



WHITEPAPER

THE TOP SIX MICROSERVICES PATTERNS



CONTENTS

CONTENTS	2
HOW TO CHOOSE THE RIGHT ARCHITECTURE FOR YOUR ORGANIZATION	3
INTRODUCTION	5
THE 6 MICROSERVICES PATTERNS	8
SECTION I: INTRODUCTORY PATTERNS	10
Fine-grained SOA.....	10
Layered APIs over fine-grained SOA	16
SECTION II: MANAGED STATE PATTERNS	21
Message-oriented state management over layered APIs.....	22
Event-driven state management over layered APIs.....	27
Isolating state in layered APIs	32
Replicating state in layered APIs (event sourcing)	38
SECTION III: FOUNDATIONAL BEST PRACTICES FOR ESTABLISHING THE MICROSERVICES PATTERNS	43
CONCLUSION: ESTABLISH A MICROSERVICES PATTERN THAT'S RIGHT FOR YOU	50
ABOUT MULESOFT	51

HOW TO CHOOSE THE RIGHT ARCHITECTURE FOR YOUR ORGANIZATION

For the last several years, microservices has been a buzzword in the IT press. Technology consulting firm Thoughtworks has declared that “a microservices architecture as a programming model” is a rising trend¹, whereas ZDNet said it was a “technology to watch.”² The press is expressing its endorsement of microservices, which might make architects and IT executives feel a fear of missing out on the next exciting trend.

The problem with the supposed imperative to adopt microservices is that there are many people who feel that it is a prescriptive architecture; it must be done one certain way — like Netflix, for example — or it simply can’t be done. But adopting microservices this way is not feasible for many organizations and can lead to failure. For organizations that have particular structures and cultures, a purist view of microservices can only go so far because of various legal, technological, and cultural constraints. Organizations will fail if they follow an overly prescriptive point of view in the microservices space if their needs are not compatible with the purist approach.

For example, MuleSoft works with a global food and beverage provider, who, because of the nature of their business, will not do 100 updates a day like a Netflix might. At most, they will do an update every 10 days,

1 Krill, Paul. “[Docker, machine learning are top tech trends for 2017.](#)” November 7, 2016.

2 Hinchcliffe, Dion. “[The enterprise technologies to watch in 2016.](#)” May 29, 2016

because of the different nature of their business. So why do they have to deploy using the purist approach when they don't have the same needs? Large enterprises with a lot of legacy systems have different needs than younger companies, so their culture might not fit with a purist approach to microservices either. Finally, some enterprises, particularly in highly regulated industries, have different security or compliance needs than a Netflix or a Spotify. In that case, constant software updates might not be feasible or even legal, further necessitating a departure from “pure” microservices.

In reality, trying to adapt to a microservices architecture needs to be done pragmatically, in a way that makes sense for your organization. Not every company has to be a Netflix or a Spotify. You can adopt microservices in a way that meshes with your culture, your goals, and your own organization.

Instead of adopting microservices as a singular approach — which would defy the point of the architecture — we propose considering microservices as a series of overlapping patterns that you can pick and choose depending on your organization's needs. There's no need to adopt every aspect of microservices all at once; it can be adopted in a pragmatic way in a manner that makes sense for your organization.

INTRODUCTION

In 1987, chef Ferran Adrià heard the seminal phrase “creativity means not copying”³ this simple, self-evident statement drove a major movement within the culinary world. We can state a similarly simple, yet potentially profound statement with microservices: “microservices are not a monolith.”

Originally created as an alternative to a monolith, microservice patterns should inherently be non-monolithic: change is easy, units are small, scalability is semi-infinite. This also means that microservices are not just one thing. Rather, we posit that microservices are a category of grouped, related patterns that share the same set of goals. This is analogous to database systems; they all share similar goals — typically with different priorities like scalability or maintainability. Yet, their specifics differ significantly. For example, RDBMSes, NoSQL Stores, Time Series Databases, and Big Data Stores, etc. all share a similarity — they offer the ability to store and query data — yet the specifics of their individual trade-offs couldn't be more different.

This whitepaper will highlight a taxonomy of microservices patterns that have been practiced in the wild. Each of these patterns cannot be seen as better or worse than another, but rather as choosing a specific set of trade-offs that prioritize different things along the way. They all obey the microservices architectural goals (speed, scalability, cohesion)

3 Adrià, Ferran. [The History of El Bulli](#).

but arrive at them in different ways. It can even be argued that some of the patterns are “steps along the journey,” where they move an architecture forward towards the microservices goals. In and of themselves they tend to produce predictable failure conditions that then encourage a team working in this way to seek alternatives, to make harder trade-offs, and then to evolve to a different, more specialized pattern based on their growing experience.

This is certainly how microservices architecture came to be. It started with a conceptual approach and set of goals that resulted in a first iteration, then it was iterated relentlessly towards the more specialized (many would view as more “extreme”) patterns. This is not a bad thing. Organizations that are implementing or living with these early-stage patterns have not done something retrograde, but rather have improved their current state to the stage where they now can clearly see the benefits of taking the next step. This is one way that microservices architectures are very “real-world”; they follow the pattern of human learning. Implementers (at least for now) are all humans and need to learn the value of more specific patterns and trade-offs through personal experience, rather than by slavishly following any particular text or blog post.

Microservice architectures should be evolutionary. Because they are defined by an ability to change rapidly, the approach to building a system develops over many iterations. It is relatively uncommon in most organizations to have the luxury of implementing a microservices architecture from the ground up (and many consider it an anti-pattern to do so). As such, there will usually be a continuum of architectures,

patterns, and technologies in play based on maturity, need, and timing. Some of this is driven by a team's ability to absorb change, some of it depends on setting up enabling infrastructure, and some of it depends on organizational change and restructuring teams around, for example, DevOps practices. The ability to focus on what can be done right now while enabling teams to keep evolving towards a more productive state is one of the key strengths of microservice architectures.

THE 6 MICROSERVICES PATTERNS

Below is a taxonomy of the 6 microservices patterns available to make adopting this architecture an easier task for your organization. Each of these patterns is not a general pattern; rather, each organization can choose to make a set of trade-offs that prioritize particular things over others. It is our hope that users who implement a particular pattern will use this as a reference in order to plan their architecture following these patterns, rather than attempting to over-apply any single pattern and therefore break it through compromise. The intended state is to design and implement a microservices architecture using a mixture of these patterns, rather than choose one and migrate towards it.

No single one of these patterns is better than the others, but each is uniquely better at a particular task. We attempt to make it clear what that task is when describing each pattern.

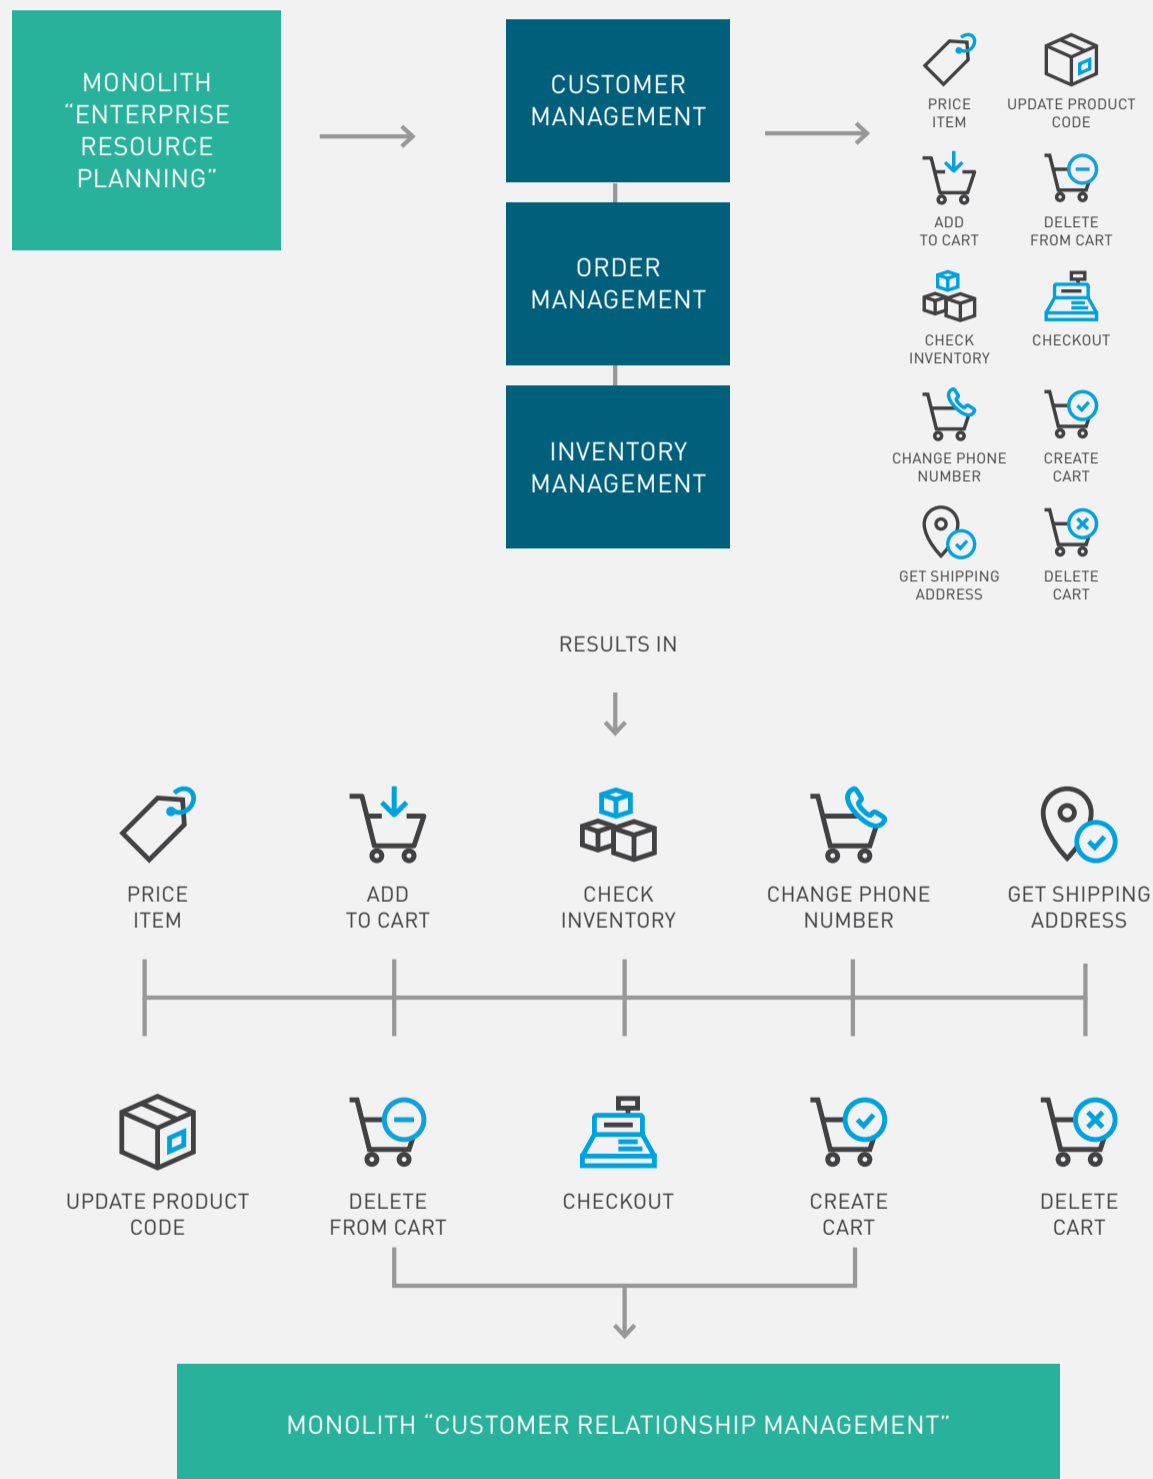
As a final note, we should emphasize that microservices architecture is not appropriate for every use case. If you depend heavily on an ERP application, for example, and have no intention to remove it, microservices probably aren't for you. It's unlikely, however, that there aren't parts of your business that would benefit from speed, scale, and cohesion. Those are the areas where you should look to take advantage of this rapidly growing ecosystem.

These patterns don't apply specifically to any one size or type of organization. It's true that the patterns that address state management

specifically tend to be more useful to large-scale, fast-moving companies. However, they can be applied to any kind of business or company that has an appropriate need and is willing to balance the needs of their architecture correctly. In other words, “where you stop” is entirely dictated by the problems you’re trying to solve, rather than any linear progression through these patterns.

SECTION I: INTRODUCTORY PATTERNS

FINE-GRAINED SOA



Fine-grained SOA is arguably the “big bang” of microservices. To many, Netflix is the origin of this style of architecture, and this was their own

articulation of the pattern at the beginning. Fine-grained SOA is somewhat self-explanatory (at least, to SOA practitioners); it reduces the issues experienced with SOA, and applies the same principles, but breaks down the design into smaller, more fine-grained pieces. Small services are, in fact, what most people still think of as the only microservices pattern.

However, along the way, Netflix and others following this path have experienced a number of basic issues with this pattern. When you make things smaller, and attempt to scale them, some difficulties emerge:

Where you used to make a single network call, now you must make tens or even hundreds of calls. This is inefficient.⁴

Where you used to manage a small number of things, now you manage hundreds or thousands or more. Your management tools therefore no longer work as well for you.

When tens, hundreds, or thousands of things all can query or modify the state of your central applications, it becomes almost impossible to trust the consistency of that store as accessed through any one of these microservices.

The overlooked challenges of truly living with microservices include inefficient inter-process communications, overarching monitoring, management, and governance; and consistent state.

4 Note that high efficiency inter-process communications frameworks are very available now (gRPC, Avro, Hystrix, Fingert) so this isn't a long-term concern. Rather, it should be something to flag. If you plan to do microservices, raw HTTP/JSON or XML will eventually become too inefficient for the pattern at high scale.

In most cases, this pattern is applied as an extension of service-oriented integration, where the point of each service is to provide connectivity to external systems. This forms tight dependencies to those external stores, making speed of change drag, and making the cohesion of the system reflect the internal state of those applications.

When you begin to implement fine-grained SOA, be aware of what you're getting into. These pain points will emerge if you succeed, and you should be ready for them. When the pain becomes significant enough, you will inevitably need to seek out other patterns.

PROBLEM:

Coarse-grained services are too difficult to change without side effects, and the monolithic nature of traditional architecture is “holding the teams back.”

SOLUTION:

Break up services into finer-grained pieces, which reduces the scope of any given change, allowing it to happen more easily.

APPLICATION:

Services are broken up into finer-grained microservices, typically each with a single purpose.

Impacts:

- ◆ Traffic increases.
 - ◆ The number of services managed becomes large.
 - ◆ Traditional monitoring solutions become insufficient.
 - ◆ The automation of integration, testing, and deployment become critical.
 - ◆ The orchestration of microservices becomes necessary.
 - ◆ The ability to change rapidly is improved.
-

GOALS:

Speed of change.

Scalability: You can theoretically improve scalability, but it practically tends to decrease unless supporting automation infrastructure is put in place.

KEY CHARACTERISTICS:

It works well at low scale, but pain emerges at high scale.

The focus is on small services, but “-ilities” (e.g. scalability, reliability, etc.) typically have not been considered.

It tends to be integration-centric, with each microservice depending on external systems.⁵

There is inefficient inter-process communication in order to achieve high speed of change.

Data consistency and state management is poor⁶ but allows existing systems to be leveraged.

Due to similarity to existing patterns, existing problems tend to happen.

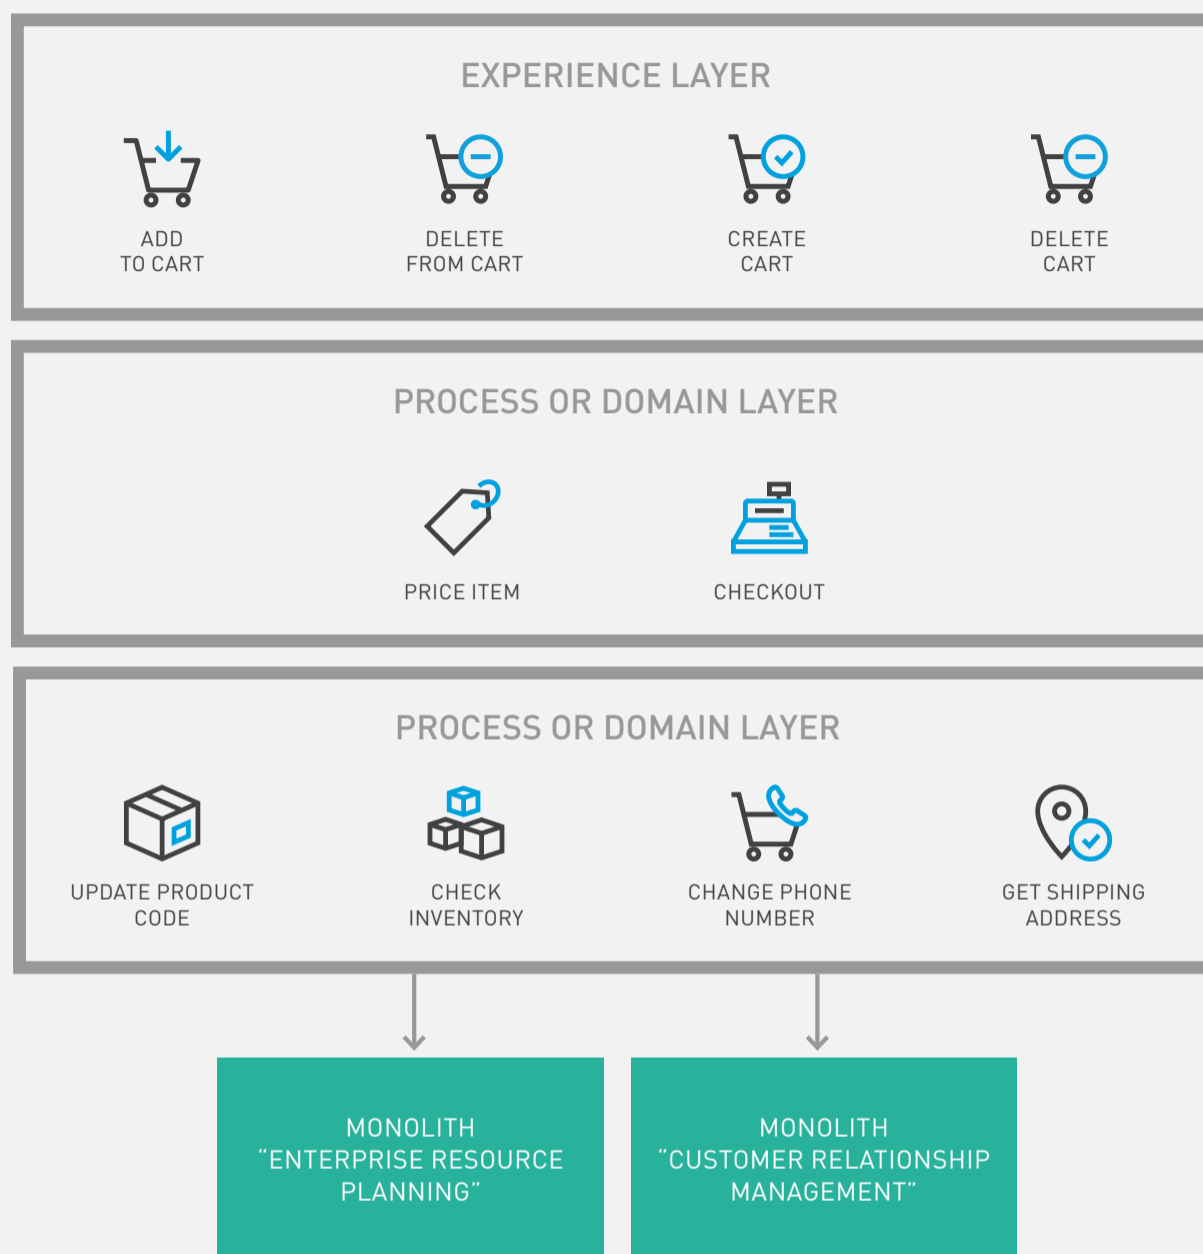
HOW DOES THIS COEXIST WITH EXISTING SYSTEMS, SOA, OR APIs?

This pattern is very similar to SOA and API-led approaches and leverages more value out of existing systems. The coexistence model is simple and, as such, it causes most of the friction with this model since the low speed of change on existing assets can cause a drag on the microservices approach.

5 Note that, generally speaking, microservices and APIs in general are increasingly treated as “external” assets—whether they are or not. This is what we at MuleSoft often refer to as “productized” APIs, where, regardless of context, every point of interconnectivity is treated as something that is robust and ready to be used across any permissible context.

6 There are multiple data masters, but the architecture doesn’t typically address the existence of a source of truth or use a specific consistency model. There are often multiple ways to make a change to data that can have cascading impact, which is difficult to understand or respond to.

LAYERED APIs OVER FINE-GRAINED SOA



MuleSoft advocates an approach termed “API-led connectivity.” Simply speaking, this can be thought of as a layered approach to API design (at least for the purposes of this paper). System APIs expose systems, Process APIs orchestrate them, and Experience APIs provide end-user experiences. This approach is well-aligned with fine-grained SOA and, often, the two either can co-exist or else fine-grained, layered APIs become an evolutionary pattern that follows fine-grained SOA.

This approach gives some structure to a fine-grained API approach, allowing some ability to consistently manage and reason with the APIs or

microservices. However, there are some deep issues with this approach that are similar to fine-grained SOA (which are intuitive):

Where fine-grained SOA makes a single network call, now you must make multiple calls through the layers. This can be inefficient from the perspective of “network hops,” however, the existence of layers does not mandate them. Calls directly across layers are completely valid; the goal of layering is to increase flexibility while also structuring the architecture in a way that separates concerns well.

Where fine-grained SOA manages a large number of things, with layered APIs you manage multiple layers of a large number of fine-grained things. Your management tools no longer work as well as they used to, as they may not have ways of visualizing complex microservices views.

End-application datastore consistency is actually improved because the set of things that modify or query them are organized and focused (i.e. System APIs).

Ultimately, this a good pattern for most enterprises, but, like fine-grained SOA, there will be pain along the way. However, this pain exhibits mostly at large scale, so one should only plan for other patterns when high scale is expected or experienced.

PROBLEM:

Microservices architectures without some amount of structure are difficult to rationalize and reason with, as there is no obvious way to categorize and visualize the purpose of each microservice.

SOLUTION:

By creating layers of microservices that are grouped by purpose (systems, processes or domain models, and experience), you can manage the complexity of the architecture more easily.

APPLICATION:

Microservices are categorized into layers. Often, standards can be put in place to make microservices that have similar purposes behave in similar ways, which further rationalizes complexity.

IMPACTS:

The ability to change rapidly is improved through standardization and further decomposition.

The number of inter-process calls can be increased because of the more specialized microservices structures.

Monitoring and visualization tools may need to be reviewed for their ability to work with the layered structure correctly.

GOALS

Speed of change.

Scalability: it can theoretically improve scalability, but practically tends to decrease unless supporting automation infrastructure is put in place.

KEY TRADE-OFFS:

There is inefficient IPC in order to achieve high speed of change.

Data consistency and state management is poor but allows high degrees of reuse. Reuse itself trades off against speed of change.

Due to similarity to existing patterns, existing problems tend to emerge.

It works well at low scale, pain emerges at high scale.

There is high cohesion due to structured architectural approach.

The focus is on the small services, but “-ilities” typically have not been considered.

It tends to be integration-centric, with each system microservice depending on external systems. This reduces speed of change.

HOW DOES THIS PATTERN COEXIST WITH EXISTING SYSTEMS LIKE SOA OR APIs?

This approach tends to be the best way to coexist with existing assets. Because the layering reduces the scope of each microservice as well as focuses its purpose, it is able to connect and leverage existing systems best without causing significant slowdown.

However, coordinating changes that ripple through a fine-grained and also layered design can be challenging to work with. You may need to employ technologies to manage the contracts between all the different pieces or employ thorough automated testing technology to ensure that changes don't break things.

SECTION II: MANAGED STATE PATTERNS


The following patterns are all focused on managing state. State is ultimately one of the most challenging aspects to a distributed architecture, because traditional system design favors consistent data queries and mutations even though consistency is difficult (if not arguably impossible) in a sufficiently distributed architecture.

Because many of our customers' microservice designs are focused on integration use cases, the issue of managing state comes front and center after even a small amount of success. This is because providing connectivity or integration means that the microservices system is inherently either querying or mutating state (or both). They may not be the only way to do so for a given entity or piece of information. This means that in order to avoid data corruption or unexpected results, you need to consider one of two strategies:

Explicitly declare state and use a strategy to deal with the side effects of mutating and querying it.

By doing so, you can achieve a greater level of physical autonomy for each component, allowing a faster rate of change.

MESSAGE-ORIENTED STATE MANAGEMENT OVER LAYERED APIs



This is usually the first pattern implemented as a way to avoid the side effects of accessing and mutating state. By providing an asynchronous queue as the primary mechanism to communicate state changes (by command or event) or to query other microservices, we allow each microservice the time necessary to converge events and therefore provide a consistent external view.

This pattern is commonly used by teams experienced with SOA and ESBs, as it is an intuitive path to follow given their prior experience.

We recommend that this pattern is used as a transitory state; this should be used by an organization beginning the microservices journey, but it may not suit your needs as your approach becomes more mature. By decoupling components temporally using a queue, the implementation and behavior of each microservice becomes obscured, meaning that often the side effects of the design become even more prominent than when simply exposing existing systems to a microservices design. It is not uncommon, for example, to see messages used to propagate events, commands, batches, or even to just stream data, all appearing at a high level to be the same thing. This can make systems unpredictable, which in more traditional environments (e.g. ESBs) is manageable, but when moving to microservices scale (thousands of services, hundreds of thousands of message types), it becomes rapidly impractical to reason with. This pattern is therefore often a transitory step, where success translates into teams realizing that at a basic level queueing is useful, but in order to really succeed they need to establish some standards around what is passed over the queues.

PROBLEM:

In order to ensure data integrity, there is a need to replicate the state of key business data between microservices or data stores.

SOLUTION:

Using a message queue allows state to be asynchronously and reliably sent to different locations.

APPLICATION:

When a change in data occurs, it is sent as a message over a queue or ESB to any other microservice or store that needs to be notified of the change.

IMPACTS:

It can increase complexity as it provides a new way for state to change and move.

It does not offer any standard patterns, so the implementation can be inconsistent unless standards are agreed and applied.

It does not offer any specific opinions as to how to deal with data conflicts or to rebuild state in the case of failures or outages.

GOALS:

Scalability: using message queues provides the means to scale task processors independently from task producers, with a reliability layer between them.

KEY TRADE-OFFS:

There is an inefficient IPC due to asynchronicity.

Data consistency and state management are made worse by the unpredictability of behavior and the potential side effects of a given message. Reuse is actually hampered by the inability to predict the use patterns.

Due to similarity to existing patterns, existing problems tend to emerge.

It works well at high scale technically but tends to become operationally unpredictable.

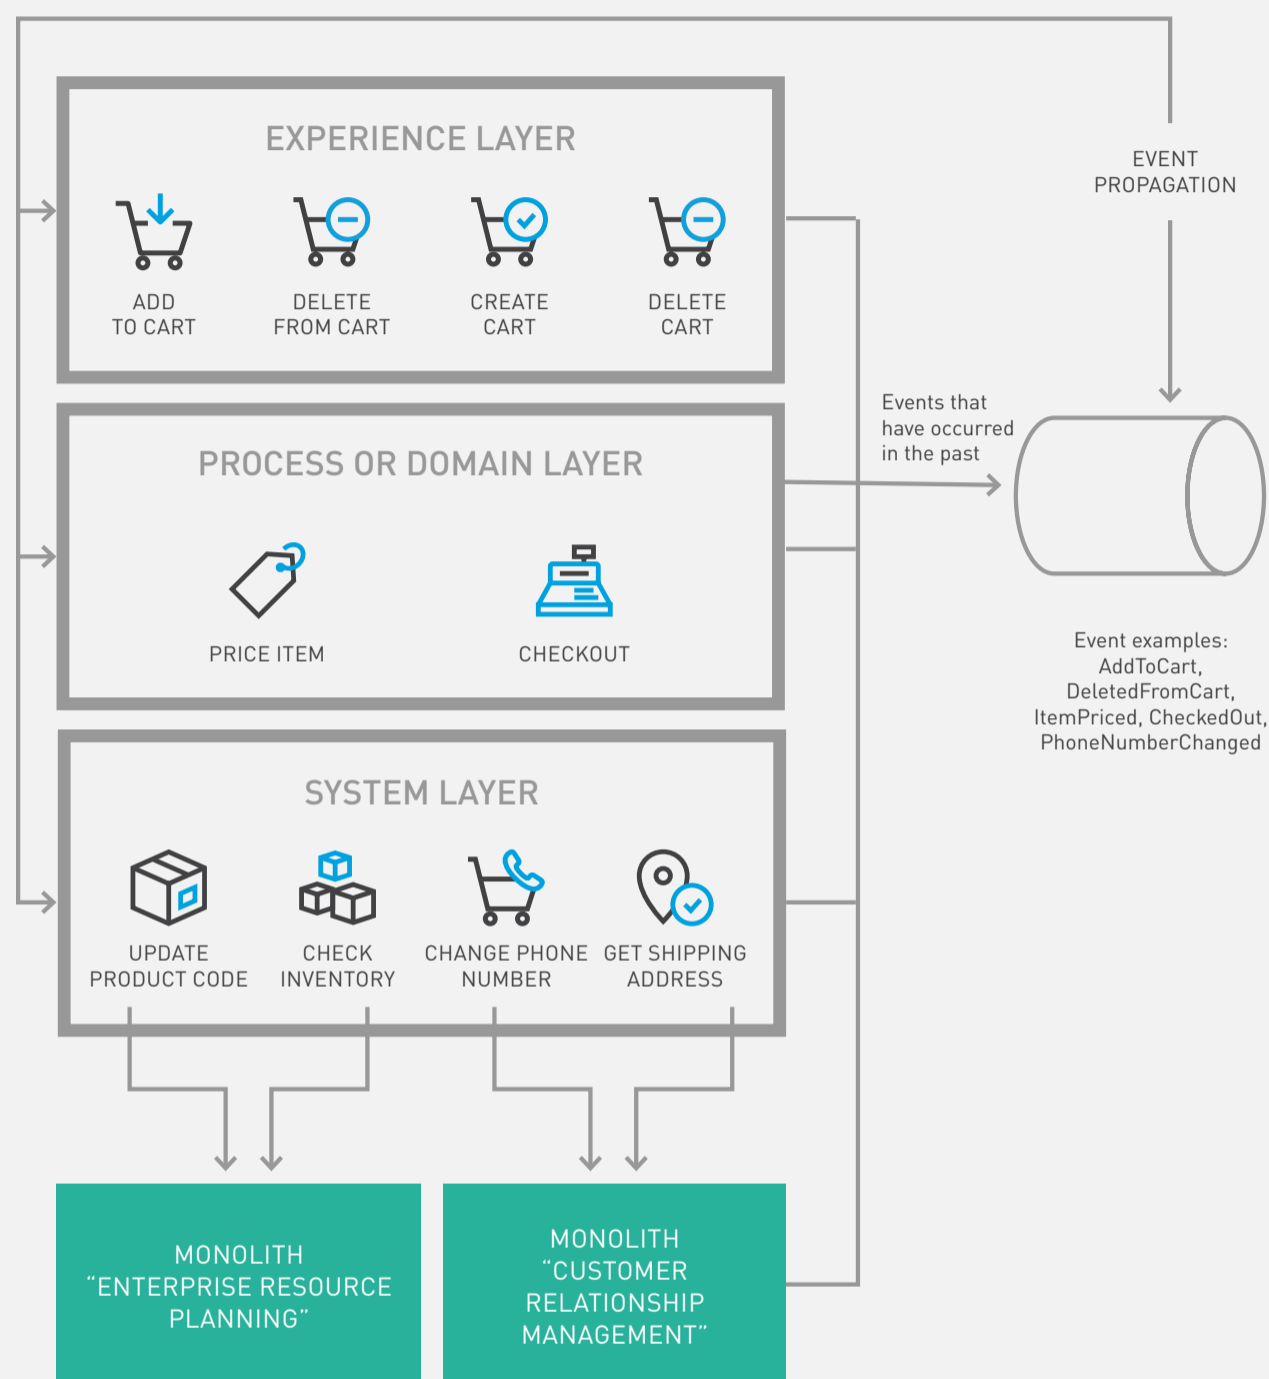
There is poor cohesion due to lack of design standards.

HOW DOES THIS CO-EXIST WITH EXISTING SYSTEMS, SOA, OR APIs?

This approach may be the best way for existing assets and connectivity to coexist with minimal change. By decoupling at the interface level using messages, you gain a great deal of flexibility and reliability. The ability to transform and route messages in the message layer means that changes can be kept isolated to their source, rather than happening at the interface level.

This increased flexibility comes with a lack of visibility and may require new management tools and troubleshooting techniques. For example, a common step is to add tracking data into message headers to understand the path a given message has taken, for troubleshooting and debugging purposes.

EVENT-DRIVEN STATE MANAGEMENT OVER LAYERED APIs



Event-driven architectures are nothing new. Mule ESB, for example, was originally designed as an event-driven system. But when overlaid on microservice patterns they provide some powerful abstractions. Event-driven systems usually use a queue of some kind (like message-oriented systems) but enforce a standard around the design and behavior of what is passed over the queue; specifically, the concept of an event.

People confuse this pattern with other patterns and, as a result, it covers a wide array of designs. Strictly, an event is something that occurred in the past with an associated representative state and timestamp. This event allows any service receiving it to reconstruct a materialized view of the state by replaying the events in order. However, in many implementations, the concept is muddied, where events (e.g. something happened) are mixed with commands (e.g. make something happen) and without the distinction, the predictability of the design is flawed. That said, this approach is undeniably better than message-orientation (due to its more specific design), but tends to have problems in implementation due to a lack of consistency. Teams that articulate and enforce a consistent standard will find this pattern tends to work very well in microservices architectures.

PROBLEM:

In order to ensure data integrity, there is a need to replicate key business events to synchronize between microservices or data stores.

SOLUTION:

Use a common event abstraction to represent the unit of change in the architecture.

APPLICATION:

When something changes in the business, an event encapsulating it in the past tense is sent to interested parties. Changes in the business are the product of these events being sent and processed.

IMPACTS:

It can increase complexity as it provides a new way for state to change and move.

It does not offer any standard patterns, so the implementation can be inconsistent unless standards are agreed and applied.

It does not offer any specific opinions as to how to deal with data conflicts or to rebuild state in the case of failures or outages.

GOALS:

Cohesion: this architecture is very easy to work with and understand due to its standardized nature.

Scalability: it requires deeper technical decisions (how do you send/process/store events? What about retransmission?) but is achievable.

Speed of change: this will occur due to the more cohesive architecture, but without dependency analysis, the tooling relies a little too much on tribal knowledge and luck.

KEY TRADE-OFFS:

There is efficient IPC due to asynchronicity.

The flexibility of design is lost in favor of predictable behavior.

Data consistency and state management are improved through a specific model of consistency; any given state is merely a reconstruction of events.

Due to similarity to message-orientation, confusion can occur where events are mixed up with commands.

There is an effective scaling model with reasonable operational oversight.

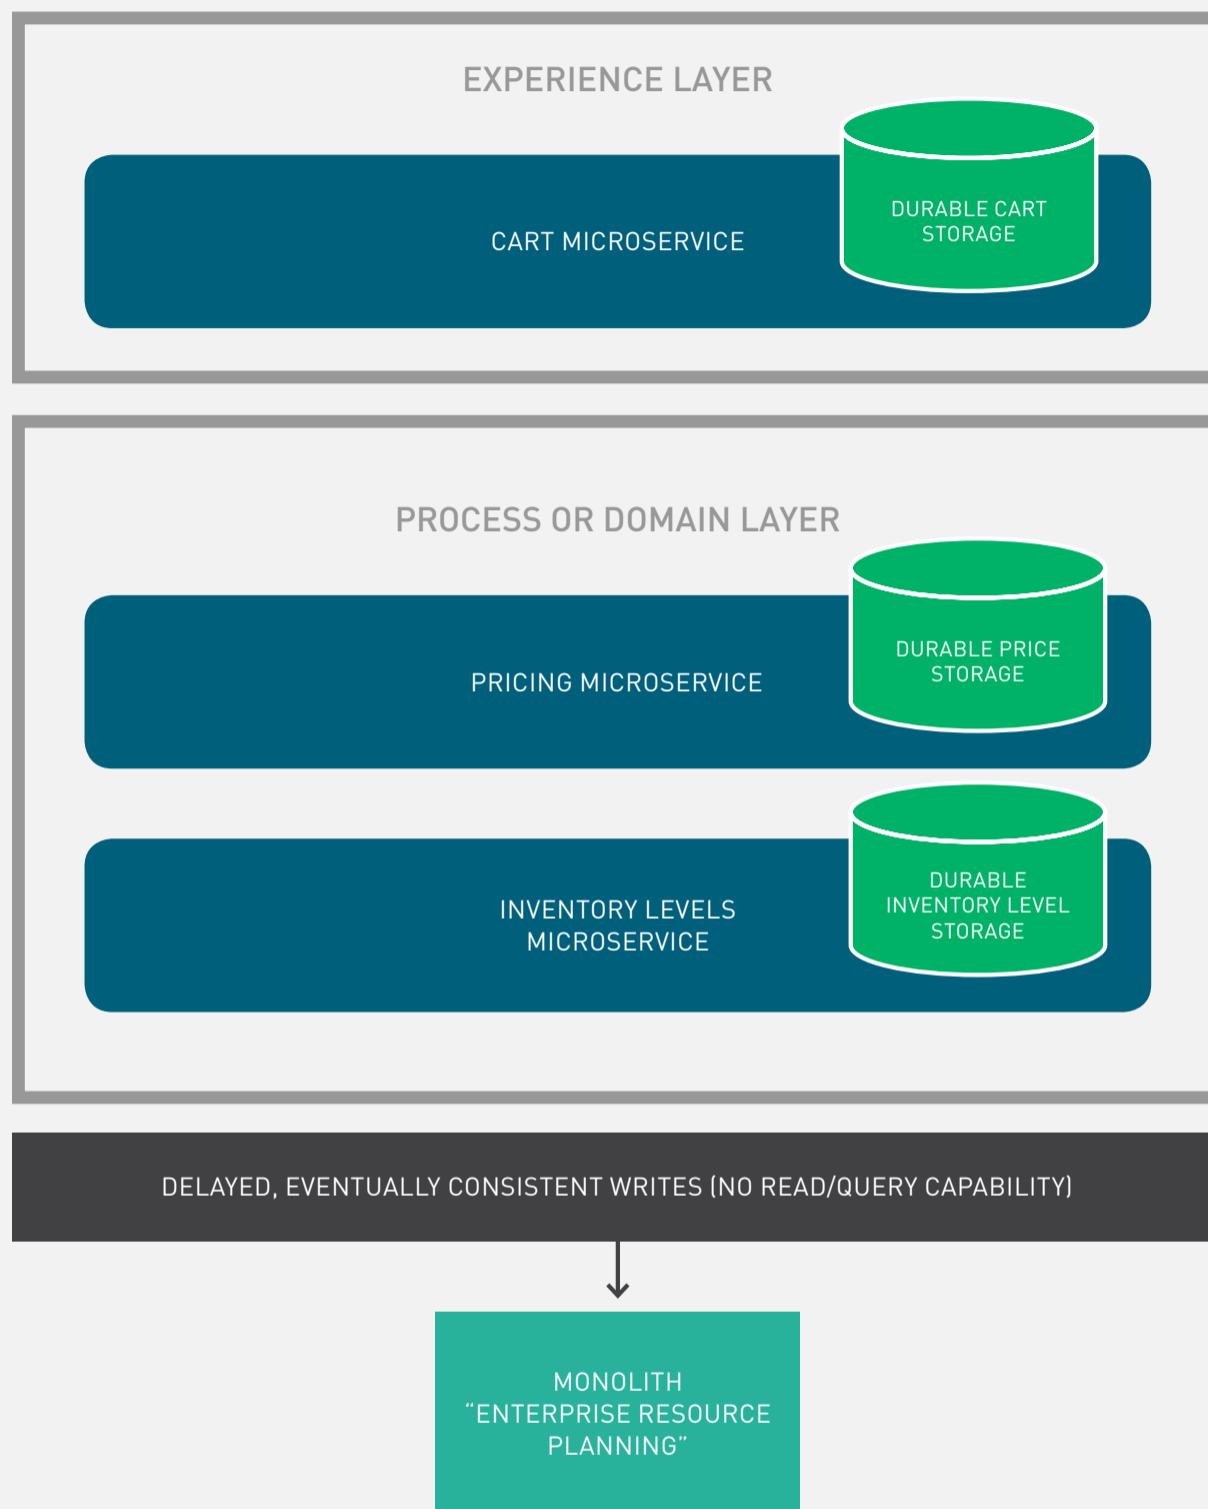
There is strong cohesion when applied consistently, but cohesion tends to drift over time.

HOW DOES THIS COEXIST WITH EXISTING SYSTEMS, SOA, OR APIs?

Event-driven systems can coexist with existing systems, but they tend to require a “translation layer” across the boundary of the event-driven parts of the architecture with those that aren’t event-driven. In essence, the event-driven system talks a consistent language internally, and anything on the outside needs to be converted (in or out) to participate.

This makes for a clean way to separate the event-driven parts of the architecture from traditional integration and enterprise systems, but it does mean that you tend to create “new” functionality with event-driven microservices that update out of band or sync with the systems and APIs outside of the boundary.

ISOLATING STATE IN LAYERED APIs



An alternative to coalescing the exchange pattern of a microservices architecture (for example, into events) is to coalesce the internal consistency of each microservice. Rather than expect consistency in the interchange, expect consistency at the time of query.

This is done most commonly by isolating state or, in other words, “each microservice contains its own state.” In this pattern, each microservice contains an internal data store that it constantly reconciles with external stores (be they an event log or an enterprise asset) so that it becomes the “single source of truth.” This can be difficult because single source of truth patterns tend to echo the complexity of master data management and its associated challenges.

However, using an external store as a single source of truth with microservices is much more practical, because of the typically single-purpose nature of a given microservice. It’s difficult to isolate the state of a customer, for example, but it’s not so difficult to isolate the state of the customer’s email address. This pattern, therefore must, by design, favor extremely granular microservices in order to succeed. It also typically requires asynchronous event propagation as a means to pass state change from one point to another. This pattern can also be thought of as something of a “distributed database,” with each microservice almost representing a column in a traditional RDBMS design.

Microservices contain a data store that is the source of truth for the entity they represent. For example, a “product” microservice could contain a MySQL database that contains all information about the product and is the only way to query or update that concept in the organization.

Unlike more SOA-oriented patterns, reuse doesn’t tend to be a priority in these designs. Each microservice has a “use,” certainly, often accessed from different contexts, but each microservice isn’t designed

with reuse in mind. If reuse occurs, it's accidental, not with the intention that comes with SOA.⁷

PROBLEM:

It is difficult to achieve data integrity when there are multiple sources of truth.

SOLUTION:

Nominate a microservice that represents the single source of truth for each given business entity, and encapsulate the state inside the microservice.

APPLICATION:

Microservices contain a data store that is the source of truth for the entity they represent. For example, a “product” microservice could contain a MySQL database that contains all information about the product, and is the only way to query or update that concept in the organization.

IMPACTS:

It requires some governance to ensure that data is not copied or has other modes of access.

⁷ SOA emphasizes designing any one thing to have many uses, to be “reusable.” By contrast, more and more specific microservice patterns emphasize “do one thing and one thing well.” Don’t cater to every possible use case, but be focused and simple. Ironically, “reuse” is more common in these designs, because a simple thing that is reasonable tends to be useful in many contexts. Reuse is more successful when driven by simplicity than by intention.

It can be difficult to enforce when working with existing assets such as ERPs, where a “strangler” pattern must be employed to replace the existing system’s data stores with the new microservices architecture piece by piece.

It sidesteps the issue of data synchronization, so if data entities get out of sync there is no easy fallback position.

GOALS:

Cohesion: this architecture is very easy to work with and understand due to its standardized nature.

Scalability: it is very scalable (each small component can implement its own scaling model).

Speed of change: it is good due to the more cohesive architecture, but requires governance to ensure the architecture is not breached.

KEY TRADE-OFFS:

There is efficient IPC due to asynchronicity.

It has a very flexible design, so speed of change is high.

The data consistency is good: there is a single source of truth.

Scalability may pose challenges, as scaling a process also requires scaling a data store with it.

It’s difficult, at scale, to divide a data model into completely independent pieces. At some stage, consistency between the views becomes important.

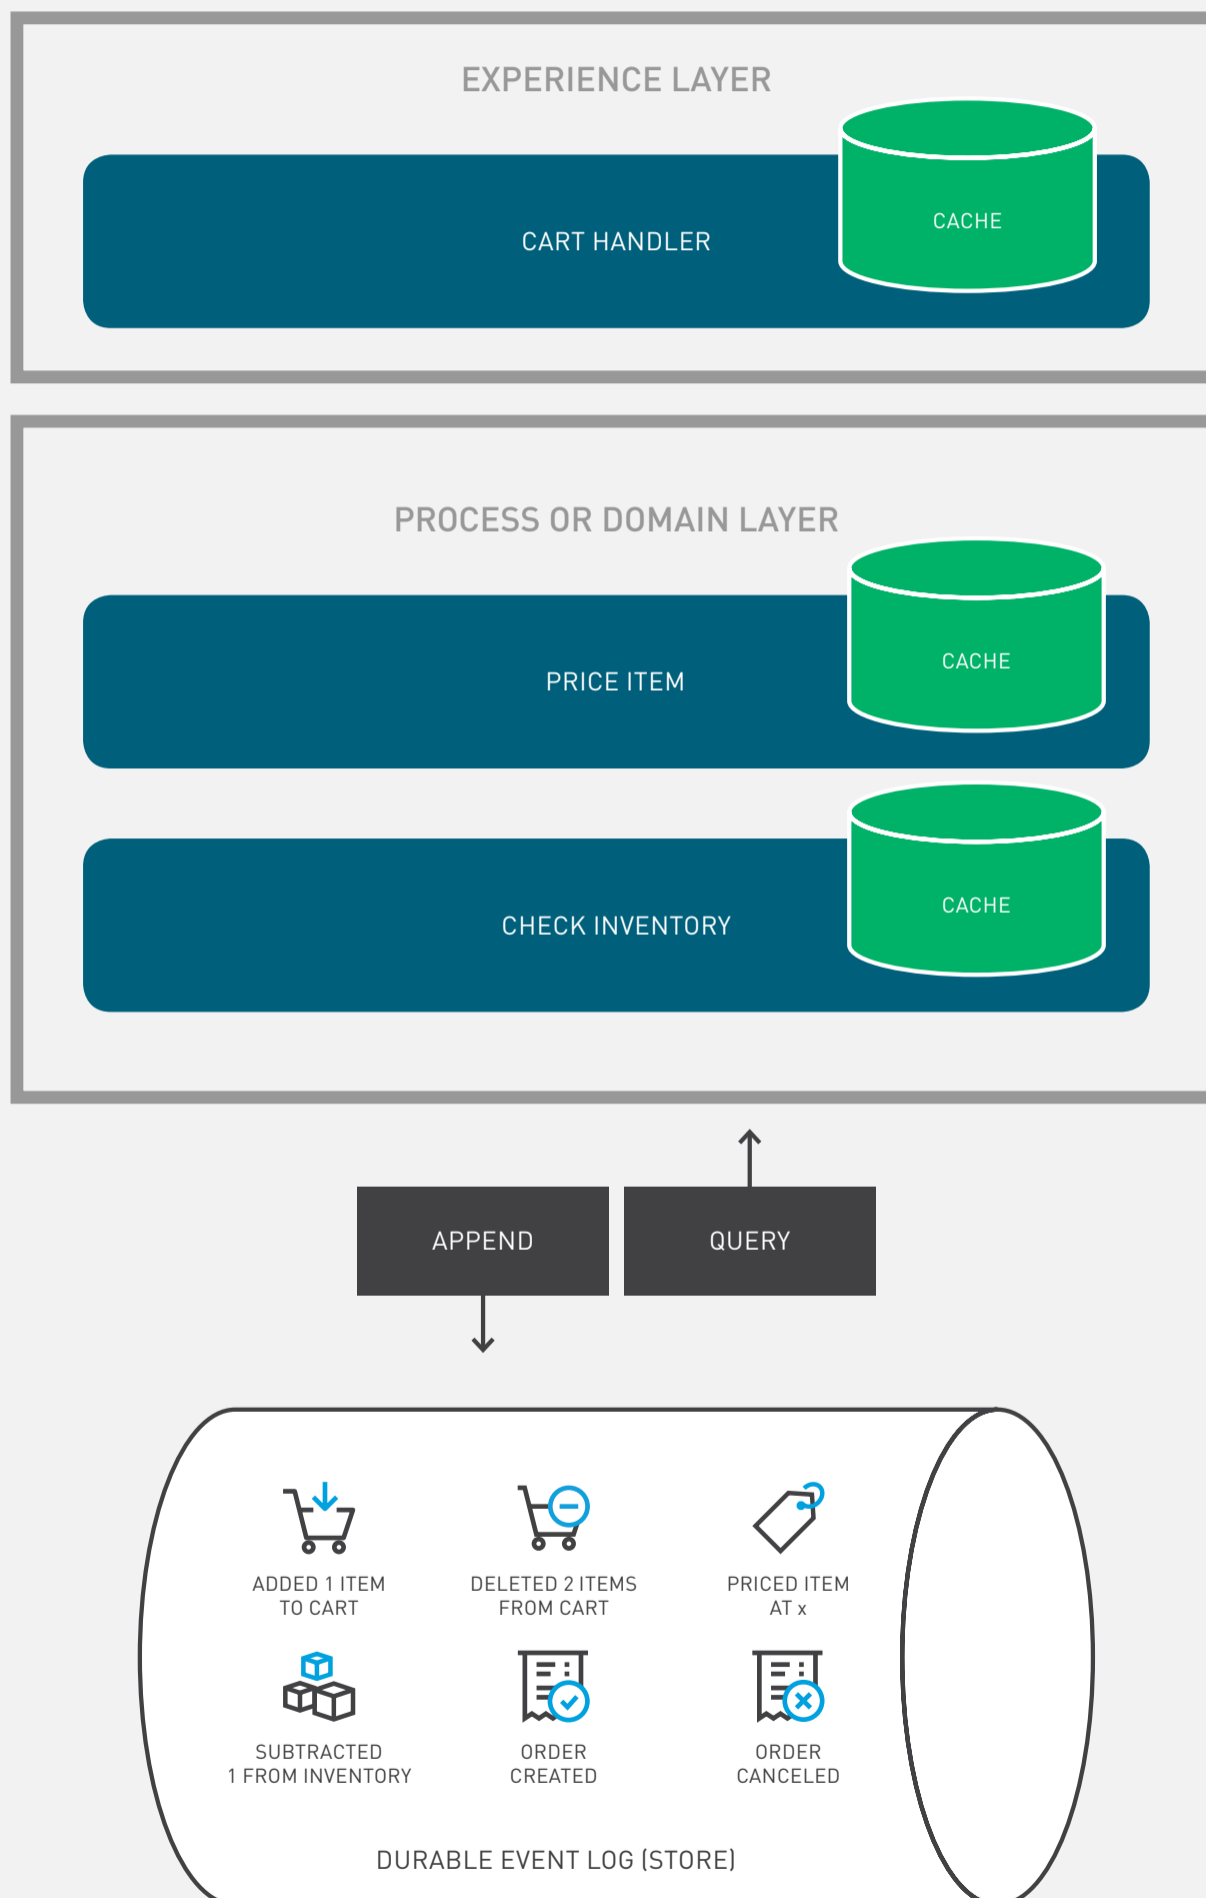
HOW DOES THIS COEXIST WITH EXISTING SYSTEMS, SOA, OR APIs?

It doesn't. If you build a microservice with isolated state, it's important that it is the "source of truth and function" for its stated purpose. If you have an existing system that also deals with the same data or function, you will need to synchronize it out of band, and it's typically a bad pattern to implement two-way sync here.

This approach typically pairs well with the "strangler" pattern, where you seek to reduce the use of a given enterprise application or another system that does not give you the time to value you need. Over time, you replace its functions with these isolated microservices and you deactivate those particular functions in the original system.

To put it another way, this is not an integration pattern. The goal here is to create a new, fast-moving partial implementation using microservices. This will allow you to gain speed and scale advantages that existing technology has not been able to provide you.

REPLICATING STATE IN LAYERED APIs (EVENT SOURCING)



Replicating state is essentially the antidote to the problems that emerge from isolating state; specifically, that consistency is required. A simple example is if we imagine a Catalog, Pricing, and Currency micro-service. If each of those contains an isolated state of each thing, they become interdependent. And failure or change in one can cause the function of the other(s) to fail.

This problem is addressed by replicating state; in other words, providing a single place to store all state mutations that each isolated micro-service can rebuild its internal state from. Often, this is coexistent with event sourcing, where event-driven microservices communicate exclusively via an immutable event log — providing a separate single source of truth that is consistent but difficult to query. The microservices that provide the ability to query state, therefore, have done the work of “materializing a view” of the event log.

This design is, by nature, eventually consistent. While this may seem like a problem in traditional transactional design, it is ameliorated by insight into the nature of the design. For example, one might think of a debit to a bank account as inherently transactional, but most modern banks have realized that it’s easier to create an eventually consistent debit (debit if the account exists, and then ensuring against the possibility the funds are not available) than to expend effort on ensuring every single transaction is consistent. This represents a newer way of thinking about IT systems, but enables greater freedom and speed of change, and, therefore, faster time to value.

Replicating state, of course, is challenging. It requires a deep understanding of the state being managed, and the behaviors of each micro-service in order to be predictable. However, it also directly addresses the problems that emerge from other patterns and, as such, can be seen as a very specific trade-off. This means eventual — rather than direct — consistency, cohesion over top-down design, and speed of change over predictability.

PROBLEM:

It is difficult to achieve data integrity when there are multiple sources of truth.

SOLUTION:

Keep a single source of truth of all changes to data, and replicate the data as needed.

APPLICATION:

Send all changes as events to a permanent Event Log. When needing to query data, build a materialized view by computing all the changes from the event log.

This is often streamlined by creating snapshots along the way of the views so that full recomputation is not required every time.

IMPACTS:

It creates a very cohesive architecture.

It is extremely scalable, due to the inherent Command-Query Request Separation in the design.

It can be difficult to visualize and understand logical dependencies (physical dependencies have been explicitly reduced).

GOALS:

Cohesion: this architecture is very easy to work with and understand due to its standardized nature.

Scalability: it is very scalable.

Speed of change: it is excellent.

KEY TRADE-OFFS:

There is efficient IPC due to asynchronicity.

There is a very flexible design, so speed of change is high.

The data consistency is good with caveats, but there is a single source of truth (typically the event log).

Scalability is effective; this design prioritizes the ability to scale each piece independently.

The autonomy is very high at the expense of a complex model.

HOW DOES THIS COEXIST WITH EXISTING SYSTEMS, SOA, OR APIs?

Ironically, this approach can coexist very well with existing systems, with only one key change: the event log must become the source of

truth for anything it contains. This means that existing systems and APIs can continue to be used, as long as they update the event log and are updated from it.

This approach can also be used in a strangler pattern way, by migrating event-by-event to this approach for the services you need to display speed of change and high scale. This replaces the existing implementation with a stable pattern to keep the existing system in sync if you wish.

SECTION III: FOUNDATIONAL BEST PRACTICES FOR ESTABLISHING THE MICROSERVICES PATTERNS

After reviewing the microservices patterns, and selecting which ones make the most sense to adopt for the rest of your organization, you might be tempted to stop there. However, in order to make the architecture work, there are a number of foundational best practices your organization needs to adopt in order to get microservices functional. These best practices are outlined below.

ANTI-FRAGILE SOFTWARE⁸

It is important to create a sense of predictability to operate a complex infrastructure at high scale. It is critical to ensure that your software is designed to be robust in the face of failure on all fronts. You can't rely on your infrastructure to be resilient.

This leads microservices developers towards practices that encourage uninterrupted operation in the face of broad systems failure. While the concept of “Chaos Engineering” (also pioneered by Netflix) has existed for some time, the key lessons weren't that broadly adopted until the advent of microservices. “Anti-fragility” is a combination of mindset and specific practice, and some of the key practices include:

⁸ There are multiple data masters, but the architecture doesn't typically address the existence of a source of truth or use a specific consistency model. There are often multiple ways to make a change to data that can have cascading impact, which is difficult to understand or respond to.

12-factor app principles are useful, but shouldn't be over-applied. Pivotal Cloud Foundry was originally based on these principles, but experience with the myriad of possibilities in real-world scale has recently led to relaxing some of the more extreme aspects. Nonetheless, most of these basic principles apply in any situation (e.g. logging as event streams).

Use intelligent defaults. If configuration files or environment variables are unavailable for some reason, software should initialize into an operational state by using reasonable default values.

Manage working directories and temp files. How many times have you had to fix an error caused when an app tries to write to a directory that doesn't exist or the app doesn't have privileges for? If your app needs access to a given file or directory, make sure that the code that accesses it checks for access and creates the file if it needs to.

Avoid race conditions and orchestrated startup. In most cases, software is simple enough that it should be launchable in isolation, and then seek to converge with other components. Many apps today require a specific start order (a specific example is first the database, then the app server), and this isn't necessary. Instead of erroring if a connection is unavailable, start up anyway and try to reconnect on a backing-off timeout.

Make bootstrap bulletproof. You can probably sense a theme from these items. The goal is to make sure that software components can start without error (regardless of the runtime environment) and, over time, seek to reach an ideal state. The benefits of this approach are

broad and obvious, but the basic guideline should be that operators never need to understand the internals of the software.

Use circuit breakers. These are a pattern. In microservices deployments, they are often implemented as part of a high-performance inter-process communications framework like Hystrix or Finagle that detects failures and provides logic to prevent them from reoccurring. In other words, it detects an error condition and prevents components from attempting to retry the “doomed” action until the error condition resolves.

Use timeouts. Similar to using intelligent defaults, any external communication should include a timeout. Further, if many services are involved in an end-to-end scenario, it is important to think of timeout “budgets.” For example, the timeout value for the third call in a chain should not be the same or longer as the timeout for the first call, or else you can have components still processing “down the chain” for calls that have already terminated at the edge.

HEALTHZ

A common pattern with microservices is “healthZ,” or, in other words, apps exposing a known endpoint (/health in Spring Boot, /healthz as a common pattern) that returns a simple health check. This check should indicate two things:

Does this app think that it is healthy? (Yes: 200, No: 5xx).

What does it believe its current state to entail? It will return a basic set of information (e.g. JSON) that describes the current state of its internal dependencies. These should be readable to an operator — not a dev — and be self-explanatory - for example “database: connecting”.

Note that health and correct operating state are different: it is common in container frameworks that a component can be running and healthy but not yet “ready” for serving traffic (e.g. the database isn’t connected yet). HealthZ endpoints should indicate when something is really wrong, not that there is a temporary operational blip.

“INFINITE” (LINEAR) SCALE

Microservices patterns focus a lot on scalability, often with the term “infinite” scale being bandied about. Of course, this doesn’t mean truly infinite scale. Instead, it is a shorthand for the idea of having a clear understanding of how it would be possible with a given piece of software to achieve a linear scaling model. Often, this focuses on the hard limits to scale: storage of data and state management. With microservices, this encourages developers to think about the patterns of data and storage they are using, and seek a way to perform the task that can support high scale. For example, this can be achieved by using eventually consistent clustered storage instead of relying on transactional boundaries provided by RDBMSes. Generally, this is a good practice. Although there is no need to overthink this point, there is value in identifying what your storage layers are being used for and making sure you use the “best tool for the job.”

MINIMIZE DEPENDENCIES

When you're deploying many changes frequently, it is important to ensure that your component has minimal dependencies on external systems (for example, by using queues for communication rather than synchronous request/reply patterns. While microservice patterns make this almost mandatory, it's useful in general to make each component as self-sufficient as possible. Again, don't overthink this point, but a good rule of thumb is to try to make every unit of deployment as self-contained as you can.

MONITOR EVERYTHING

A sufficiently complex infrastructure requires visibility. Microservices practitioners learn the hard way that it's essential to have an effective monitoring solution baked into their software, and it's hard to argue that this shouldn't be the default state for any well-made app. In an ideal state, it is possible to instrument software without requiring changes in the application code, but following basic practices (e.g. logging as event streams to `STDOUT/STDERR`) make it easier to put monitoring infrastructure in place. This is only half the battle, of course; more monitoring data creates a need to comprehend that data better. The topic of how to effectively design monitors and to manage monitoring infrastructure is out of the scope of this paper, but there are many specific texts on this topic, such as James Turnbull's "The Art of Monitoring."

REDUCE BATCH SIZE

In Lean Manufacturing, we have learned that there is significant value in reducing the batch size. Microservices as a set of patterns explicitly leverage this approach; the smaller the unit you work with, the simpler each unit is to operate. Of course, the larger the number of units you have, the more complex the management gets.

However, even without adopting microservice patterns, you can reduce the size of your unit of deployment by making changes more frequently and in smaller amounts. Even when deploying a change to a monolith, the disciplines that come with implementing smaller, more frequent changes are valuable.

CONTAINERIZATION

The use of containers to structure, isolate, and manage units of deployment is generally useful, even when applied to monolithic software. With recent data showing that the use of containers brings no performance or security overhead, it is easy to recommend the use of Docker as a basic unit of deployment. Better yet, microservices patterns generally require containerization techniques, so adopting containers is a good step towards being ready for a microservices journey.

FOCUSED (DOMAIN-DRIVEN) DESIGN

Domain-driven design has risen to a new popularity due to its high applicability to microservices. While, in our experience, overuse of DDD can be a distraction in a high scale microservices environment, it's inarguable that the practices proposed by DDD are of high value and make

a lot of sense to apply in general. Further, we might even argue that DDD is more appropriate to a monolithic design than it is to microservices patterns and, as such, adopting domain thinking within your software will both be helpful and pave an easier transition towards microservices when or if) you need to make the change.

CONCLUSION: ESTABLISH A MICROSERVICES PATTERN THAT'S RIGHT FOR YOU

Now that you have read through the different microservices patterns, it's important for you to assess which ones work best for your organization. It could be one, it could be a set, or you might decide to stay with a monolithic architecture.

The important thing is to remember that microservices are not a cure-all that will solve all of your problems. It is an architecture designed to overcome obstacles that, when deployed correctly, will produce certain desired results.

It has yet to be seen how the patterns discussed here will evolve, and how software may rise up to simplify the challenges of this highly complex approach. It is clear, however, that many enterprises will move to adopt microservice patterns, and we hope that by providing some definition of these patterns we may help many of these organizations to avoid the most obvious and painful mistakes along the way.

For more information on establishing a microservices architecture in your organization, take a look at our whitepaper, [Microservices Best Practices](#), or take a look at the demos in our [microservices webinar](#).



ABOUT MULESOFT

MuleSoft's mission is to help organizations change and innovate faster by making it easy to connect the world's applications, data and devices. With its API-led approach to connectivity, MuleSoft's market-leading Anypoint Platform™ is enabling over 1,000 organizations in more than 60 countries to build application networks. For more information, visit mulesoft.com.

MuleSoft is a registered trademark of MuleSoft, Inc. All other marks are those of respective owners