


Aggregator

Aggregator

This applies for Camel version 2.2 or older. If you use a newer version then the Aggregator has been rewritten from Camel 2.3 on and you should use this [Aggregator2](#) link instead.


The [Aggregator](#) from the [EIP patterns](#) allows you to combine a number of messages together into a single message.




A correlation [Expression](#) is used to determine the messages which should be aggregated together. If you want to aggregate all messages into a single message, just use a constant expression. An [AggregationStrategy](#) is used to combine all the message exchanges for a single correlation key into a single message exchange. The default strategy just chooses the latest message; so its ideal for throttling messages.


For example, imagine a stock market data system; you are receiving 30,000 messages per second; you may want to throttle down the updates as, say, a GUI cannot cope with such massive update rates. So you may want to aggregate these messages together so that within a window (defined by a maximum number of messages or a timeout), messages for the same stock are aggregated together; by just choosing the latest message and discarding the older prices. (You could apply a delta processing algorithm if you prefer to capture some of the history).

Using the aggregator correctly

 Torsten Mielke wrote a nice [blog entry](#) with his thoughts and experience on using the aggregator. Its a well worth read. AggregationStrategy changed in Camel 2.0

 In Camel 2.0 the [AggregationStrategy](#) callback have been changed to also be invoked on the very first Exchange.

On the first invocation of the [aggregate](#) method the [oldExchange](#) parameter is [null](#). The reason is that we have not aggregated anything yet. So its only the [newExchange](#) that has a value. Usually you just return the [newExchange](#) in this situation. But you still have the power to decide what to do, for example you can do some alternation on the exchange or remove some headers. And a more common use case is for instance to count some values from the body payload. That could be to sum up a total amount etc. [BatchTimeout](#) and [CompletionPredicate](#)

 You cannot use both [batchTimeout](#) and [completionPredicate](#) to trigger a completion based on either on reaching its goal first. The batch timeout will always trigger first, at that given interval.

Using the Fluent Builders

The following example shows how to aggregate messages so that only the latest message for a specific value of the [cheese](#) header are sent. Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```
// in this route we aggregate all from direct:state based on the header id cheese
from("direct:start").aggregate(header("cheese")).to("mock:result");

from("seda:header").setHeader("visited", constant(true)).aggregate(header("cheese")).to("mock:result");

// in this sample we aggregate using our own strategy with a completion predicate
// stating that the aggregated header is equal to 5.
from("direct:predicate").aggregate(header("cheese"), new MyAggregationStrategy()).
    completionPredicate(header("aggregated").isEqualTo(5)).to("mock:result");

// this sample is similar to the one above but it also illustrates the use of batchSize
// to send exchanges to mock:endpoint in batches of 10.
from("direct:outBatchPredicate").aggregate(header("cheese"), new MyAggregationStrategy()).
    completionPredicate(header("aggregated").isEqualTo(5)).outBatchSize(10).to("mock:result");
```

If you were using [JMS](#) then you may wish to use the [JMSTDestination](#) header as the correlation key; or some custom header for the stock symbol (using the above stock market example).

```
from("activemq:someReallyFastTopic")
    .aggregator(header("JMSTDestination"))
    .to("activemq:someSlowTopicForGuis");
```

You can of course use many different [Expression](#) languages such as [XPath](#), [XQuery](#), [SQL](#) or various [Scripting Languages](#). Here is an example using **XPath**:

```
//aggregate based on the message content using an XPath expression
//example assumes an XML document starting with <stockQuote symbol='...'>
//aggregate messages based on their symbol attribute within the <stockQuote> element
from("seda:start").aggregate().xpath("/stockQuote/@symbol", String.class).batchSize(5).to("mock:result");

//this example will aggregate all messages starting with <stockQuote symbol='APACHE'> into
//one exchange and all the other messages (different symbol or different root element) into another exchange.
from("seda:start").aggregate().xpath("name(/stockQuote[@symbol='APACHE'])", String.class).batchSize(5).to("mock:result");
```

For further examples of this pattern in use you could look at the [junit test case](#)

Using the [Spring XML Extensions](#)

The `correlationExpression` element is in Camel 2.0. For earlier versions of Camel you will need to specify your expression without the enclosing `correlationExpression` element.

```
<aggregator>
  <simple>header.cheese</simple>
  <to uri="mock:result"/>
</aggregator>
```

The following example shows how to create a simple aggregator using the XML notation; using an [Expression](#) for the correlation value used to aggregate messages together. Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <aggregate>
      <correlationExpression>
        <simple>header.cheese</simple>
      </correlationExpression>
      <to uri="mock:result"/>
    </aggregate>
  </route>

  <route>
    <from uri="seda:header"/>
    <process ref="setHeaderProcessor"/>
    <to uri="direct:temp"/>
  </route>

  <route>
    <from uri="direct:temp"/>
    <aggregate>
      <correlationExpression>
        <simple>header.cheese</simple>
      </correlationExpression>
      <to uri="mock:result"/>
    </aggregate>
  </route>

  <route>
    <from uri="direct:predicate"/>
    <aggregate strategyRef="myAggregatorStrategy">
      <correlationExpression>
        <simple>header.cheese</simple>
      </correlationExpression>
      <to uri="mock:result"/>
      <completionPredicate>
        <method bean="myAggregatorStrategy" method="isCompleted"/>
      </completionPredicate>
    </aggregate>
  </route>

  <route>
    <from uri="direct:outBatchPredicate"/>
    <aggregate strategyRef="myAggregatorStrategy" outBatchSize="10">
      <correlationExpression>
        <simple>header.cheese</simple>
      </correlationExpression>
      <to uri="mock:result"/>
      <completionPredicate>
        <method bean="myAggregatorStrategy" method="isCompleted"/>
      </completionPredicate>
    </aggregate>
  </route>

  <!-- This route turns off in batching by setting batchSize to 1 to run unit test for out batching.
       Normal use cases may not want to disable in batching
  -->
  <route>
    <from uri="direct:outBatchNoInBatching"/>
    <aggregate strategyRef="myAggregatorStrategy" batchSize="1" outBatchSize="10">
      <correlationExpression>
        <simple>header.cheese</simple>
      </correlationExpression>
      <to uri="mock:result"/>
      <completionPredicate>
        <method bean="myAggregatorStrategy" method="isCompleted"/>
      </completionPredicate>
    </aggregate>
  </route>
</camelContext>

```

You can specify your own **AggregationStrategy** if you prefer as shown in the following example
 Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
  <route>
    <from uri="direct:start"/>
    <aggregate strategyRef="aggregatorStrategy">
      <correlationExpression>
        <simple>header.cheese</simple>
      </correlationExpression>
      <to uri="mock:result"/>
    </aggregate>
  </route>
</camelContext>

<bean id="aggregatorStrategy" class="org.apache.camel.spring.processor.MyAggregator"/>

```

Notice how the `strategyRef` attribute is used on the `<aggregator>` element to refer to the custom strategy in Spring.

Exchange Properties

The following properties is set on each Exchange that are aggregated:

header	type	description
<code>org.apache.camel.Exchange.AggregatedCount</code>	int	Camel 1.x: The total number of Exchanges aggregated in this combined Exchange.
<code>CamelAggregatedSize</code>	int	Camel 2.0: The total number of Exchanges aggregated into this combined Exchange.
<code>CamelAggregatedIndex</code>	int	Camel 2.0: The current index of this Exchange in the batch.

Batch options

The aggregator supports the following batch options:

Option	Default	Description
<code>batchSize</code>	100	The in batch size. This is the number of incoming exchanges that is processed by the aggregator and when this threshold is reached the batch is completed and send. Camel 1.6.2/2.0: You can disable the batch size so the Aggregator is only triggered by timeout by setting the <code>batchSize</code> to 0 (or negative). In Camel 1.6.1 or older you can set the <code>batchSize</code> to a very large number to archive the same.
<code>outBatchSize</code>	0	Camel 1.5: The out batch size. This is the number of exchanges currently aggregated in the <code>AggregationCollection</code> . When this threshold is reached the batch is completed and send. By default this option is disabled. The difference to the <code>batchSize</code> options is that this is for outgoing, so setting this size to e.g. 50 ensures that this batch will at maximum contain 50 exchanges when its sent.
<code>batchTimeout</code>	1000L	Timeout in millis. How long should the aggregator wait before its completed and sends whatever it has currently aggregated.
<code>groupExchanges</code>	false	Camel 2.0: If enabled then Camel will group all aggregated Exchanges into a single combined <code>org.apache.camel.impl.GroupedExchange</code> holder class that holds all the aggregated Exchanges. And as a result only one Exchange is being sent out from the aggregator. Can be used to combine many incoming Exchanges into a single output Exchange without coding a custom <code>AggregationStrategy</code> yourself.
<code>batchConsumer</code>	false	Camel 2.0: This option is if the exchanges is coming from a Batch Consumer . Then when enabled the <code>Aggregator</code> will use the batch size determined by the <code>Batch Consumer</code> in the message header <code>CamelBatchSize</code> . See more details at Batch Consumer . This can be used to aggregate all files consumed from a <code>File</code> endpoint in that given poll.
<code>completionPredicate</code>	null	Allows you to use a Predicate to signal when an aggregation is complete. See warning in top of this page.

AggregationCollection and AggregationStrategy

This aggregator uses a `AggregationCollection` to store the exchanges that is currently aggregated. The `AggregationCollection` uses a correlation `Expression` and an `AggregationStrategy`.

- The correlation `Expression` is used to correlate the incoming exchanges. The default implementation will group messages based on the correlation expression. Other implementations could for instance just add all exchanges as a batch.
- The strategy is used for aggregate the old (lookup by its correlation id) and the new exchanges together into a single exchange. Possible implementations include performing some kind of combining or delta processing, such as adding line items together into an invoice or just using the newest exchange and removing old exchanges such as for state tracking or market data prices; where old values are of little use.

Camel provides these implementations:

- `DefaultAggregationCollection`
- `PredicateAggregationCollection`
- `UseLatestAggregationStrategy`

Examples

Default example

By default Camel uses `DefaultAggregationCollection` and `UseLatestAggregationStrategy`, so this simple example will just keep the latest received exchange for the given correlation [Expression](#):`Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'`

```
// our route is aggregating from the direct queue and sending the response to the mock
from("direct:start")
    // aggregated by header id
    // as we have not configured more on the aggregator it will default to aggregate the
    // latest exchange only
    .aggregate().header("id")
    // wait for 0.5 seconds to aggregate
    .batchTimeout(500L)
    .to("mock:result");
```

Using PredicateAggregationCollection

The `PredicateAggregationCollection` is an extension to `DefaultAggregationCollection` that uses a [Predicate](#) as well to determine the completion. For instance the [Predicate](#) can test for a special header value, a number of maximum aggregated so far etc. To use this the routing is a bit more complex as we need to create our `AggregationCollection` object as follows:`Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'`

```
// create the aggregation collection we will use.
// - we will correlate the received message based on the id header
// - as we will just keep the latest message we use the latest strategy
// - and finally we stop aggregate if we receive 2 or more messages
AggregationCollection ag = new PredicateAggregationCollection(header("id"),
    new UseLatestAggregationStrategy(),
    property(Exchange.AGGREGATED_SIZE).isEqualTo(3));

// our route is aggregating from the direct queue and sending the response to the mock
from("direct:start")
    // we use the collection based aggregator we already have configured
    .aggregate(ag)
    // wait for 0.5 seconds to aggregate
    .batchTimeout(500L)
    .to("mock:result");
```

In this sample we use the predicate that we want at most 3 exchanges aggregated by the same correlation id, this is defined as:

```
header(Exchange.AGGREGATED_COUNT).isEqualTo(3)
```

Using this the aggregator will complete if we receive 3 exchanges with the same correlation id or when the specified timeout of 500 msecs has elapsed (whichever criteria is met first).

Using Custom Aggregation Strategy

In this example we will aggregate incoming bids and want to aggregate the highest bid. So we provide our own strategy where we implement the code logic:`Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'`

```
private static class MyAggregationStrategy implements AggregationStrategy {

    public Exchange aggregate(Exchange oldExchange, Exchange newExchange) {
        if (oldExchange == null) {
            // the first time we only have the new exchange so it wins the first round
            return newExchange;
        }
        int oldPrice = oldExchange.getIn().getBody(Integer.class);
        int newPrice = newExchange.getIn().getBody(Integer.class);
        // return the "winner" that has the highest price
        return newPrice > oldPrice ? newExchange : oldExchange;
    }
}
```

Then we setup the routing as follows:`Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'`

```

// our route is aggregating from the direct queue and sending the response to the mock
from("direct:start")
    // aggregated by header id and use our own strategy how to aggregate
    .aggregate(new MyAggregationStrategy()).header("id")
    // wait for 0.5 seconds to aggregate
    .batchTimeout(500L)
    .to("mock:result");

```

And since this is based on an unit test we show the test code that send the bids and what is expected as the `winners`:Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

MockEndpoint result = getMockEndpoint("mock:result");

// we expect to find the two winners with the highest bid
result.expectedMessageCount(2);
result.expectedBodiesReceived("200", "150");

// then we sent all the message at once
template.sendBodyAndHeader("direct:start", "100", "id", "1");
template.sendBodyAndHeader("direct:start", "150", "id", "2");
template.sendBodyAndHeader("direct:start", "130", "id", "2");
template.sendBodyAndHeader("direct:start", "200", "id", "1");
template.sendBodyAndHeader("direct:start", "190", "id", "1");

assertMockEndpointsSatisfied();

```

Using Custom Aggregation Collection

In this example we will aggregate incoming bids and want to aggregate the bids in reverse order (this is just an example). So we provide our own collection where we implement the code logic:Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

class MyReverseAggregationCollection extends AbstractCollection<Exchange> implements AggregationCollection {

    private List<Exchange> collection = new ArrayList<Exchange>();
    private Expression correlation;
    private AggregationStrategy strategy;

    public Expression getCorrelationExpression() {
        return correlation;
    }

    public void setCorrelationExpression(Expression correlationExpression) {
        this.correlation = correlationExpression;
    }

    public AggregationStrategy getAggregationStrategy() {
        return strategy;
    }

    public void setAggregationStrategy(AggregationStrategy aggregationStrategy) {
        this.strategy = aggregationStrategy;
    }

    public boolean add(Exchange exchange) {
        return collection.add(exchange);
    }

    public Iterator<Exchange> iterator() {
        // demonstrate the we can do something with this collection, so we reverse it
        Collections.reverse(collection);

        return collection.iterator();
    }

    public int size() {
        return collection.size();
    }

    public void clear() {
        collection.clear();
    }

    public void onAggregation(Object correlationKey, Exchange newExchange) {
        add(newExchange);
    }
}

```

Then we setup the routing as follows:Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

// our route is aggregating from the direct queue and sending the response to the mock
from("direct:start")
    // use our own collection for aggregation
    .aggregate(new MyReverseAggregationCollection())
    // wait for 0.5 seconds to aggregate
    .batchTimeout(500L)
    .to("mock:result");

```

And since this is based on an unit test we show the test code that send the bids and what is expected as the expected reverse order:Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

MockEndpoint result = getMockEndpoint("mock:result");

// we expect 5 messages since our custom aggregation collection just gets it all
// but returns them in reverse order
result.expectedMessageCount(5);
result.expectedBodiesReceived("190", "200", "130", "150", "100");

// then we sent all the message at once
template.sendBodyAndHeader("direct:start", "100", "id", "1");
template.sendBodyAndHeader("direct:start", "150", "id", "2");
template.sendBodyAndHeader("direct:start", "130", "id", "2");
template.sendBodyAndHeader("direct:start", "200", "id", "1");
template.sendBodyAndHeader("direct:start", "190", "id", "1");

assertMockEndpointsSatisfied();

```

Custom aggregation collection in Spring DSL

You can also specify a custom aggregation collection in the Spring DSL. Here is an example for Camel 2.0:Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

<camelContext id="camel" xmlns="http://camel.apache.org/schema/spring">
    <route>
        <from uri="direct:start"/>
        <aggregate batchTimeout="500" collectionRef="aggregatorCollection">
            <to uri="mock:result"/>
        </aggregate>
    </route>
</camelContext>

<bean id="aggregatorCollection" class="org.apache.camel.processor.aggregator.MyReverseAggregationCollection"/>

```

In Camel 1.5.1 you will need to specify the aggregator as:

```

<aggregator batchTimeout="500" collectionRef="aggregatorCollection">
    <expression/>
    <to uri="mock:result"/>
</aggregator>

```

Using Grouped Exchanges

Available as of Camel 2.0

You can enable grouped exchanges to combine all aggregated exchanges into a single `org.apache.camel.impl.GroupedExchange` holder class that contains all the individual aggregated exchanges. This allows you to process a single Exchange containing all the aggregated exchange. Lets start with how to configure this in the router:Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

// our route is aggregating from the direct queue and sending the response to the mock
from("direct:start")
    // aggregate all using same expression
    .aggregate().constant(true)
    // wait for 0.5 seconds to aggregate
    .batchTimeout(500L)
    // group the exchanges so we get one single exchange containing all the others
    .groupExchanges()
    .to("mock:result");

```

And the next part is part of an unit code that demonstrates this feature as we send in 5 exchanges each with a different value in the body. And we will only get 1 exchange out of the aggregator, but we can access all the individual aggregated exchanges from the List which we can extract as a property from the Exchange using the key `Exchange.GROUPED_EXCHANGE`.Error rendering macro 'code': Invalid value specified for parameter 'java.lang.NullPointerException'

```

MockEndpoint result = getMockEndpoint("mock:result");

// we expect 1 messages since we group all we get in using the same correlation key
result.expectedMessageCount(1);

// then we sent all the message at once
template.sendBody("direct:start", "100");
template.sendBody("direct:start", "150");
template.sendBody("direct:start", "130");
template.sendBody("direct:start", "200");
template.sendBody("direct:start", "190");

assertMockEndpointsSatisfied();

Exchange out = result.getExchanges().get(0);
List<Exchange> grouped = out.getProperty(Exchange.GROUPED_EXCHANGE, List.class);

assertEquals(5, grouped.size());

assertEquals("100", grouped.get(0).getIn().getBody(String.class));
assertEquals("150", grouped.get(1).getIn().getBody(String.class));
assertEquals("130", grouped.get(2).getIn().getBody(String.class));
assertEquals("200", grouped.get(3).getIn().getBody(String.class));
assertEquals("190", grouped.get(4).getIn().getBody(String.class));

```

Using Batch Consumer

Available as of Camel 2.0

The [Aggregator](#) can work together with the [Batch Consumer](#) to aggregate the total number of messages that the [Batch Consumer](#) have reported. This allows you for instance to aggregate all files polled using the [File](#) consumer.

For example:

```

from("file://inbox")
    .aggregate(xpath("//order/@customerId"), new AggregateCustomerOrderStrategy()).batchConsumer().batchTimeout(60000).to("bean:processOrder");

```

When using `batchConsumer` Camel will automatic adjust the `batchSize` according to reported by the [Batch Consumer](#) in this case the file consumer. So if we poll in 7 files then the aggregator will aggregate all 7 files before it completes. As the timeout is still in play we set it to 60 seconds.

Using This Pattern

If you would like to use this EIP Pattern then please read the [Getting Started](#), you may also find the [Architecture](#) useful particularly the description of [Endpoint](#) and [URIs](#). Then you could try out some of the [Examples](#) first before trying this pattern out.

See also

- The [Loan Broker Example](#) which uses an aggregator
- [Blog post by Torsten Mielke](#) about using the aggregator correctly.