubiquitous Language. Consider the differences in the Ubiquitous Language when we discuss the side effects from relocating a customer, the event makes the concept explicit where as previously the changes that would occur within an aggregate or between multiple aggregates were left as an implicit concept that needed to be explored and defined. As an example, in most systems the fact that a side effect occurred is simply found by a tool such as Hibernate or Entity Framework, if there is a change to the side effects of a use case, it is an implicit concept. The introduction of the event makes the concept explicit and part of the Ubiquitous Language; relocating a customer does not just change some stuff, relocating a customer produces a `CustomerRelocatedEvent` which is explicitly defined within the language.

In terms of code, an event is simply a data holding structure as can be seen in Listing 1.

```
public class InventoryItemDeactivatedEvent {
    public readonly Guid InventoryItemId;
    public readonly string Comment;

    public InventoryItemDeactivatedEvent(Guid id, string comment) {
        InventoryItemId = id;
        Comment = comment;
    }
}
```

Listing 1 A Simple Event

The code listing looks very similar to the code listing that was provided for a Command the main differences exist in terms of significance and intent. Commands have an intent of asking the system to perform an operation where as events are a recording of the action that occurred.

Other Definitions and Discussion

There is a related concept to a Domain Event in this description that is defined in Streamlined Object Modeling (SOM). Many people use the term “Domain Event” In SOM when discussing “The Event Principle”

Model the event of people interacting at a place with a thing with a transaction object. Model a point-in-time interaction as a transaction with a single timestamp; model a time-interval interaction as a transaction with multiple timestamps. (Jill Nicola, 2002II, p. 23)

Although many people use the terminology of a Domain Event to describe this concept the terminology is not having the same definition as a Domain Event in the context of this document. SOM uses another terminology for the concept that better describes what the object is, a Transaction. The concept of a transaction object is an important one in a domain and absolutely deserves to have a name. An example of such a transaction might be a player swinging a bat, this is an action that occurred at a given point in time and should be modeled as such in the domain, this is not however the same as a Domain Event.

This also differs from Martin Fowler’s example of what a Domain Event is.

“Example: I go to Babur’s for a meal on Tuesday, and pay by credit card. This might be modeled as an event, whose type is “Make Purchase”, whose subject is my credit card, and whose occurred date is Tuesday. If Babur’s uses an old manual system and doesn’t transmit the transaction until Friday, then the noticed date would be Friday.” (Fowler)

Further along

“By funneling inputs of a system into streams of Domain Events you can keep a record of all the inputs to a system. This helps you to organize your processing logic, and also allows you to keep an audit log of the system” (Fowler)

The astute reader may pick up on the fact that what Martin is actually describing here is a Command as was discussed previously when discussing Task Based UIs. The language of “Make Purchase” is wrong. A purchase was made. It makes far more sense to introduce a PurchaseMade event. Martin did actually make a purchase at the location, they did actually charge his credit card, and he likely ate and enjoyed his food. All of these things are in the past tense, they have already happened one cannot undo that they have happened.

An example such as the sales example given also tends to lead towards a secondary problem when built within a system. The problem is that the domain may be responsible for filling in parts of the event. Consider a system where the sale is processed by the domain itself, how much is the sales tax? Often the domain would be calculating this as part of its calculations. This leads to a dual definition of the event, there is the event as is sent from the client without the sales tax then the domain would receive that and add in the sales tax, it causes the event to have multiple definitions, as well as forcing mutability on some attributes. One can bypass this by having dual events (one for the client with just what it provides and another for the domain including what it has enriched the event from the client with) but this is basically the command event model and the linguistic problems still exist.

A further example of the linguistic problems involved can be shown in error conditions. How should the domain handle the fact that a client told it to do something that it cannot? This condition can exist for many reasons but let’s imagine a simple one of the client simply not having enough information to be able to source the event in a known correct way. Linguistically the command/event separation makes much more sense here as the command arrives in the imperative “Place Sale” while the event is in the past tense “SaleCompleted”. It is quite natural for the domain to reject a client attempting to “Place a sale”; it is not natural for the domain to tell the client that something in the past tense no longer happened. Consider the discussion with a domain expert; does the domain have a time machine? Parallel realities are far too complex and costly to model in most business systems.

These are exactly the problems that have led to the separation of the concepts of Commands and Events. This separation makes the language much clearer and although subtle it tends to lead developers towards a clearer understanding of context based solely on the language being used. Anytime one ends up with dual definitions of a concept there is a weight placed on the developer to recognize and distinguish context, this weight can translate into both ramp up time for new developers on a project and another thing a member of the team needs to “remember”. Anytime a team member needs to remember something to distinguish context there is a higher probability that it will be overlooked or mistook for another context. Being explicit in the language and avoiding dual definitions helps make things clearer both for domain experts, the developers, and anyone who may be consuming the API.

Events as a Mechanism for Storage

When most people consider storage for an object they tend to think about it in a structural sense. That is when considering how the “sale” discussed above should be stored they think about it as being stored as a “Sale” that has “Line Items” and perhaps some “Shipping Information” associated with it. This is not

however the only way that one can conceptualize the problem and other solutions offer different and often interesting architectural properties.

Consider for a moment the creation of a small Order object for a web based sale system. Most developers would envision something similar to what is represented in Figure 1. That is a structural viewpoint of what the Order is. An Order has n Line Items and Shipping Information. Of course this is an overly simplified view of what an Order is but one can see that the focus is upon the structure of the order and its parts.

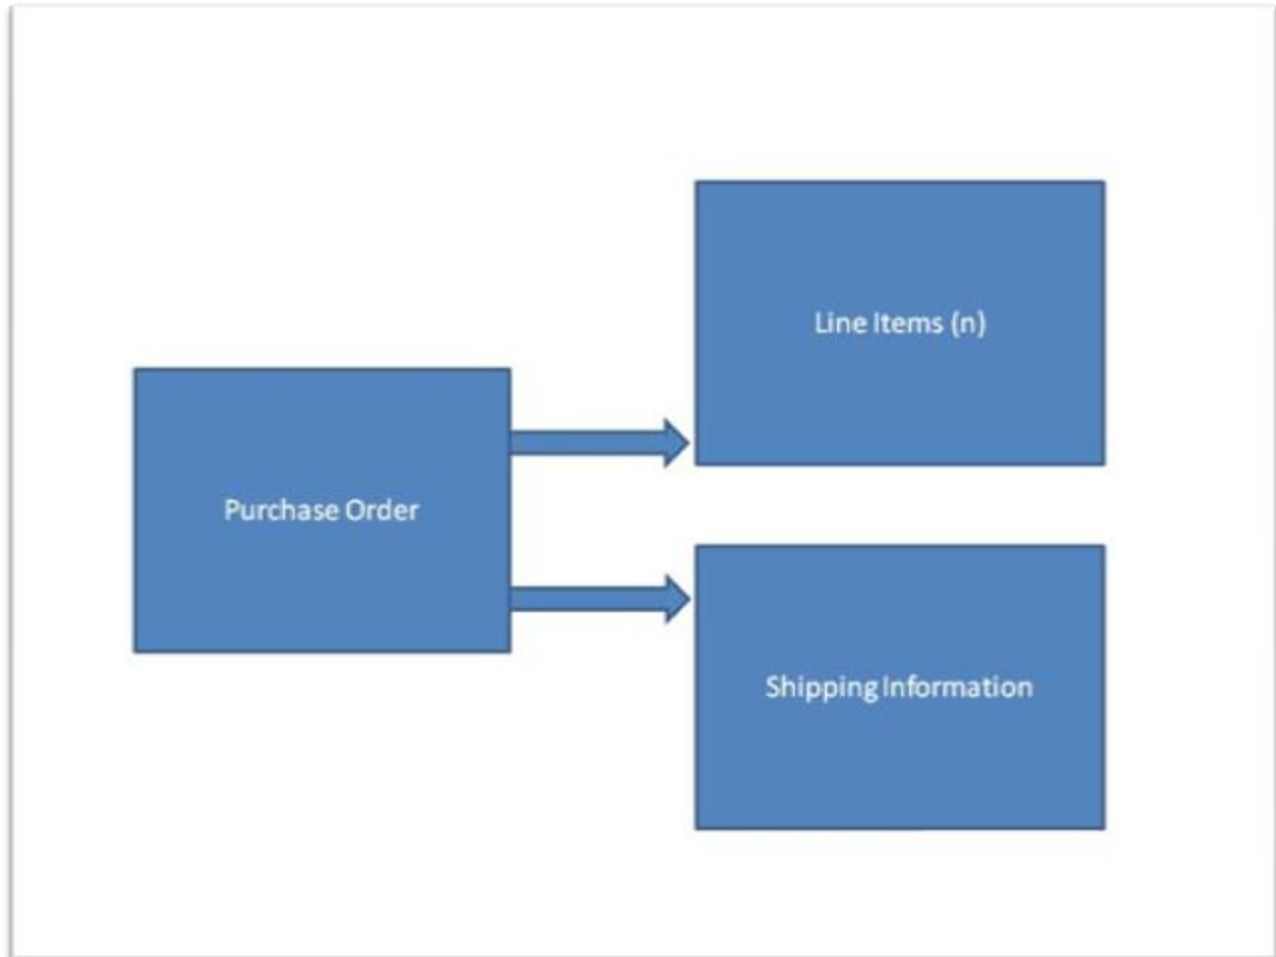


Figure 4 A Structural View of an Order

This is not the only way that this data can be viewed. Previously in the area of discussions there was a discussion about the concept of a transaction. Developers deal with the concept of transactions regularly, they can be viewed as representing the change between a point and the next subsequent point. They are also regularly called “Deltas”. One can always define what the delta is between two states but more often than not this is left to be an implicit concept, usually relegated to a framework such as Hibernate in the Java world or Entity Framework in the Microsoft world. These frameworks save the original state and then calculate the differences with the new state and update the backing data

model accordingly. The making of these deltas explicit can be highly valuable both in terms of technical benefits and more importantly in business benefits.

One can see the usage of deltas in many mature business models. The canonical example of delta usage is in the field of accounting. When looking at a ledger such as in Figure 2 one can see that each transaction or delta is being recorded. Next to it is a denormalized total of the state of the account at the end of that delta. In order to calculate this number one must just apply the delta and use the last known value. One can trust the last known value because at any given point one could re-run all the transactions from the “beginning of time” for that account and reconcile the validity of that value. In other words one has a verifiable audit log.

Date	Comment	Change	Current Balance
1/1/2000	Deposit from 1372	+10000.00	10000.00
1/3/2000	Check 1	-4000.00	6000.00
1/4/2000	Purchase Coffee	-3.00	5997.00
1/6/2000	Purchase Internet	-5.00	5992.00
1/8/2000	Deposit from 1373	+1000.00	6992.00

Figure 5 A Simplified Ledger

Because one has all of the transactions or deltas associated with the account one can step through each transaction. The “Current Balance” at any point can be derived either by looking at the “Current Balance” or by adding up all of the “Changes” since the beginning of time for the account. The second property is obviously valuable in a domain such as accounting as accountants are dealing with money and the ability to check that calculations were performed correctly is extremely valuable, it was even more valuable before computers when it was common place to have an exhausted accountant make a mistake in a calculation at 3 am when they should be sleeping instead of working with the books.

There are however some other interesting properties to this mechanism of representing state, one can go back and look at what a state was at a given point in time. Consider for example that the account was allowed to reach a balance of below zero and there is a rule that says it is not supposed to. One could look at the account as it was just prior to processing that transaction that put it into the invalid state and see what state it was in, making it far easier to reproduce what often times end up as heisenbugs in other circumstances.

These types of benefits are not only limited to naturally transaction based domains though. In fact **every** domain is a naturally transaction based domain when one is applying Domain Driven Design. When applying Domain Driven Design one has a heavy focus on behaviors, normally coinciding with use cases. One is interested in how users **use** the system.

Returning to the Order example from Figure 1, one could represent the same order in the form of a transactional model as shown in Figure 3.

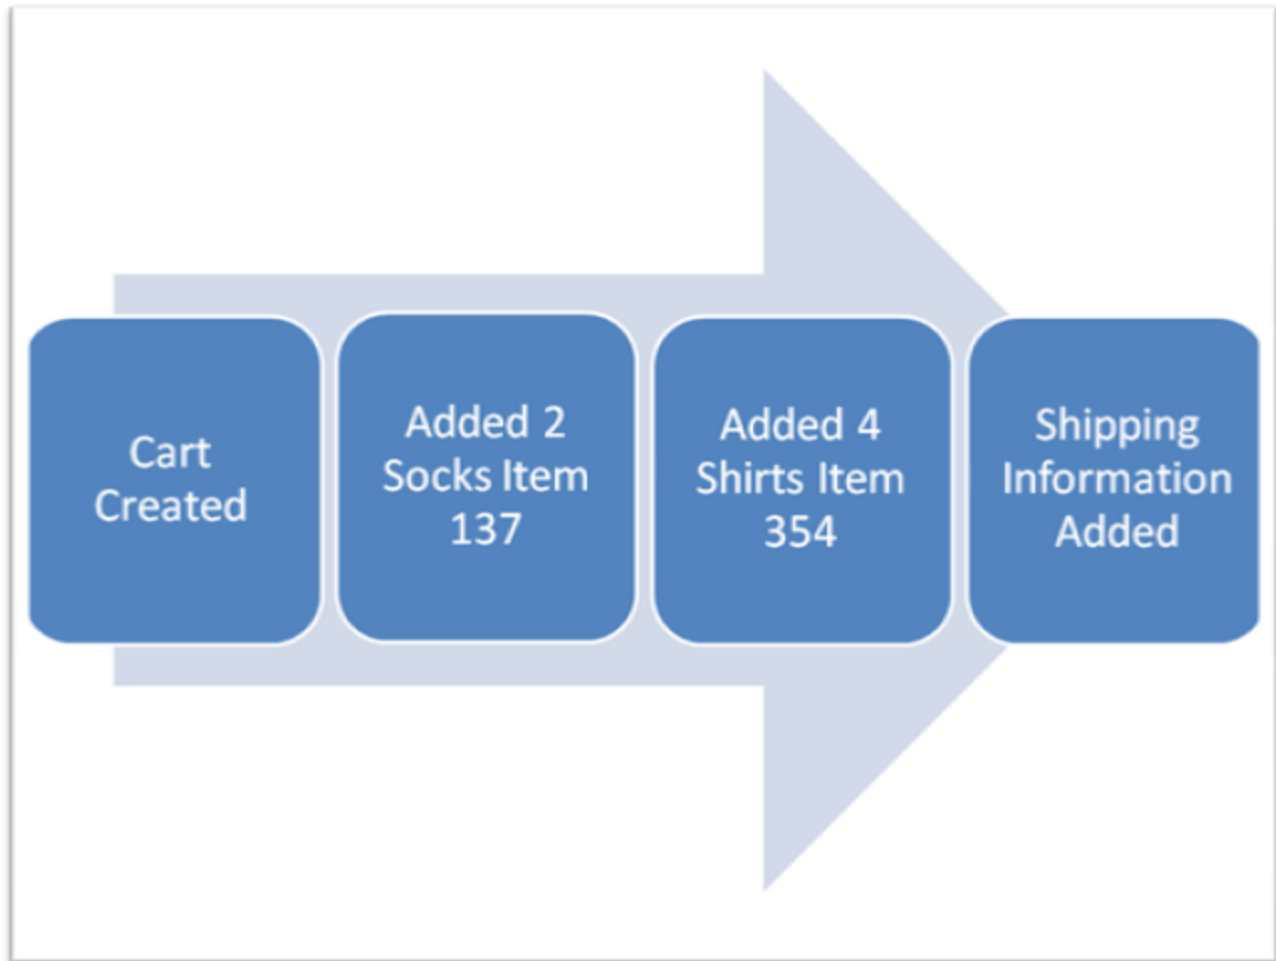


Figure 6 Transactional View of Order

One can apply this type of model to any object. By replaying through the events that have occurred one can return the object to the last known state. There is a structural representation of the object, but it exists only by replaying previous transactions to return the structure to its last known state, data is not persisted in a structure but as a series of transactions. One very interesting possibility here is that unlike when one stores current state in a structural way there is no coupling between the representation of current state in the domain and in storage, one can vary the way one represents current state in the domain without thought of the persistence mechanism.

It is vitally important to note the language in Figure 3. All of the verbs are in the past tense. These are Domain Events. Consider what would happen if the language were in the imperative tense, “Add 2 socks item 137”, “Create Cart”. What if there were behaviors associated with adding an item (such as reserving it from an inventory system via a webservice call), would one run these behaviors when reconstituting an object? What if logic has changed so that one could no longer add this item given the context? This is one of many examples where one would require dual contexts between Commands and Events, there is a contextual difference between returning to a given state and attempting to transition to a new one.

There is no Delete

A common question that arises is how one would delete information. One cannot as mentioned previously jump into the time machine and say that an event never happened (eg: delete a previous event). As such one would need to model a delete explicitly as a new transaction as shown in Figure 4. Further discussion on the business value of handling deletes in this mechanism can be found in “Business Value of the Event Log”.

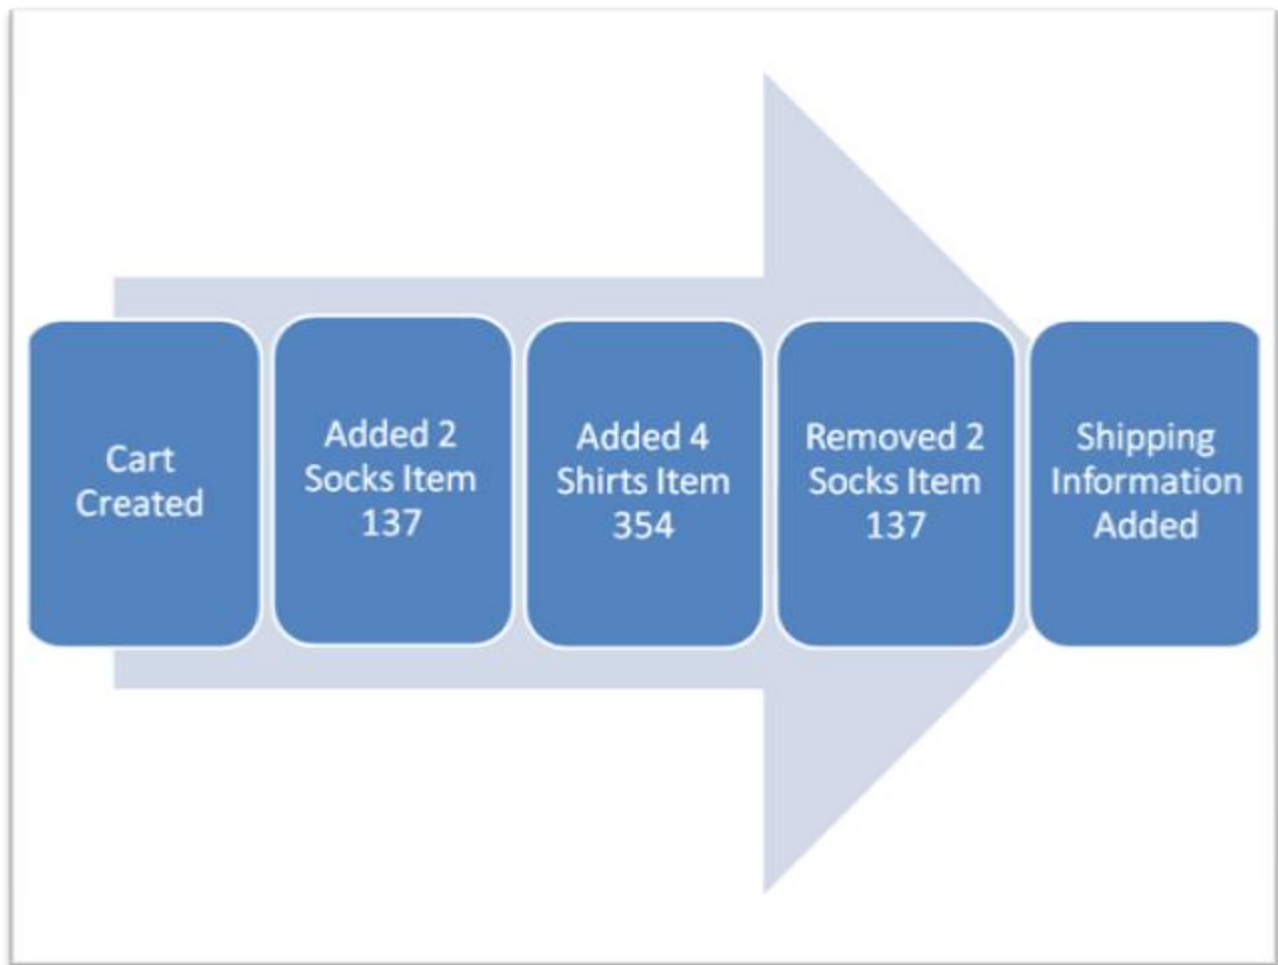


Figure 7 Transactional View of Order with Delete

In the event stream in Figure 4 the two pairs of socks were added then later removed. The end state is equivalent to not having added the two pairs of socks. The data has not however been deleted, new data has been added to bring the object to the state as if the first event had not happened, this process is known as a Reversal Transaction.

By placing a Reversal Transaction in the event stream one not only returns the object to the state as if the item had not been added, one leaves a trail that shows that the object had been in that state at a given point in time.

There are also architectural benefits to not deleting data. The storage system becomes an additive only architecture, it is well known that append-only architectures distribute more easily than updating architectures because there are far fewer locks to deal with.

Performance and Scalability

As an append-only model storing events is a far easier model to scale. There is however other benefits in terms of performance and scalability especially compared with a stereotypical relational model. As an example, the storage of events offers a much simpler mechanism to optimize as it is limited to a single append-only model. There are many other benefits.

Partitioning

A very common performance optimization in today's systems is the use of Horizontal Partitioning. With Horizontal Partitioning the same schema will exist in many places and some key within the data will be used to determine in which of the places the data will exist. Some have renamed the term to "Sharding" as of late. The basic idea is that you can maintain the same schema in multiple places and based on the key of a given row place it in one of many partitions. <TO DO GET BETTER DEFINITION AND CITING>

One problem when attempting to use Horizontal Partitioning with say a Relational Database is that one needs to define the key with which the partitioning should operate. This problem goes away when using events. **Aggregates are the only partition point in the system.** No matter how many aggregates you add or how you might change structures, the Aggregate Id associated with events is the only partition point in the system.

Horizontally Partitioning an Event Store is a very simple process.

Saving Objects

When dealing with a stereotypical system utilizing a relation data storage it can be quite complex to figure out what has changed within the Aggregate. Again many tools have been built to help alleviate the pain that arises from this often complex task but **is the need for a tool a sign of a bigger problem?**

Most ORMs can figure out the changes that have occurred within a graph. They do this generally by maintaining two copies of a given graph, the first they hold in memory and the second they allow other code to interact with. When it becomes time to save a complex bit of code is run, walking the graph the code has interacted with and using the copy of the original graph to determine what has changed while the graph was in use by the code. These changes will then be saved back to the data storage system.

In a system that is Domain Event centric, the aggregates are themselves tracking strong events as to what has changed within them. There is no complex process for comparing to another copy of a graph, one would simply ask the aggregate for its changes. The operation to ask for changes is far more efficient than having to figure out what has changed.

Loading Objects

A similar issue exists when loading objects. Consider the work that is involved with loading a graph of objects in a stereotypical relational database backed system. Very often there are many queries that

must be issued to build the aggregate. In order to help minimize the latency cost of these queries many ORMs have introduced a heuristic of Lazy Loading also known as Delayed Loading where a proxy is given in lieu of the real object. The data is only loaded when some code attempts to use that particular object.

Lazy Loading is useful because quite often a given behavior will only use a certain portion of data out of the aggregate and it prevents the developer from having to explicitly represent which data that is while amortizing the cost of the loading of the aggregate. It is this need for amortization of cost that shows a problem.

Aggregates are considered as a whole represented by the Aggregate Root. Conceptually an Aggregate is loaded and saved in its entirety. (Evans, 2001).

Conceptually it is much easier to deal with the concept of an Aggregate being loaded and saved in its entirety. The concept of Lazy Loading is not a trivial one when added and is especially not trivial when one begins to try to optimize use cases. The heuristic is needed because loading full aggregates from a relational database is operationally too slow.

When dealing with events as a storage mechanism things are quite different. There is but one thing being stored, events. One simply loads all of the events for an Aggregate and replays them. There can only ever be a single query on the system, there is no need to attempt to implement things like Lazy Loading. This is bad for people who want to build complex and quite often impressive frameworks for managing things like Lazy Loading but it is good for development teams who no longer need to learn these frameworks.

Many would quickly point out that although it requires more queries in a relational system, when one stores events one may have a huge number of events for some aggregates. This can happen quite often and a relatively simple solution exists for the problem.

Rolling Snapshots

A Rolling Snapshot is a denormalization of the current state of an aggregate at a given point in time. It represents the state when all events to that point in time have been replayed. Rolling Snapshots are used as a heuristic to prevent the need to load all events for the entire history of an aggregate. Figure 5 shows a typical Event Stream. The way that one would process the event stream is to replay the events from the beginning of time until one reached the end of the event stream.




Figure 8 An Event Stream

The problem that exists is that there may be a very large number of events between the beginning of time and the current point. One can fairly easily imagine an event stream with a million or more events that have occurred, such an event stream would be quite inefficient to load.

The solution is to use a Rolling Snapshot, to place a denormalization of the state at a given point in time. One would then only need to play the events from that point in time forward in order to load the Aggregate.




Figure 9 Event Stream with Snapshot

Figure 6 shows an Event Stream with a Rolling Snapshot placed within it. The process for rebuilding an Aggregate changes when one uses Rolling Snapshots. Instead of reading from the beginning of time forward, one would read backwards putting the events on to a stack until either there were no more events left or a snapshot was found. The snapshot would then if found be applied and the events would be popped off the stack and applied until the stack was empty.

It is important to note that although this is an easy way to conceptualize how Rolling Snapshots work, one would not want to implement this exact solution in a production system for various reasons. Further discussion on the implementation of Rolling Snapshots can be found in “Building an Event Storage”.

The snapshot itself is nothing more than a serialized form of the graph at that given point in time. **By having the state of that graph at that point in time one can avoid having to replay all the events prior to that snapshot.** Snapshots can be taken asynchronously by a process monitoring the Event Store.

Introducing Rolling Snapshots allows you to control the worst case that you will run into when loading from events. The maximum number of events that would be processed can be tuned to optimize performance for the system in question. With the introduction of Rolling Snapshots it is a relatively trivial process to achieve one to two orders of magnitude of performance gain on the two operations that the Event Storage supports. It is important though to remember that Rolling Snapshots are just a heuristic and that conceptually one still views the event stream in its entirety.

Impedance Mismatch

Using events as a storage mechanism also offers very different properties when compared to a typical relational model when one analyzes the impedance mismatch that exists between a typical relational model and the object oriented domain model. Scott Ambler describes the problem in an essay on agiledata.org as

“Why does this impedance mismatch exist? The object-oriented paradigm is based on proven software engineering principles. The relational paradigm, however, is based on proven mathematical principles. Because the underlying paradigms are different the two technologies do not work together seamlessly. The impedance mismatch becomes apparent when you look at the preferred approach to access: With the object paradigm you traverse objects via their relationships whereas with the relational paradigm you join the data rows of tables. This fundamental difference results in a non-ideal combination of object and relational technologies, although when have you ever used two different things together without a few hitches?” (Ambler)

The impedance mismatch between the domain model and the relational database has a large cost associated with it. There are many tools that aim to help minimize the effects of the impedance mismatch such as Object Relational Mappers (ORM). They tend to work well in most situations but there is still a fairly large cost associated to the impedance mismatch even when using tools such as ORMs.

The cost is that a developer really needs to be intimately with both the relational model and the object oriented model. They also need to be familiar with the many subtle differences between the two models. Scott identifies this with

“To succeed using objects and relational databases together you need to understand both paradigms, and their differences, and then make intelligent tradeoffs based on that knowledge.” (Ambler)

Some of these subtle differences can be found in Wikipedia under the “Object-Relational Impedance Mismatch” page but to include some of the major differences.

Declarative vs. imperative interfaces — Relational thinking tends to use data as interfaces, not behavior as interfaces. It thus has a declarative tilt in design philosophy in contrast to OO's behavioral tilt. (Some relational proponents propose using triggers, stored procedures, etc. to provide complex behavior, but this is not a common viewpoint.) (Object-Relational Impedance Mismatch)

Structure vs. behaviour — OO primarily focuses on ensuring that the structure of the program is reasonable (maintainable, understandable, extensible, reusable, safe), whereas relational systems focus on what kind of behaviour the resulting run-time system has (efficiency, adaptability, fault-tolerance, liveness, logical integrity, etc.). Object-oriented methods generally assume that the primary user of the object-oriented code and its interfaces are the application developers. In relational systems, the end-

users' view of the behaviour of the system is sometimes considered to be more important. However, relational queries and "views" are common techniques to re-represent information in application- or task-specific configurations. Further, relational does not prohibit local or application-specific structures or tables from being created, although many common development tools do not directly provide such a feature, assuming objects will be used instead. This makes it difficult to know whether the stated non-developer perspective of relational is inherent to relational, or merely a product of current practice and tool implementation assumptions. (Object-Relational Impedance Mismatch)

Set vs. graph relationships - The relationship between different items (objects or records) tend to be handled differently between the paradigms. Relational relationships are usually based on idioms taken from set theory, while object relationships lean toward idioms adopted from graph theory (including trees). While each can represent the same information as the other, the approaches they provide to access and manage information differ. (Object-Relational Impedance Mismatch)

There are many other subtle differences such as data types, identity, and how one transactions work. The object-relational impedance mismatch can be quite a pain to deal with and it requires a very large amount of knowledge to deal with effectively.

There is not an impedance mismatch between events and the domain model. The events are themselves a domain concept, the idea of replaying events to reach a given state is also a domain concept. The entire system becomes defined in domain terms. Defining everything in domain terms not only lowers the amount of knowledge that developers need to have, it also limits the number of representations of the model one would need as the events are directly tied to the domain model itself.

Business Value of the Event Log

It needs to be made clear at the very start of this section that the value of the Event Log is directly correlated with places that you would want to use Domain Driven Design in the first place. Domain Driven Design should be used in places where the business derives competitive advantage. Domain Driven Design itself is very difficult and expensive to apply; a company will however receive high ROI on the effort if the domain is complex and if they derive competitive advantage from it. Using an Event Log similarly will have high ROI when dealing with an area of competitive advantage but may have negative ROI in other places.

Storing only current state forces one to only be able to ask certain kinds of questions of the data. For example consider orders in the stock market. They can change for a few reasons, an order can change the amount of volume that they would like to buy/sell, the trading system can automatically adjust the volume of an order, or a trade could occur lowering the volume available on the current order.

If one wants to ask a question regarding current liquidity such as the price for a given number of shares in the market, it really does not matter which of these items occurred, it does not really matter **how** the data got the way it was, it matters **what** it is at a given point in time. A vast majority of queries even in the business world are focused on the **what**, labels to send customers mails, how much was sold in April, how many widgets are in the warehouse.

There are however other types of queries that are becoming more and more popular in business, they focus on the **how**. Examples can commonly be seen in the latest buzzword “Business Intelligence”. Perhaps there is a correlation between people having done an action and their likelihood of purchasing some product? These types of questions generally focus on how something came into being as opposed to what it came out to be.

It is best to go through an example. There is a development team at a large online retailer. In an iteration planning meeting a domain expert comes up with an idea. He believes that there is a correlation between people having added then removed an item from their cart and their likelihood of responding to suggestions of that product by purchasing it at a later point. The feature is added to the following iteration.

The first hypothetical team is utilizing a stereotypical current state based mechanism for storing state. They plan that in this iteration they will add tracking of items via a fact table that are removed from carts. They plan for the next iteration that they will then build a report. The business will receive after the second iteration a report that can show them information back to the previous iteration when the team released the functionality that began tracking items being removed from carts.

This is a very stereotypical process, at some organizations the report and the tracking may be released simultaneously but this is a relatively small detail in the handling. From a business perspective the domain experts are happy, they made a request of the team and the team was able to quickly fulfill the request, new functionality has been added in a quick and relatively painless way. The second team will however have quite a different result.

The second team has been storing events; they represent their current state by building up off of a series of events. They just like the first team go through and add tracking of items removed from carts via a fact table but they also run this handler from the beginning of the event log to back populate all of the data from the time that the business started. They release the report in the same iteration and the report has data that dates back for years.

The second team can do this because they have managed to store what the system actually did as opposed to what the current state of data is. One can always go back and look and interpret the old data in new and interesting ways. No one had ever thought to track what items were removed from carts or perhaps no one had thought the number of times a user removes and items from their cart was important. These are both examples of new and interesting ways of looking at data.

As the events represent every action the system has undertaken any possible model that one would want describing the system can be built from the events.

Businesses regularly come up with new and interesting ways of looking at data. One cannot with any level of confidence predict how a business will want to look at today’s data in five years. The ability for the business to look at the data in the way that it wants in five years is of an unknown but possibly extremely high value; it has already been stated that this should be done in areas where the business

derives its competitive advantage so it is relatively easy to reason that the ability to look at today's data in an unexpected way could be a competitive advantage for the business. How do you value the possible success or failure of a company based upon an architectural decision now?

How do software teams justify looking at their Magic 8 Ball to predict what the business will need in five or even ten years? Many try to use YAGNI (You Ain't Gonna Need It) (Wikipedia) but YAGNI only applies when you actually know that you won't need it, how can one predict the dynamic world of business and how they may want to look at data in five or ten years?

Is it more expensive to actually model every behavior in the system? Yes.

Is it more expensive in terms of disk cost and thought process to store every event in the system? Yes.

Are these costs worth the ROI when the business derives a competitive advantage from the data?

Works Cited

Ambler, S. W. (n.d.). *The Object Relational Mismatch*. Retrieved from agiledata.org:
<http://www.agiledata.org/essays/impedanceMismatch.html>

Evans, E. (2001). *Domain Driven Design*. Addison Wesley.

Fowler, M. (n.d.). *Domain Event*. Retrieved from EAA Dev:
<http://martinfowler.com/eeaDev/DomainEvent.html>

Jill Nicola, M. M. (2002II). *Streamlined Object Modelling*. Prentice H.

Object-Relational Impedance Mismatch. (n.d.). Retrieved from Wikipedia:
http://en.wikipedia.org/wiki/Object-relational_impedance_mismatch

Wikipedia. (n.d.). *You ain't gonna need it*. Retrieved from wikipedia:
http://en.wikipedia.org/wiki/You_ain't_gonna_need_it

Building an Event Storage

In “Events as a Storage Mechanism” the concept of rebuilding state from a series of events was looked at from a conceptual viewpoint. This chapter will focus on the implementation of an actual Event Storage and some of the issues that come up in producing an implementation.

The implementation discussed in this chapter is not intended to be a production quality Event Storage, more so it is provided as a discussion point around how to build an Event Storage. The implementation here although not highly performant could meet the needs of a large percentage of applications that are built today.

For the explanatory implementation it is easiest to build the Event Storage in an existing technology such as a RDBMS. This will alleviate many of the technical issues that can arise that are out of the scope of a basic discussion on how to build an event storage such as transaction commit models or data locality for read performance.

Structure

A basic Event Storage can be represented in a Relational Database utilizing only two tables.

Column Name	Column Type
fkAggregateId	Guid
Data	Blob
Version	Int

Figure 10 Table Layout for Events Table

This table represents the actual Event Log. There will be one entry per event in this table. The event itself is stored in the [Data] column. The event is stored using some form of serialization, for the rest of this discussion the mechanism will assumed to be built in serialization although the use of the memento pattern can be highly advantageous.

The table is shown with the minimum amount of information possible, likely one would want to add a few columns such as the time that the change was made or context information associated with the change. Examples of context information might include the user that initiated the change, the ip address they sourced the change from, or their level of permission when they sourced the change.

A version number is also stored with each event in the Events Table. This can generally be thought of as an increasing integer for most cases. Each event that is saved has an incremented version number. The version number is unique and sequential only within the context of a given aggregate. This is because Aggregate Root boundaries are consistency boundaries.

The [fkAggregateId] column is a foreign key that should be indexed; it points to the next table which is the Aggregates table.

Column Name	Column Type
pkAggregateId	Guid
Type	Varchar
Version	Int

Figure 11 Table Layout for Aggregate Table

Author comment: I have gone back and forth between calling this concept “Aggregate” in the Event Storage in lieu of another name such as “Event Provider” as “Aggregate” is really a domain concept and an Event Storage could work without a domain.

The Aggregate table is representing the aggregates currently in the system, every aggregate must have an entry in this table. Along with the identifier there is a denormalization of the current version number. This is primarily an optimization as it could be derived from the Events table but it is much faster to query the denormalization that it would be to query the Events table directly. This value is also used in the optimistic concurrency check.

I have also included a [Type] column for this exasmple, this would be the fully qualified name of the type of aggregate being stored. This can be useful for various purposes not the least of which is debugging, it is however unnecessary for the creation of a basic Event Storage.

Operations

Event Storages are far simpler that most data storage mechanisms as they do not support general purpose querying. An Event Storage at its simplest level has only two operations. Having only two operations makes an Event Storage simpler than most data storage mechanisms as well as easier to optimize.

The first operation is to get all of the events for an aggregate. It is extremely important that the events are ordered in the same order that they were written, the version number can be used for this purpose. This can all be done quite simply using an underlying RDBMS.

```
SELECT * FROM EVENTS WHERE FKAGGREGATEID=' ' ORDER BY VERSION
```

This is the only query that should be executed by a production system against the Event Storage. A possible secondary query that can be useful is to limit this result set by an actual date to see the state of an object at a point in time, but generally a production system should not be doing this.

The other operation an Event Storage must support is the writing of a set of events to an aggregate root. This can be done either in code or in a stored procedure. A stored procedure or dynamically generated SQL containing if statements is preferred as without the insert process will take multiple round trips. The pseudo-code for the insert process can be seen in Listing 1.

```
Begin
  version = SELECT version from aggregates where pkAggregateId = "
  if version is null
    Insert into aggregates
    version = 0
  end
  if expectedversion != version
    raise concurrency problem
  foreach event
    insert event with incremented version number
  update aggregates with last version number
End Transaction
```

Listing 2 Write Operation in Event Storage

The write operation is also relatively simple though there are a few subtleties to be found within it. The basic narrative is that it first checks to see if an aggregate exists with the unique identifier it is to use, if there is not one it will create it and consider the current version to be zero. It will then attempt to do an optimistic concurrency test on the data coming in if the expected version does not match the actual version it will raise a concurrency exception. Providing the versions are the same, it will then loop through the events being saved and insert them into the events table, incrementing the version number by one for each event. Finally it will update the Aggregates table to the new current version number for the aggregate. It is important to note that these operations are in a transaction as it is required to insure that optimistic concurrency amongst other things works in a distributed environment.

The contract for an Event Storage in code can be defined with the following interface.

```
public interface IEventStore {
  void SaveChanges(Guid AggregateId, int OriginatingVersion, IEnumerable<Event> events);
  IEnumerable<Event> GetEventsFor(Guid AggregateId);
}
```

Listing 3 Interface for an Event Store

Although not a trivial exercise to create a production quality Event Storage the overall concepts behind an Event Storage are relatively easy. Likely in the future there will be many off the shelf Event Storage systems available as either products or open source projects. There is however one very important optimization that was discussed in “Events as a Storage Mechanism” that really should exist in most systems and that is the concept of a “Rolling Snapshot”.

Rolling Snapshots

Rolling Snapshots are a heuristic to prevent the need to load all of the events when issuing a query to rebuild an Aggregate. They are a denormalization of the aggregate at a given point in time. In order to add the concept of a Rolling Snapshot to the basic Event Storage one will need to add a table and change the query logic slightly. Further discussion on Rolling Snapshots at a conceptual level can be found in the “Events as a Storage Mechanism” chapter.

Column Name	Column Type
fkAggregateId	Guid
SerializedData	Blob
Version	Int

Figure 12 Definition of Snapshot Table

The Snapshots table is relatively basic. It’s primary data in the blob that contains the serialized version of the aggregate at a given point in time. The serialized data could be in any one of a host of possible schemas, binary, XML, raw text, etc. The decision on how to serialize the snapshots is really dependent upon the system being built. A version number is included with the snapshot, it represents which version of the aggregate the snapshot represents.

In order to have snapshots being created one would need a process that handles the task of creating the snapshots. This process can live outside of the Application Server as a background process. There can be a single process running or many depending on needs due to throughput. All snapshots happen asynchronously. Figure 4 shows a conceptual architecture with a [SnapShotter] process introduced.

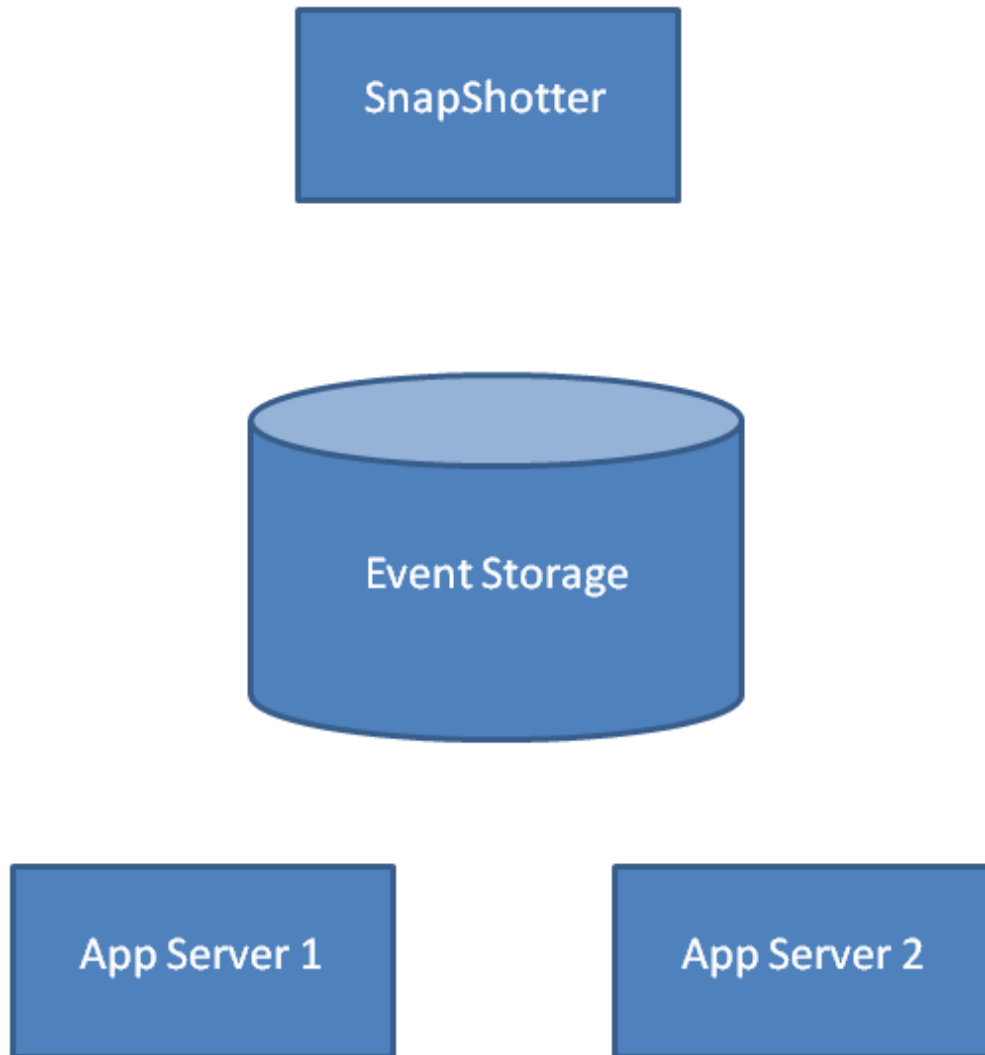


Figure 13 Introduction of a SnapShotter

The [SnapShotter] sits behind the Event Storage and periodically queries for any Aggregates that need to have a snapshot taken because they have gone past the allowed number of events. This query can be done quite easily in the simple Event Storage discussed by joining the Aggregates table to the Snapshots table on the Aggregate identifier. One would then subtract the current version from the last snapshot version with a where clause that only returned the aggregates with a value greater than some number. This query will return all of the Aggregates that a snapshot to be created. The snapshotter would then iterate through this list of Aggregates to create the snapshots (if using multiple snapshotters the competing consumer pattern works well here).

The process of creating a snapshot involves having the domain load up the current version of the Aggregate then take a snapshot of it. The creation of the snapshot can be done in many ways. Once the

snapshot has been taken, it is saved back to the snapshot table so that queries will have the snapshot available.

Many use the default serialization package available with their platform with good results though one should consider using the Memento pattern when dealing with snapshots. The reason why one would want to consider using the Memento pattern (or custom serialization) is that it is quite likely that over time the structure of the domain objects will change. If one is using the default serializer there is a versioning problem when the new structure is released (the existing snapshots must either be deleted and recreated or updated to match the new schema). The use of the Memento pattern allows the separated versioning of the snapshot schema from the domain object itself.

In “Events as a Storage Mechanism” a different, simpler mechanism was shown for the storage of snapshots. That system had the snapshots in line in the Event Log, this other mechanism although conceptually simpler has a few issues that can come up in a production system. The issues revolve around the need of ordering of the snapshot within the event log.

Consider that the Snapshotter has realized that an Aggregate Root needs to have a snapshot taken. It loads up the Aggregate and takes the snapshot. Unfortunately while it was doing this, one of the Application Servers made a change to the same Aggregate. As the snapshot is position dependent within the Event Log, it would receive an optimistic concurrency failure. The easy answer would be to simply repeat the process but what if it failed again? The snapshotter on a very busy Aggregate could end up in a situation where it would have a very low probability of actually writing the snapshot successfully.

By separating the snapshots into their own table and associating them to a version of the aggregate this problem is solved. Ordering of snapshots is not needed, the snapshot does not even need to be at the latest version, the snapshot that is taken is valid **at the version it was taken**.

Snapshots are a heuristic that will dramatically improve the performance of many systems, though not all systems need snapshotting. It is generally recommended to handle development without snapshotting as it can always be introduced later as a simple performance enhancement for the system.

Event Storage as a Queue

It has been previously discussed that the events coming out of a domain are also an [Integration Model]. Very often these events are not only saved but also published to queue where they are dispatched asynchronously to listeners either within the same system (the reporting model is a good example) or to other applications. An issue that exists with many systems publishing events is that they require a two-phase commit between whatever storage they are using (Relational or otherwise) and the publishing of their events to the queue.

The reason that the two-phase commit is needed is that a catastrophe could occur during the small period of time between when the write to the data storage commits and when the write to the queue commits. If a failure were to happen during this period the message would not be published on the queue (or if the other direction it may be published but the change may not be saved). If either case were to happen the listeners of the events would be out of sync with the producer.

The two-phase commit can be expensive but for low latency systems there is a larger problem when dealing with this situation. Generally the queue itself is persistent so the event becomes written on disk twice in the two-phase commit, once to the Event Storage and once to the persistent queue. Given for most systems having dual writes is not that important but if you have low latency requirements it can become quite an expensive operation as it will also force seeks on the disk. Figure 5 illustrates the two-phase commit between data storage and a publishing queue.

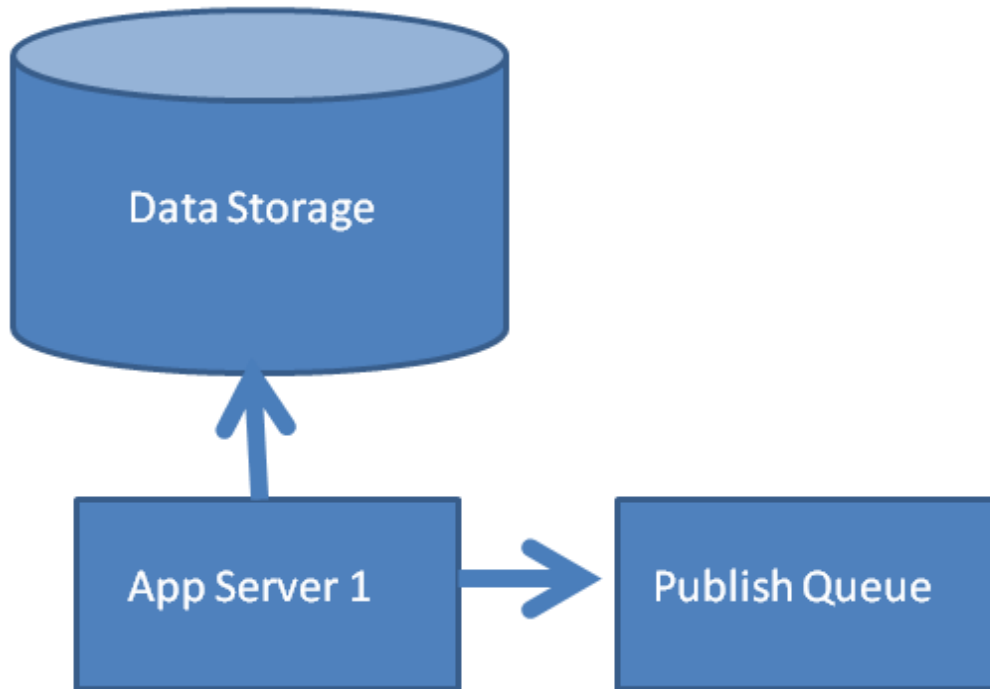


Figure 14 Two Phase Commit with Queue

Some try to get around this problem by only writing to a queue then have something on the other side of the queue update the data storage with the changes represented by the events, this however has some issues. The largest issue is that not all of the events will be able to be written to the storage, eventual consistency has been introduced and it is possible that an optimistic concurrency problem will occur on the write of the events. Dealing with this problem in a production system is non-trivial.

Instead one can do the opposite, use the event storage as a queue. If one were to add a sequence number to the Events table previously discussed one could also use the Event Storage as a queue. Figure 5 illustrates the change to the schema of the Events table.

Column Name	Column Type
fkAggregatId	Guid
Data	Blob

SequenceNumber	Long
Version	Int

Figure 15 Events Table as a Queue

The database would insure that the values of sequence number would be unique and incrementing, this can be easily done using an auto-incrementing type. Because the values are unique and incrementing a secondary process can chase the Events table, publishing the events off to ther queue. The chasing process would simply have to store the value of the sequence number of the last event it had processed, it could even update this value with a two-phase commit bringing the update and the publish to the queue into the same transaction. This process can be seen in Figure 7.




Figure 16 Event Storage as a Queue

The work has been taken off of the initial processing in a known safe way. The publish can happen asynchronously to the actual write. This allows one to lower the latency of completing the initial operation, it also will limit the number of disk writes in the processing of the initial request to one. This strategy can be extremely valuable when dealing with low latency requirements as it allows much of the work on the initial processing to be offloaded to another process asynchronously and in a safe way, there is little difference whether the publish happens as part of the initial processing or asynchronously as generally when publishing messages one is in an eventually consistent situation anyways, using the Event Store as a queue just raises the time until the message is published slightly.

TODO Put in source of a simple event storage.

User Do Things

This chapter introduces the concept of a Task Based User Interface and compares it with a CRUD style user interface. It also shows the changes that occur within the Application Server when a more task oriented style is applied to it's API.

One of the largest problems seen in "A Stereotypical Architecture" was that the intent of the user was lost. Because the client interacted by posting data-centric DTOs back and forth with the Application Server, the domain was unable to have any verbs in it. The domain had become a glorified abstraction of the data model. There were no behaviors, the behaviors that existed, existed in the client, on pieces of paper, or in the heads of the users of the software.

Many examples of such applications can be cited. Users have "work flow" information documented for them. Go to screen xyz edit foo to bar, then go to this other screen and edit xyz to abc. For many types of systems this type of workflow is fine. These systems are also generally low value in terms of the business. In an area that is sufficiently complex and high enough ROI in order to use Domain Driven Design these types of workflows become unwieldy.

One reason that is commonly cited for wanting to build a system such as described is that "the business logic and work flows can be changed at any time to anything without need of a change to the software". While this may be true one must ask at what cost. What happens when someone misses a step in the process they have in their head or you have multiple users who do it differently as is commonly the case? How do you get any reasonable information out of a system in terms of reporting?

One way of dealing with this issue is to move away from the DTO up/down architecture that was illustrated in a "Stereotypical Architecture". Figure 1 shows the client interaction side of a DTO up/down architecture.




Figure 17 Interaction in a DTO Up/Down Architecture

The basic explanation of the operation is that the UI will request a DTO, say for Customer 1234 from the Application Server. This DTO will be returned to the client and then shown on the screen. The user will interact with the DTO in some way (likely either directly or through a View Model). Eventually the client will click Save or some other trigger will occur and the client will take the DTO and send it back up to the Application Server. The Application Server will then internally map the data back to the domain model and save the changes returning a success or failure.

As discussed the intention of the user is being lost because a DTO is being sent up that just represents the current state of the object after the client's actions are completed. It is possible to bring forward the intention of the user; this will allow the Application Server to process behaviors as opposed to saving data. Figure shows an interaction capturing intent.

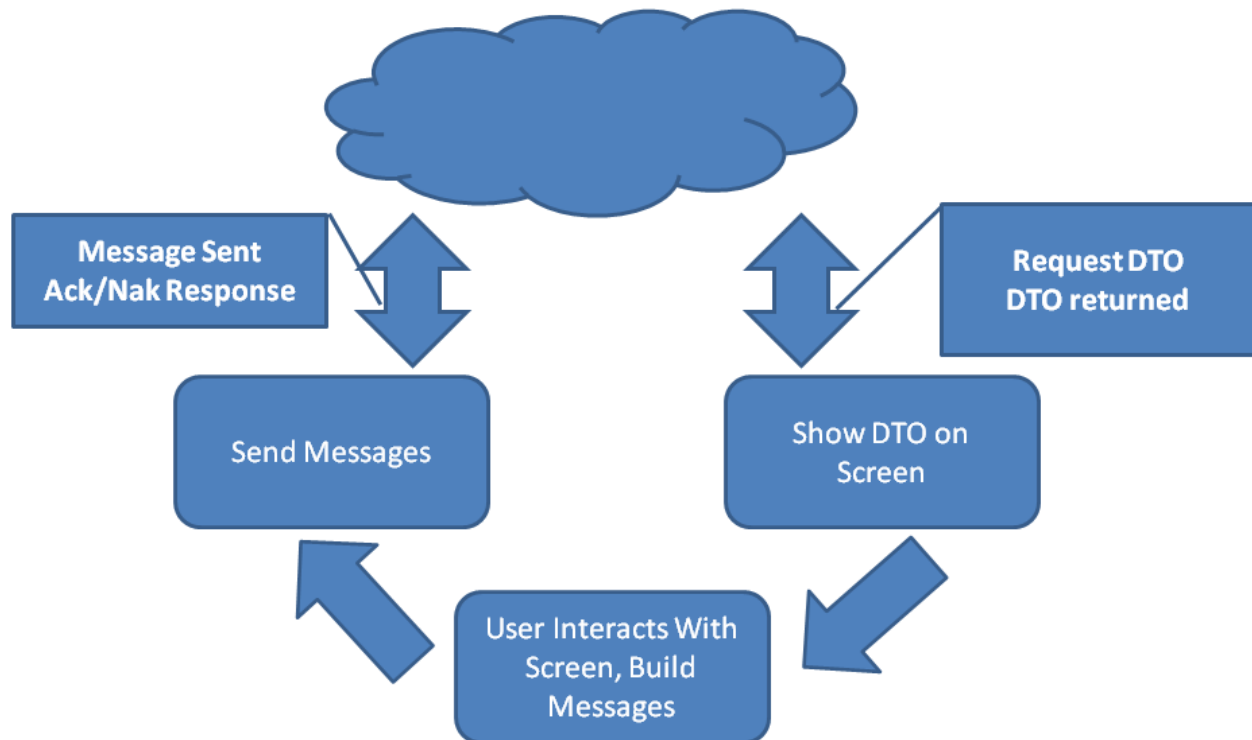


Figure 18 Behavioral Interface

Capturing intent the client interaction is very similar to the DTO up/down methodology in terms of interactions. The client first requests a DTO from the Application Server for instance Customer 1234. The Application Server returns a DTO representing the customer that is then shown on the screen for the user to interact with usually either directly or through a View Model. The similarities however stop at this point.

Instead of simply sending the same DTO back up when the user is completed with their action the client needs to send a message to the Application Server telling it to **do something**. It could be to “Complete a Sale”, “Approve a Purchase Order”, “Submit a Loan Application”. Said simply the client needs to send a message to the Application Server to have it complete the task that the user would like to complete. By telling the Application Server what the user would like to do, it is possible to know the intention of the user.

Commands

The method through which the Application Server will be told what to do is through the use of a Command. A command is a simple object with a name of an operation and the data required to perform that operation. Many think of Commands as being Serializable Method Calls. Listing 1 includes the code of a basic command.

```
public class DeactivateInventoryItemCommand {
    public readonly Guid InventoryItemId;
    public readononly string Comment;

    public DeactivateInventoryItemCommand (Guid id, string comment) {
        InventoryItemId = id;
        Comment = comment;
    }
}
```

Listing 4 A Simple Command

As a side note the example in Listing 1 includes the pattern name after the name of the Command. This is a decision that has many positives and negatives both linguistically and operationally. The choice over whether to use a pattern name in a class name is one that should not be taken lightly by a development team.

One important aspect of Commands is that they are always in the imperative tense; that is they are telling the Application Server to do something. The linguistics with Commands are important. A situation could for with a disconnected client where something has already happened such as a sale and could want to send up a “SaleOccurred” Command object. When analyzing this, is the domain allowed to say no that this thing did not happen? Placing Commands in the imperative tense linguistically shows that the Application Server is allowed to reject the Command, if it were not allowed to, it would be an Event for more information on this see “Events”.

Occasionally there exist funny examples of language in English. A perfect example of this would be “Purchase” which can be used either as a verb in the imperative or as a noun describing the result of its usage in the imperative. When dealing with these situations, ensure that the concept being pushed forward represents the imperative of the verb and not the noun. As an example a purchase should be including what to purchase and expecting the domain to possibly fill in some information like when the item was purchased as opposed to sending up a purchase DTO that fully describes the purchase.

The simple Command in Listing 1 includes two data properties. It includes an Id which represents the InventoryItem it will apply to and it includes a comment as to why the item is being deactivated. The comment is quite typical of an attribute associated with a Command, it is a piece of data that is required in order to process the behavior. There should only exist on a Command data points that are required to process the given behavior. This contrasts greatly with the typical architecture where the entire data of the object is passed back to the Application Server.

Most importantly of the data is the Id of the associated inventory item. At least one Id must exist for all commands that are updating state in some way, as all commands are intended to be routed to an object. When issuing a Create Command it is not necessary though valuable to include an Id. Having the client originate Ids normally in the form of UUIDs is extremely valuable in distributed systems.

It is quite common for developers to learn about Commands and to very quickly start creating Commands using vocabulary familiar to them such as “ChangeAddress”, “CreateUser”, or “DeleteClass”. This should be avoided as a default. Instead a team should be focused on what the use case really is.

Is it “ChangeAddress”? Is there a difference between “Correcting an Address” and “Relocating the Customer”? It likely will be if the domain in question is for a telephone company that sends the yellow pages to a customer when they move to a new location.

Is it “CreateUser” or is it “RegisterUser”? “DeleteClass” or “DeregisterStudent”. This process in naming can lead to great amounts of domain insight. To begin defining Commands, the best place to begin is in defining use cases, as generally a Command and a use case align.

It is also important to note that sometimes the only use case that exists for a portion of data is to “create”, “edit”, “update”, “change”, or “delete” it. All applications carry information that is simply supporting information. It is important though to not fall into the trap of mistaking places where there are use cases associated with intent for these CRUD only places.

Commands as a concept are not difficult but are different for many developers. Many developers see the creation of the Commands as a lot of work. If the creation of Commands is a bottleneck in the workflow, many of the ideas being discussed are likely being applied in an incorrect location.

User Interface

In order to build up Commands the User Interface will generally work a bit differently than in a DTO up/down system. Because the UI must build Command objects it needs to be designed in such a way that the user intent can be derived from the actions of the user.

The way to solve this is to lean more towards a “Task Based User Interface” also known as an “Inductive User Interface” in the Microsoft world. This style of UI is not by any means new and offers a quite different perspective on the design of user interfaces. Microsoft identified three major problems with Deductive UIs when researching Inductive UIs.

Users don't seem to construct an adequate mental model of the product. *The interface design for most current software products assumes that users will understand a conceptual model that the designers carefully crafted. Unfortunately, most users don't seem to ever acquire a mental model that is thorough and accurate enough to guide their navigation. These users aren't dumb — they are just very busy and overloaded with information. They do not have the time, energy, or desire to wonder about a conceptual model for their software.*

Even many long-time users never master common procedures. *Designers know that new users may have trouble at first, but expect these problems to vanish as users learn common tasks. Usability data indicates this often doesn't happen. In one study, researchers set up automated equipment to videotape users at home. The tapes showed that users focusing on the task at hand do not necessarily notice the procedure they are following and do not learn from the experience. The next time users perform the same operation, they may stumble through it in exactly the same way.*

Users must work hard to figure out each feature or screen. Most software products are designed for (the few) users who understand its conceptual model and have mastered common procedures. For the majority of customers, each feature or procedure is a frustrating, unwanted puzzle. Users might assume these puzzles are an unavoidable cost of using computers, but they would certainly be happier without this burden. (Microsoft Corporation, 2001)

The basic idea behind a Task Based or Inductive UI is that its important to figure out how the users want to use the software and to make it guide them through those processes.

Many commercial software applications include user interfaces in which a screen presents a set of controls, but leaves it to the user to deduce the page's purpose and how to use the controls to accomplish that purpose. (Microsoft Corporation, 2001)

The goal is to guide the user through the process. An example of the differences can be seen in the DeactivateInventoryItem example previously shown. A typical deductive UI might have an editable data grid containing all of the inventory items. It would have editable fields for various data and perhaps a drop down for the status of the inventory item, deactivated being one of them. In order to deactivate an inventory item the user would have to go to the item in the grid, type in a comment as to why they were deactivating it and then change the drop down to the status of deactivated. A similar example could be where you click to a screen to edit an inventory item but go through the same process as seen in Figure 3.

Figure 19 A CRUD screen for an Inventory Item

If the user attempts to submit an item that is “deactivated” and has not entered a comment they will receive an error saying that they must enter a comment as it is a mandatory field for a deactivated item. Some UIs might be a bit more user friendly, they may not show the comment field until the user selects deactivated from the drop down at which point it would appear on the screen. This is far more intuitive to the user as it is a cue that they should be putting data in that field but one can do even better.

Name ▲	Supplier	Active	
Super Fun Club Membership	ACME	<input checked="" type="checkbox"/>	Deactivate
Norwegian Salmon	ACME	<input type="checkbox"/>	
Giant Bean Bag	Walmart	<input checked="" type="checkbox"/>	Deactivate
Watermelons	Patata	<input checked="" type="checkbox"/>	Deactivate
TShirts	Company Corp	<input checked="" type="checkbox"/>	Deactivate

Figure 20 Listing Screen with Link

A Task Based UI would take a different approach, likely it would show a list of inventory items, next to an inventory item there might be a link to “deactivate” the item as seen in Figure 4. This link would take them to a screen that would then ask them for a comment as to why they are deactivating the items which is shown in Figure 5. The intent of the user is clear in this case and the software is guiding them through the process of deactivating an inventory item. It is also very easy to build Commands representing the user’s intentions with this style of interface.

Deactivate Inventory Item

A comment is required explaining why you are deactivating the inventory item.

Cancel
Deactivate

Figure 21 Deactivating an Inventory Item

Web, Mobile, and especially Mac UIs have been trending towards the direction of being task based. The UI guides you through a process and offers you contextually sensitive guidance pushing you in the right direction. This is largely due to the style offering the capability of a much better user experience. **There is a solid focus on how and why the user is using the software; the user’s experience becomes an integral part of the process.** Beyond this there is also value on focusing more in general on how the user wants to use the software; this is a great first step in defining some of the verbs of the domain.

Works Cited

Microsoft Corporation. (2001, Feb 9). *Microsoft Inductive User Interface Guidelines*. Retrieved from MSDN: <http://msdn.microsoft.com/en-us/library/ms997506>

More information on the DDDD/CQRS, article and source code references and updates: <http://abdullin.com/cqrs/>