



To Microservices... and back again

Alexandra Noonan
alexandra@segment.com

 anoonan



What we'll cover

1. What does Segment do?
2. Why Microservices worked
3. Moving back to a Monolith and why
4. The trade-offs we made, and how we're moving forward




— What does Segment do?



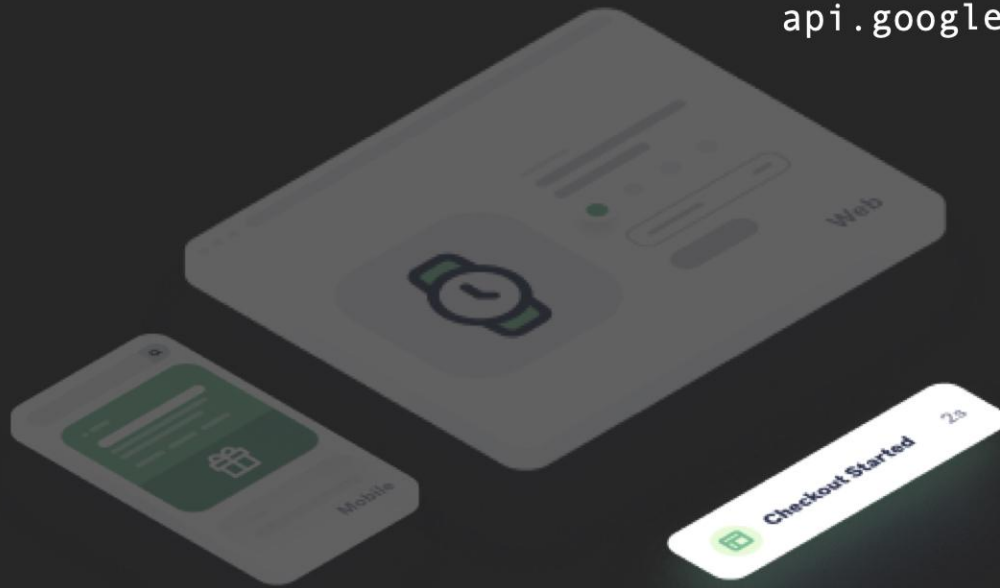
Sample event

```
{  
  "type": "identify",  
  "traits": {  
    "name": "Alex Noonan",  
    "email": "alexandra@segment.com",  
    "company": "Segment",  
    "title": "Software Engineer"  
  },  
  "userId": "97980cfea0067",  
  "timestamp": "2020-02-19T02:38:04Z"  
}
```






api.google-analytics.com



api.google-analytics.com






api.google-analytics.com

api.mixpanel.com


api.salesforce.com

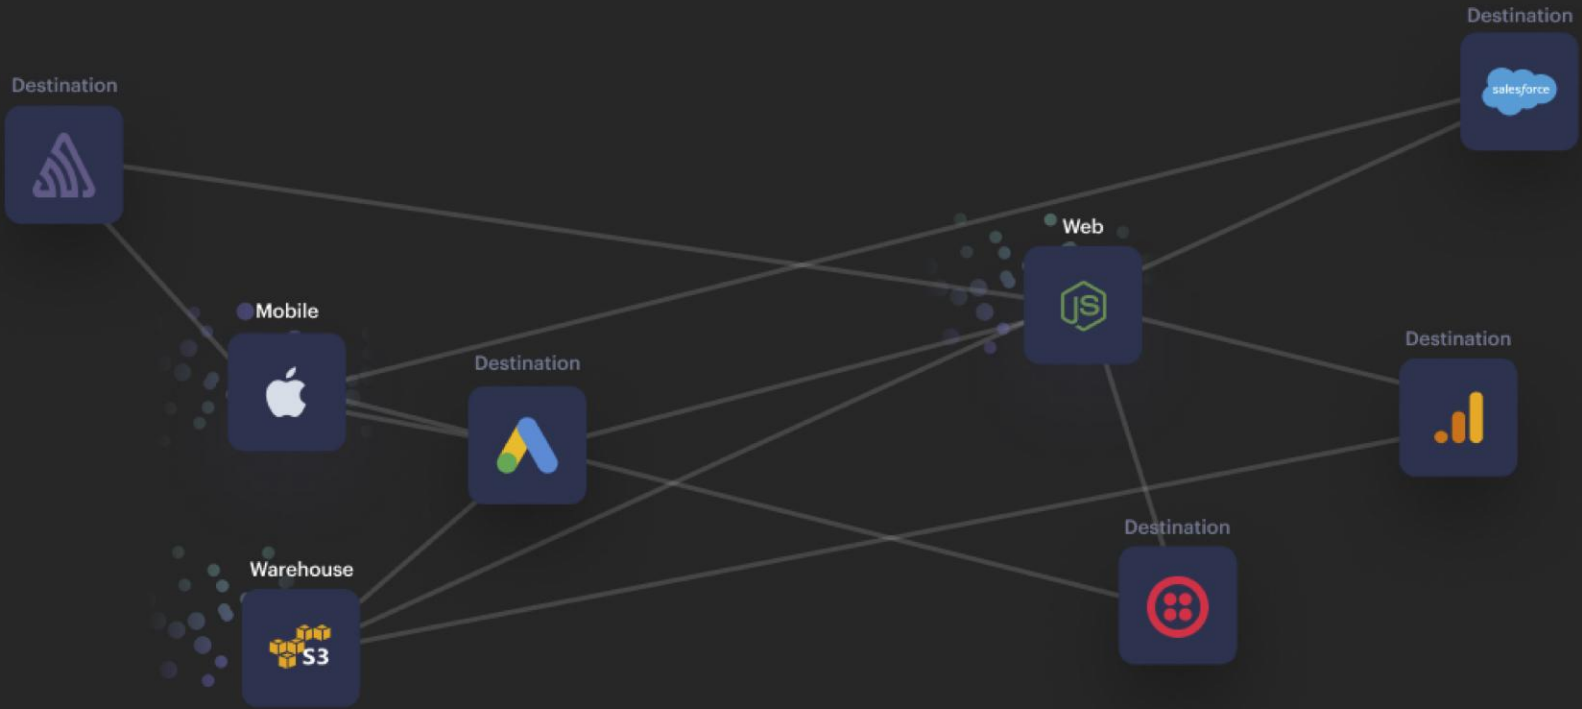


api.google-analytics.com

api.mixpanel.com

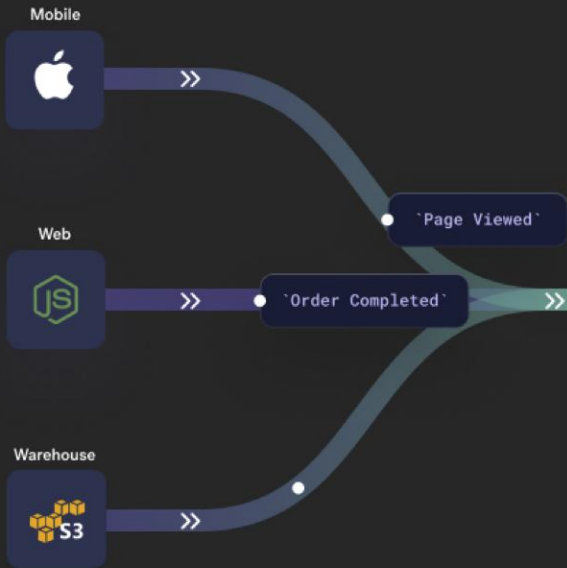
api.salesforce.com








SOURCES



DESTINATIONS



DESTINATIONS



Launched with monolith





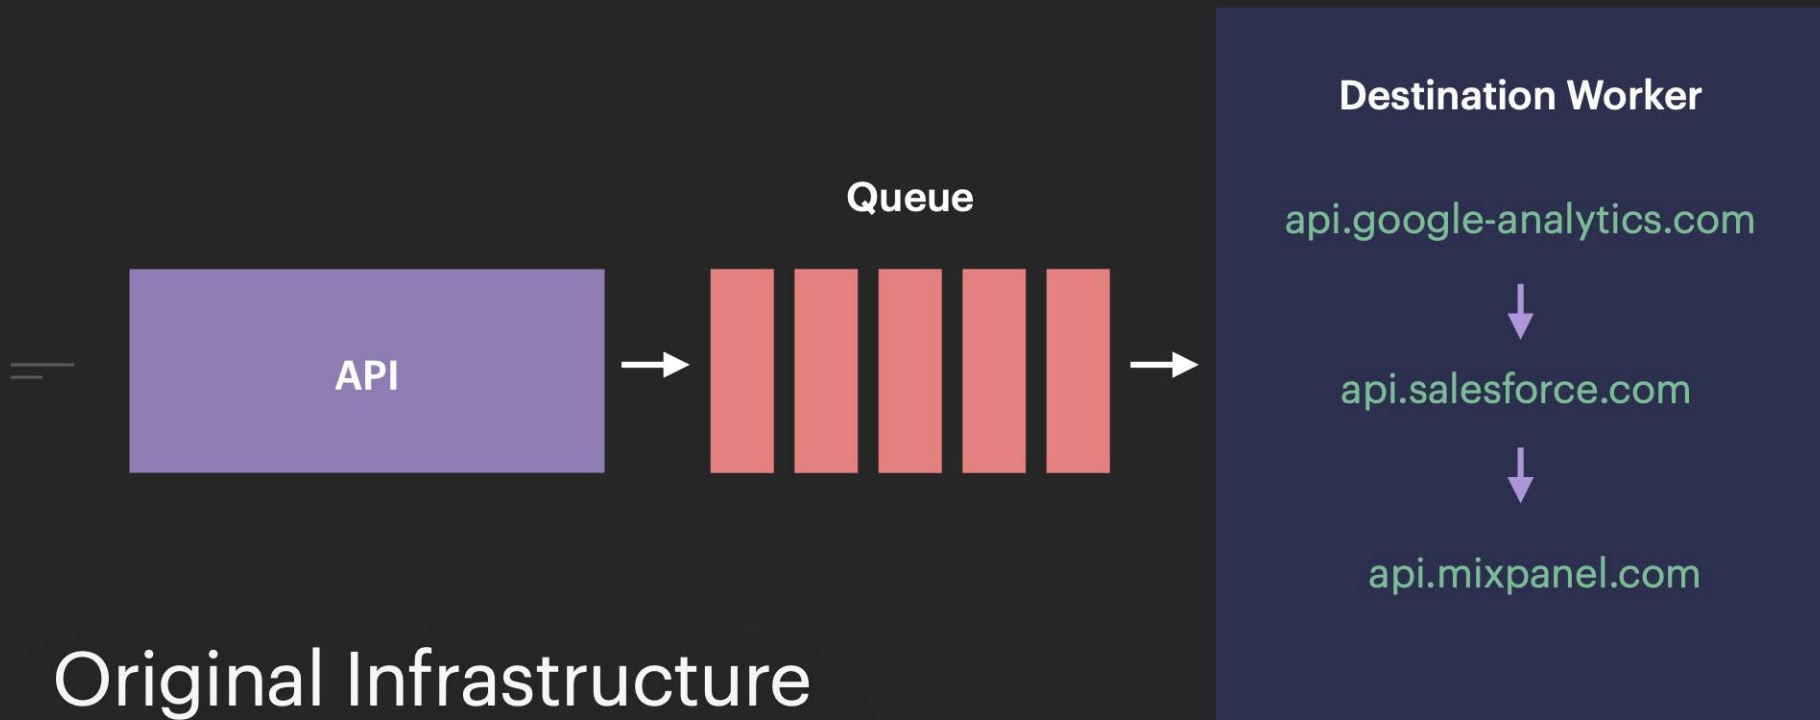
TRADE-OFFS



Low operational overhead

Launched with monolith





Original Infrastructure



Retryable Error

```
HTTP/1.1 429 Too Many Requests  
Content-Type: application/json
```

Response Body

```
{  
  "status": 429,  
  "message": "Rate limit exceeded. Try again  
in 1 second."  
}
```

Non-retryable Error

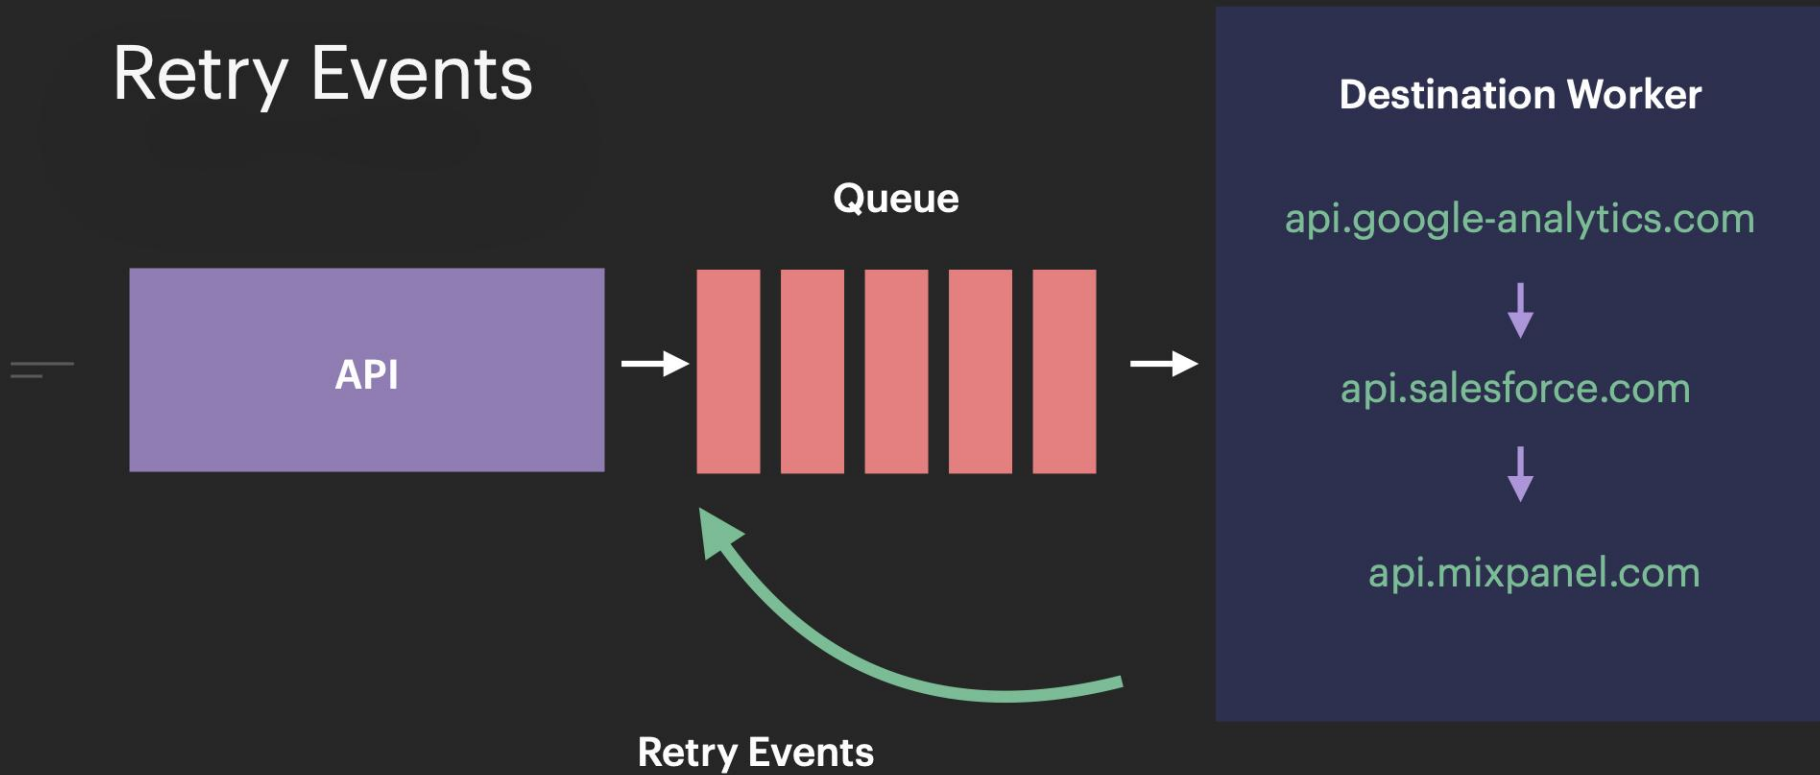
```
HTTP/1.1 400 Bad Request  
Content-Type: application/json
```

Response Body

```
{  
  "status": 400,  
  "message": "Bad Request. userId required  
in event payload. "  
}
```

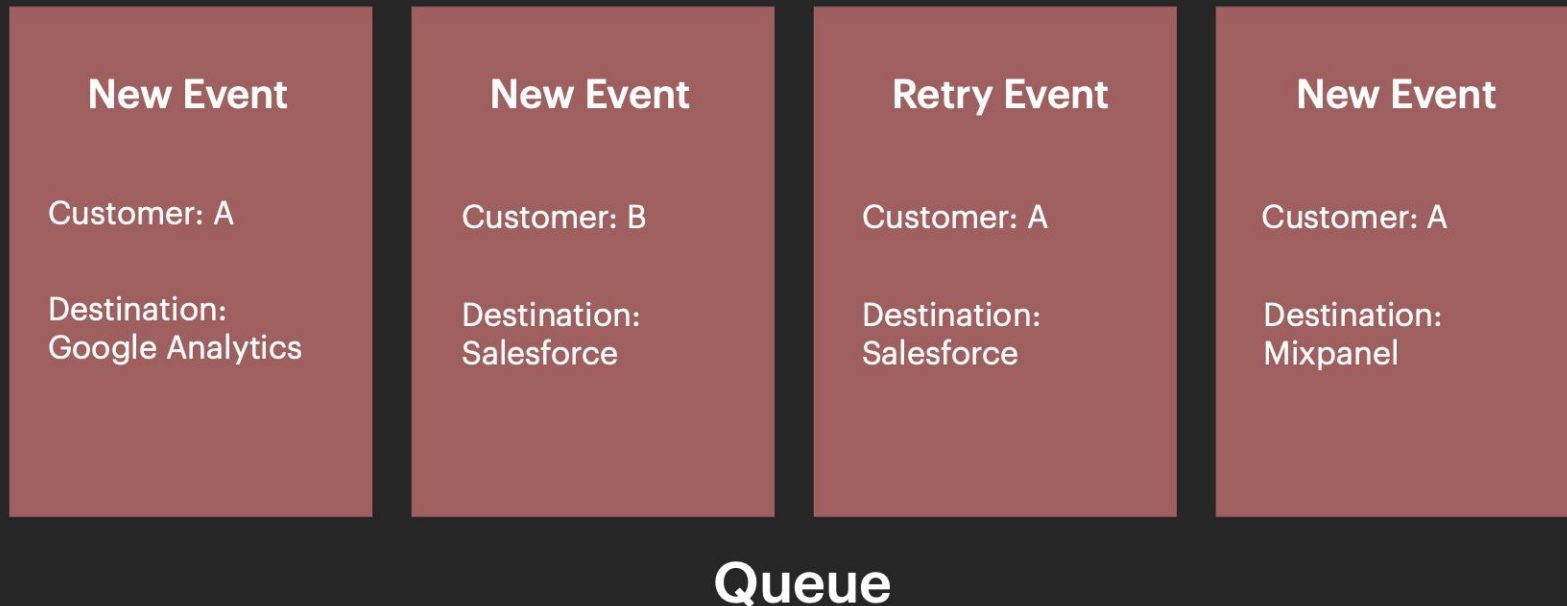


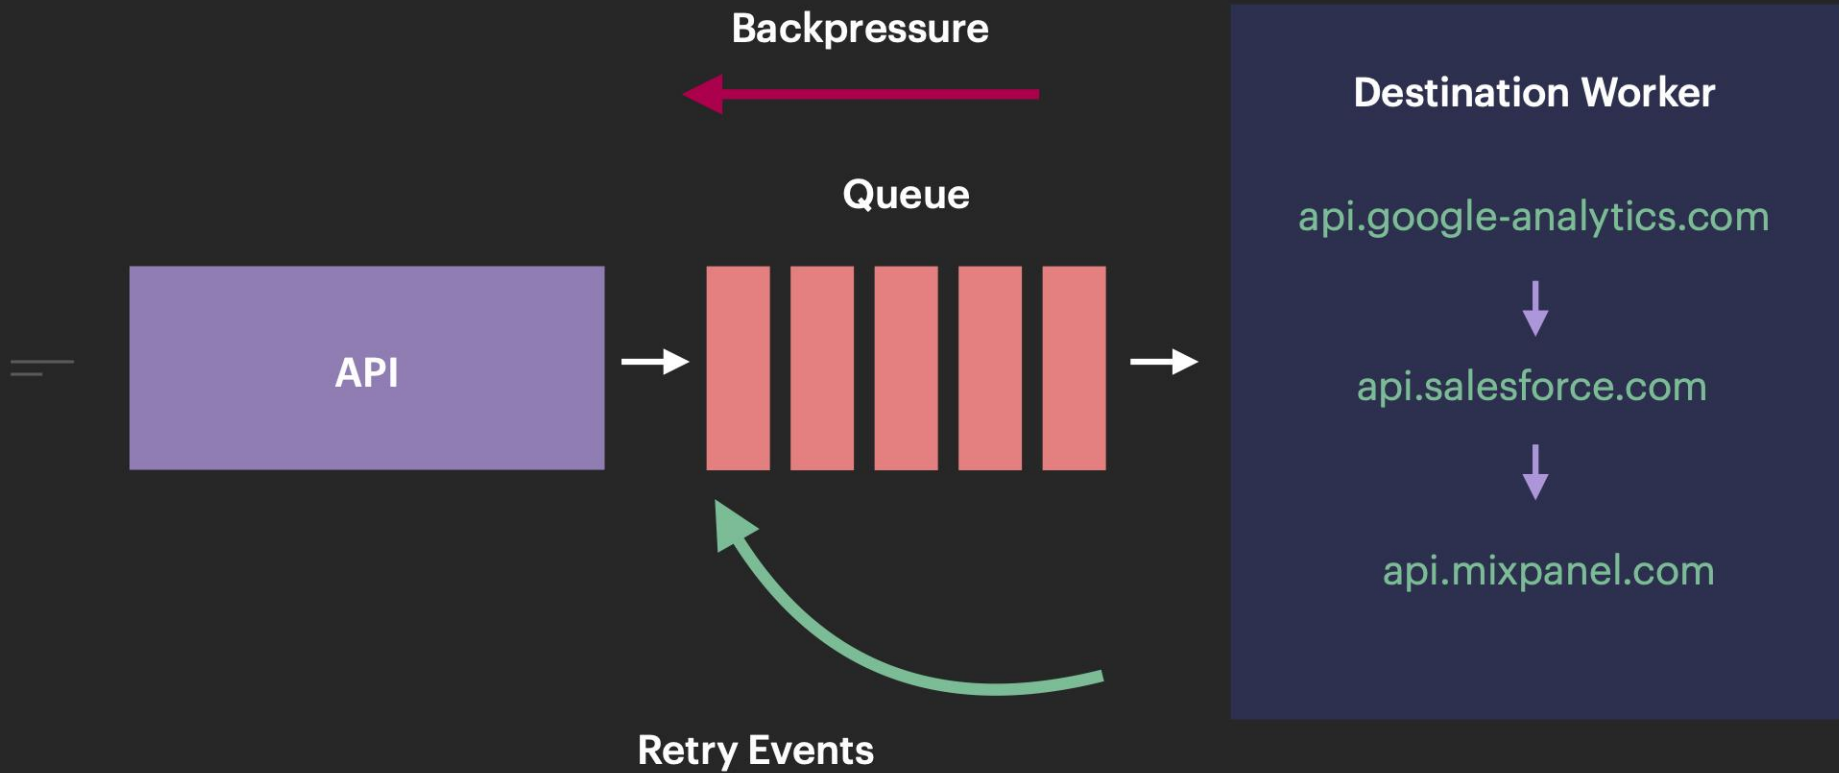
Retry Events





Head-of-line blocking







Queue

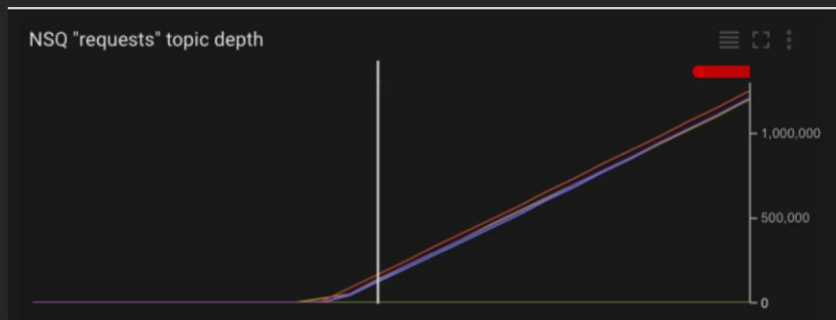


Worker Pool

Destination worker

Destination worker

Destination worker





Queue




Worker Pool

Destination worker


Destination worker

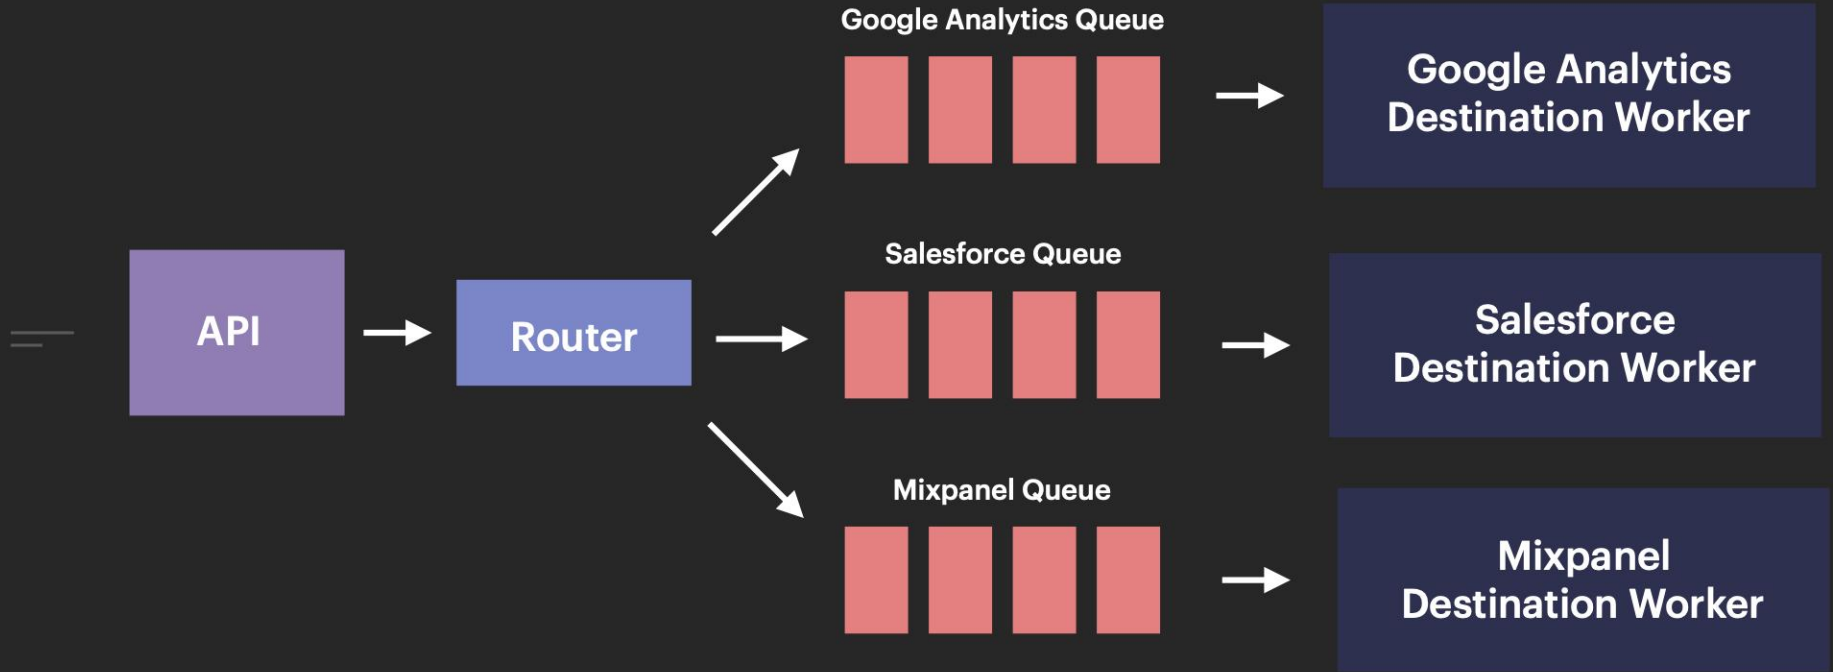
Destination worker



TRADE-OFFS

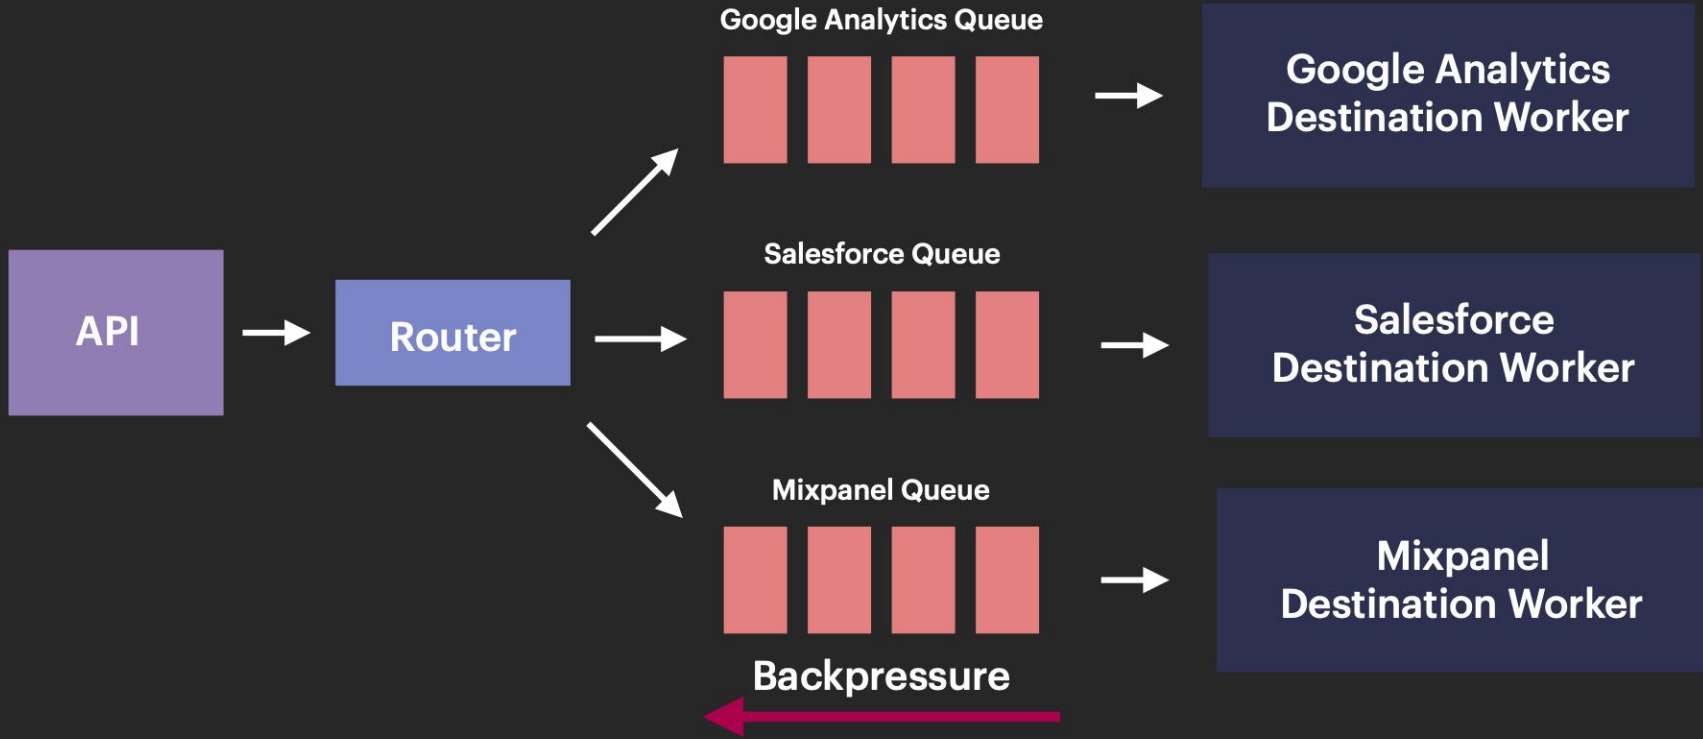
- ✓ Low operational overhead
- ✗ Environmental isolation

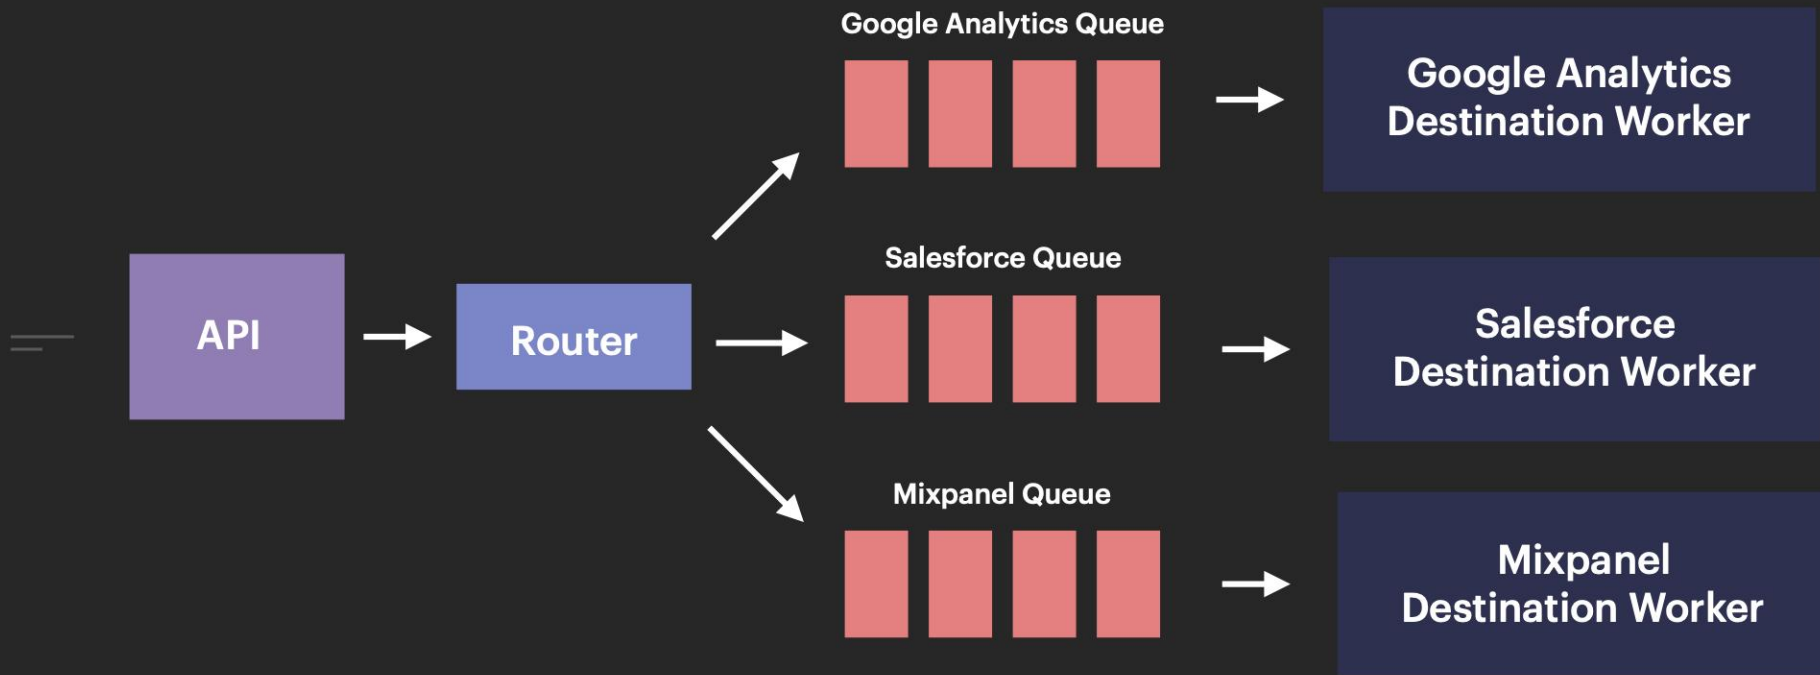


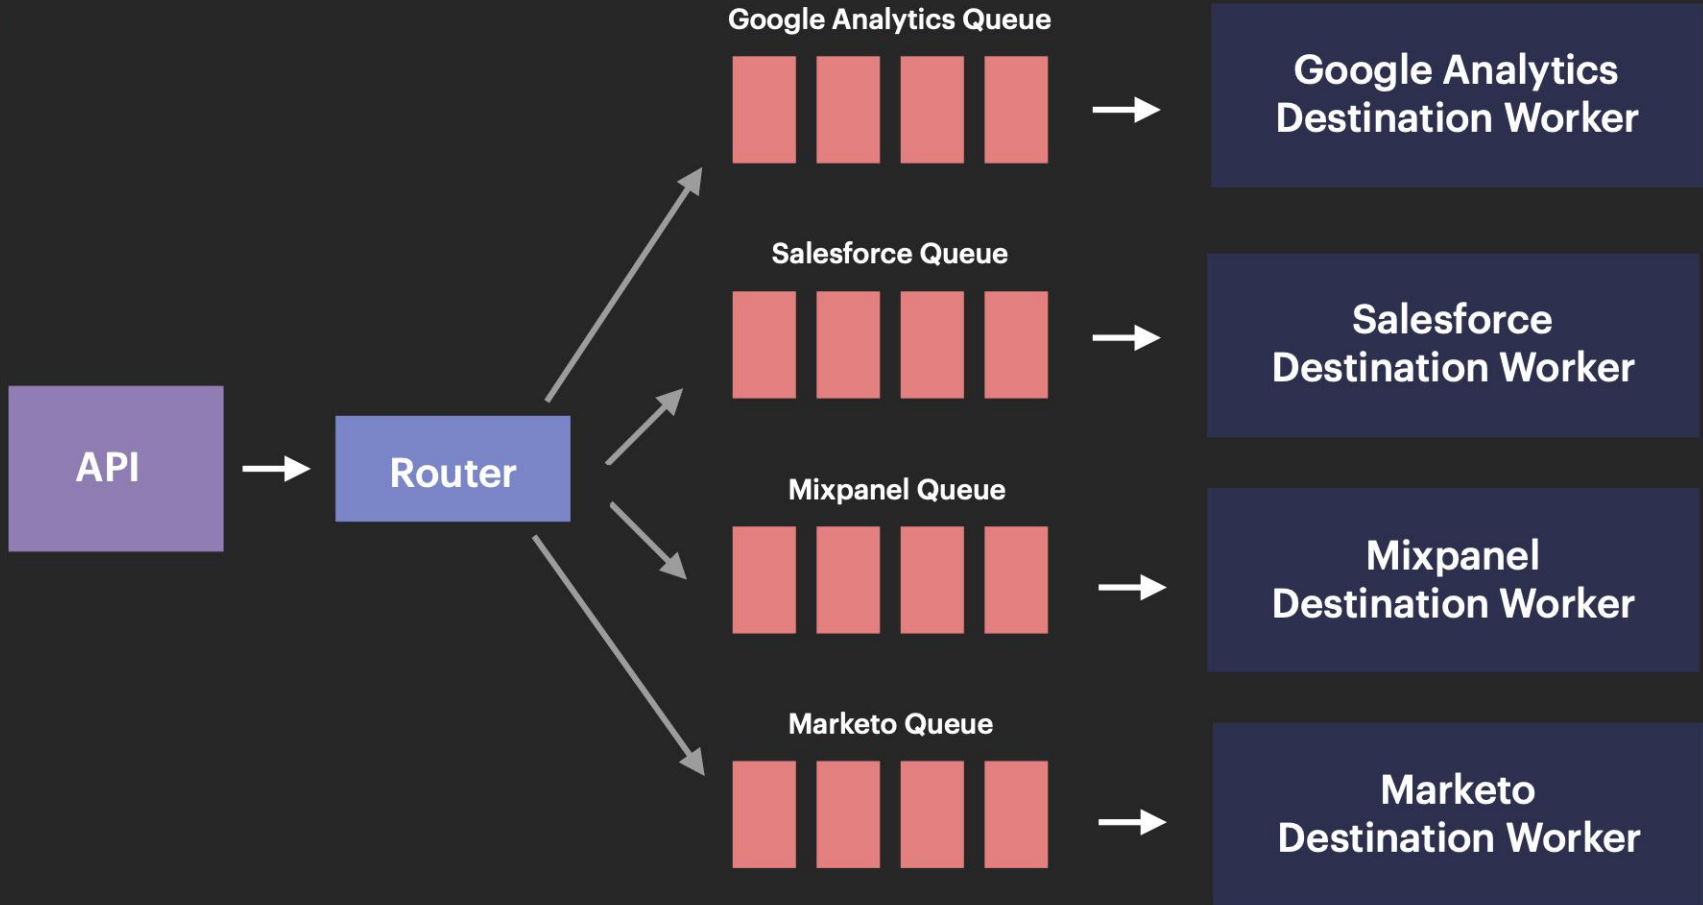




= Why microservices worked









```
212 exports.identify = function (segmentEvent) {↵
213   ··const destinationPayload = {}↵
214   ↵
215   ··destinationPayload.dateOfBirth: segmentEvent.birthdate,↵
216   ··destinationPayload.email:      segmentEvent.email,↵
217   ··destinationPayload.firstName:  segmentEvent.firstName,↵
218   ··destinationPayload.lastName:   segmentEvent.lastName,↵
219   ··destinationPayload.userId:     segmentEvent.userId↵
220   ↵
221   ··return destinationPayload↵
222 }
```



```
{
  "requestId": "e42b#14272d07d78",
  "success": true,
  "result": [
    {
      "status": "skipped",
      "reasons": [
        {
          "code": "1004",
          "message": "Not Found"
        }
      ]
    }
  ]
}
```



```
54 exports.parseResponseForErrors = function (res) {~
55   ~const success = res.body.success~
56   ~
57   ~// destination X will return success = true even if there is an error for a few of their errors so we have to make sure to catch that.~
58   ~if (success) {~
59     ~~~if (res.body.result.length > 0) {~
60       ~~~~~if (res.body.result[0].code === 1004 || res.body.result[0].code === 1007) {~
61         ~~~~~~const message = fmt('error code: %s message: %s', res.body.result[0].code, res.body.result[0].message)~
62         ~~~~~~const err = new Error(message)~
63         ~~~~~~err.status = Number(res.body.result[0].code)~
64         ~~~~~~return err~
65         ~~~~~}~
66         ~~~~~// destination X will return success = true but then status = skipped meaning it wasn't actually successful~
67         ~~~~~// sometimes the response will provide a reason why it was skipped and sometimes it won't :)~
68         ~~~~~if (res.body.result[0].status === 'skipped') {~
69           ~~~~~~const reason = res.body.result[0].reasons || res.body.result[0].errors~
70           ~~~~~~if (reason) {~
71             ~~~~~~const message = fmt('error code: %s message: %s', reason[0].code, reason[0].message)~
72             ~~~~~~const err = new Error(message)~
73             ~~~~~~err.status = Number(reason[0].code)~
74             ~~~~~~return err~
75             ~~~~~} else {~
76             ~~~~~~const err = new Error('Lead was not updated. Desintation X skipped the lead but did not return a reason why.')~
77             ~~~~~~err.status = 400~
78             ~~~~~~return err~
79             ~~~~~}~
80           ~~~~~}~
81         ~~~~~}~
82         ~~~~~return~
83       ~~~~~}~
84     ~~~if (res.body.errors) {~
85       ~~~~~const message = fmt('error code: %s message: %s', res.body.errors[0].code, res.body.errors[0].message)~
86       ~~~~~const err = new Error(message)~
87       ~~~~~err.status = Number(res.body.errors[0].code)~
88       ~~~~~return err~
89       ~~~~~}~
90     ~~~if (res.body.result) {~
91       ~~~~~const message = fmt('error code: %s message: %s', res.body.result[0].code, res.body.result[0].message)~
92       ~~~~~const err = new Error(message)~
93       ~~~~~err.status = Number(res.body.result[0].code)~
94       ~~~~~return err~
95       ~~~~~}~
96     ~~~}~
97   ~}
```



— Mono-repo vs Micro-repos



Monorepo

[/ destinations]

[google-analytics]

[marketo]

[drift]

[mixpanel]

[salesforce]

[webhooks]



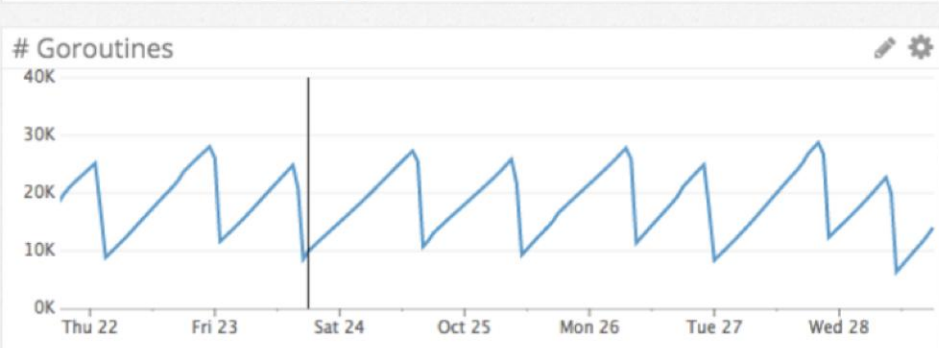

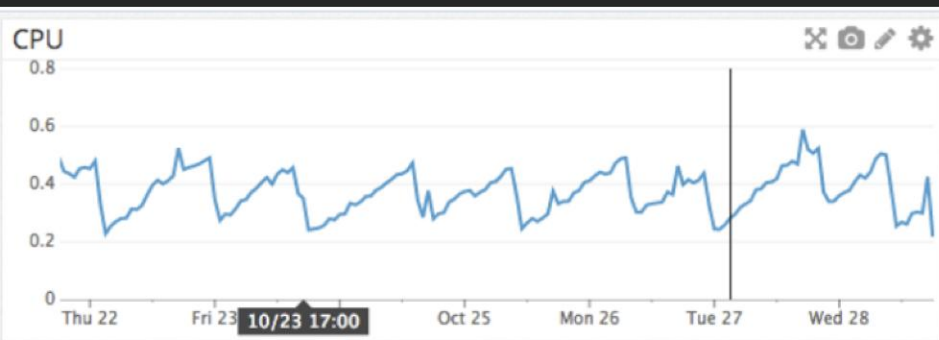
Failing tests

[/ destinations]

[google-analytics]	✓
[marketo]	✗
[drift]	✓
[mixpanel]	✓
[salesforce]	✓
[webhooks]	✓




Out of the box visibility







Queue depth

NSQ "requests" topic depth




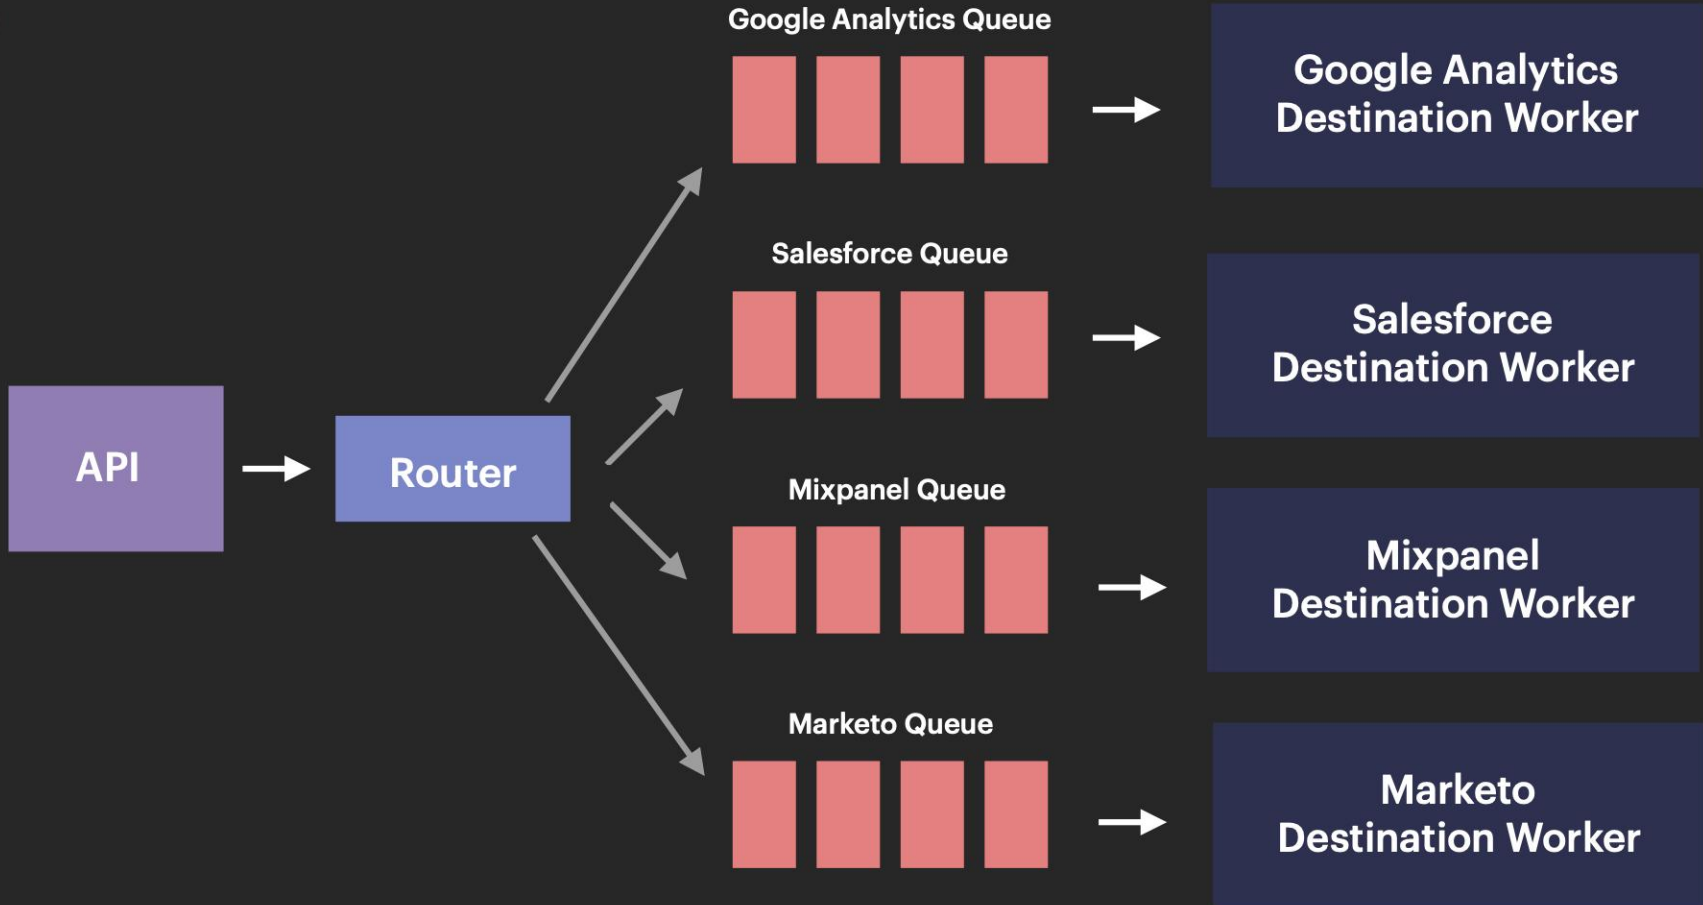


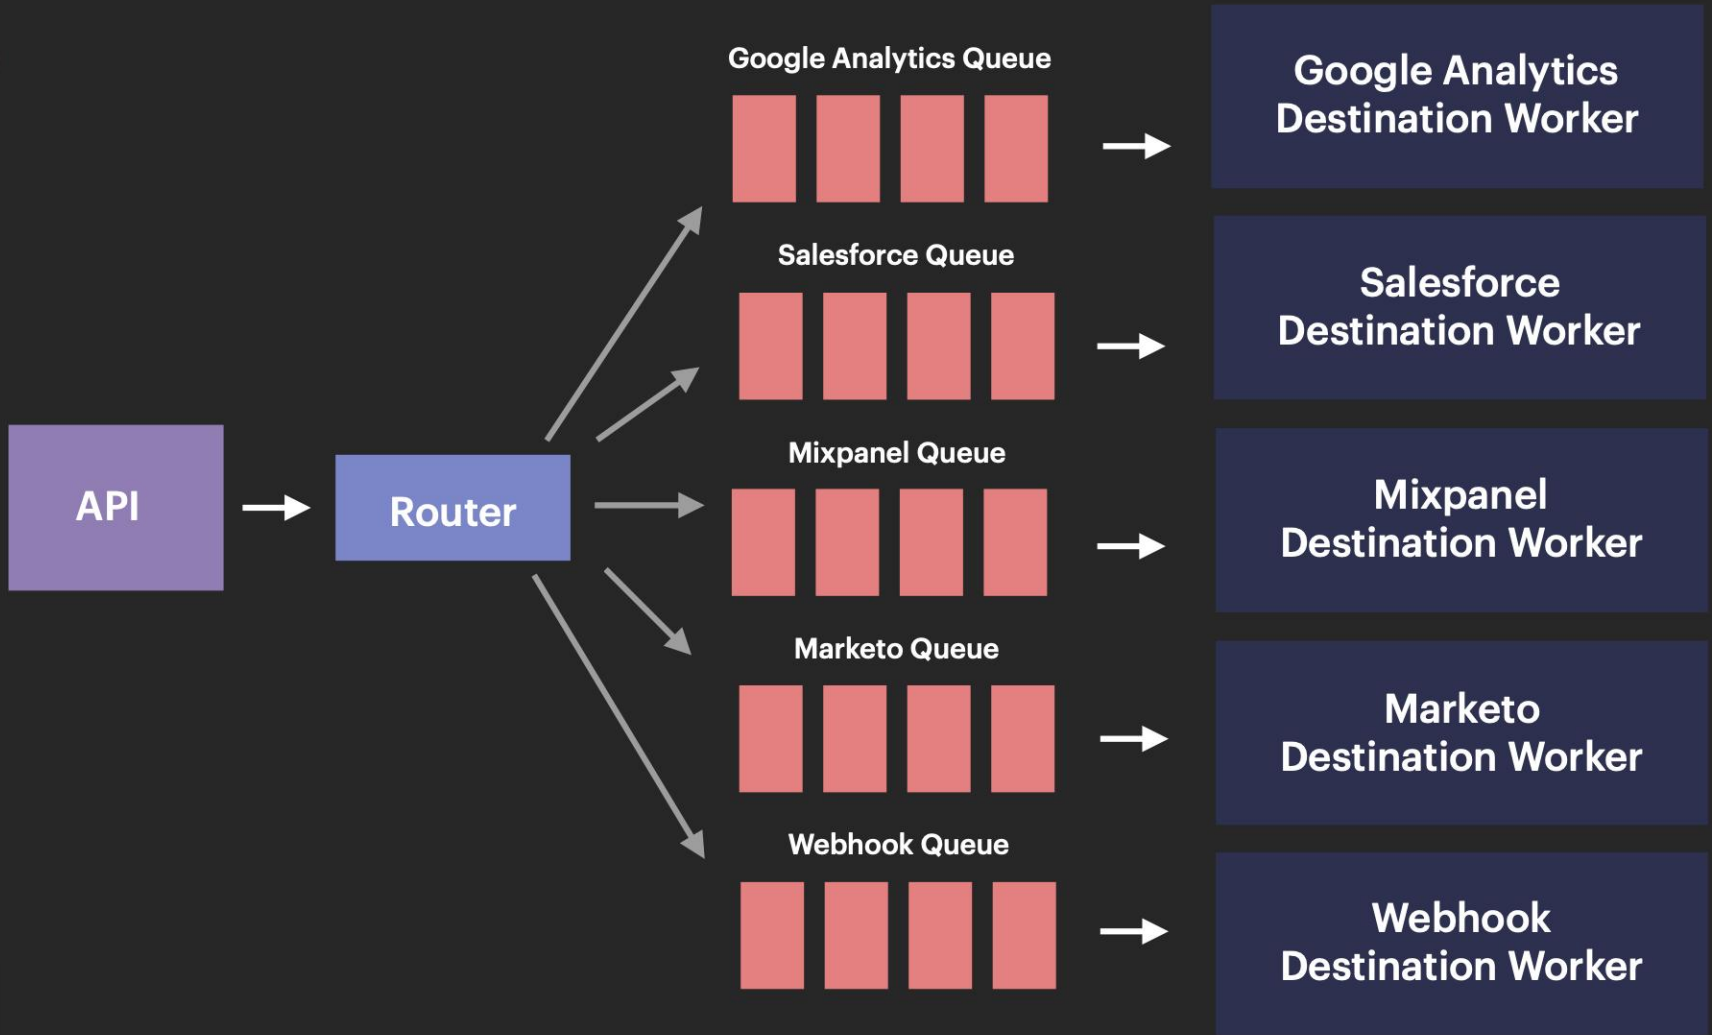


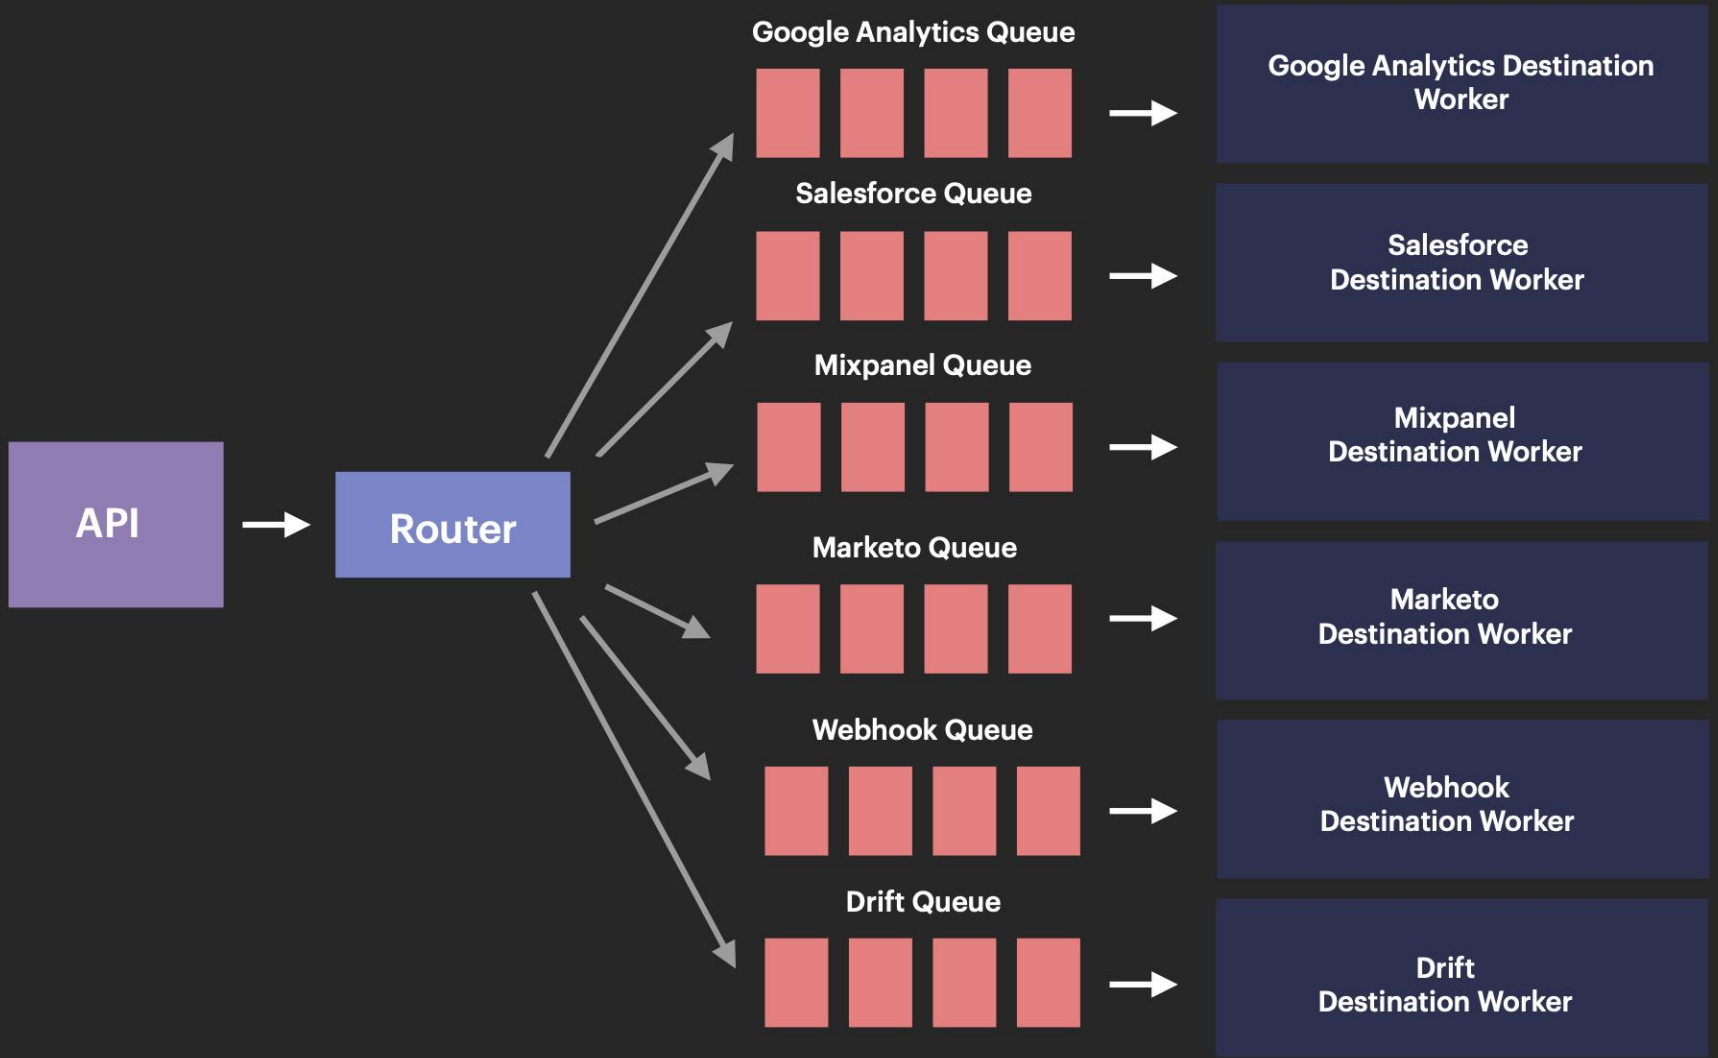
TRADE-OFFS

- ✓ Improved Modularity
- ✓ Environmental isolation
- ✓ Visibility
- Operational overhead











API



Router



Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker



```
{
  "type": "identify",
  "traits": {
    "name": "Alex Noonan"
  },
  "userId": "97980cfea0067"
}
```

```
{
  "type": "identify",
  "traits": {
    "firstName": "Alex",
    "lastName": "Noonan",
  },
  "userId": "97980cfea0067"
}
```

```
{
  "type": "identify",
  "traits": {
    "first_name": "Alex",
    "last_name": "Noonan",
  },
  "userId": "97980cfea0067"
}
```



Shared Libraries

```
174 Identify.prototype.name = function() {  
175   ..var name = this.proxy('traits.name');  
176   ..if (typeof name === 'string') {  
177     ...return trim(name);  
178   }  
179   ..  
180   ..var firstName = this.firstName();  
181   ..var lastName = this.lastName();  
182   ..if (firstName && lastName) {  
183     ...return trim(firstName + ' ' + lastName);  
184   }  
185 }
```



```
{
  "type": "identify",
  "traits": {
    "name": "Alex Noonan"
  },
  "userId": "97980cfea0067"
}
```

```
{
  "type": "identify",
  "traits": {
    "firstName": "Alex",
    "lastName": "Noonan",
  },
  "userId": "97980cfea0067"
}
```

```
{
  "type": "identify",
  "traits": {
    "first_name": "Alex",
    "last_name": "Noonan",
  },
  "userId": "97980cfea0067"
}
```



Standard Deploy Process

1. PR to bump shared library version.
2. Run test suite.
3. Release new version.
4. PR to update staging.
5. Run Terraform commands.
6. Test in staging.
7. PR to update production.
8. Run Terraform commands.
9. Test in production.
10. Monitor



API



Router



Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue



Destination Worker

Queue

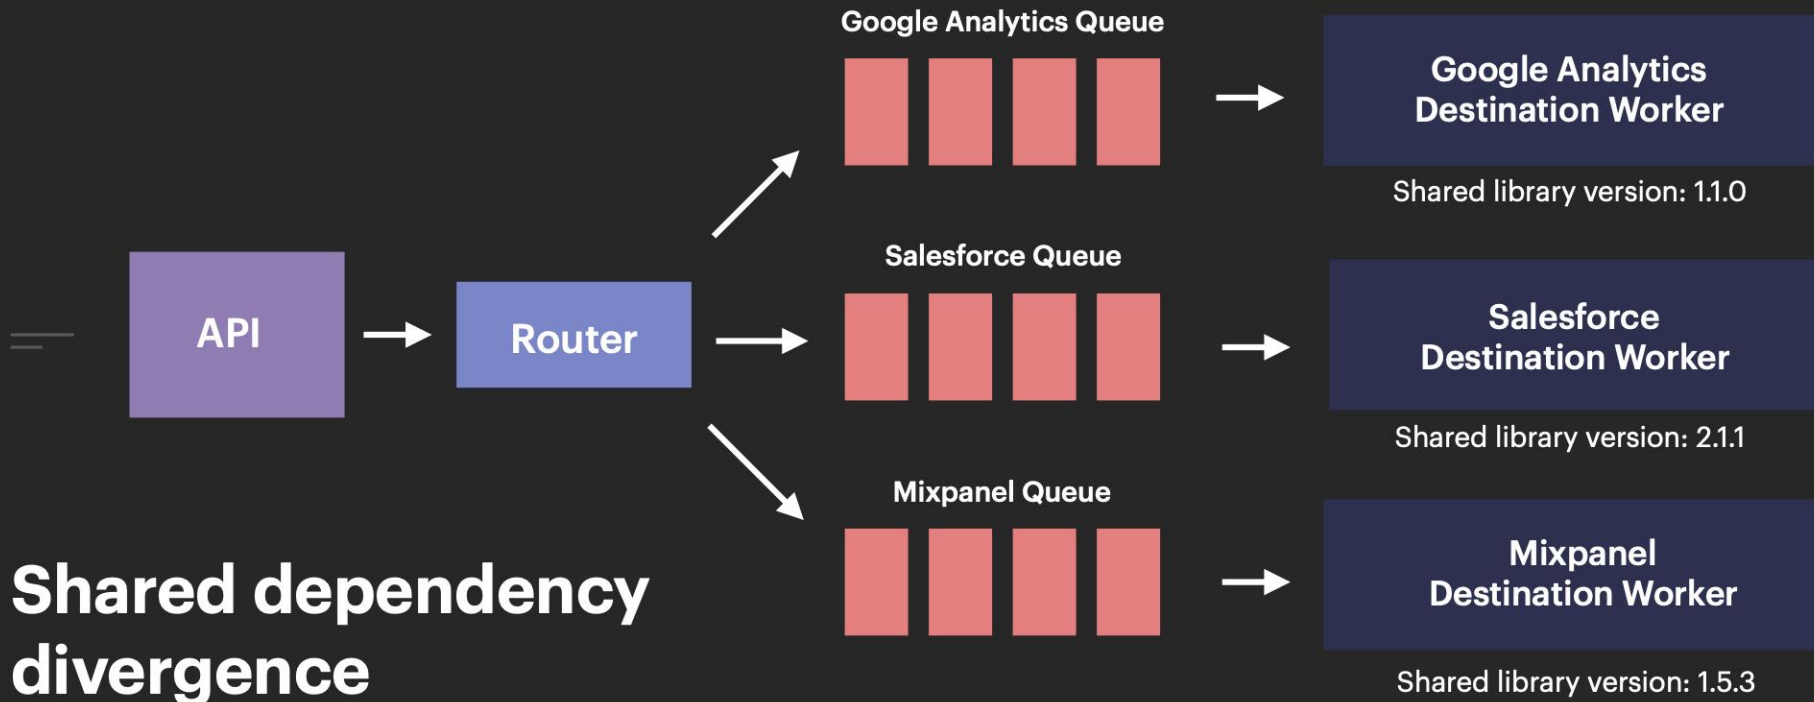


Destination Worker

Queue



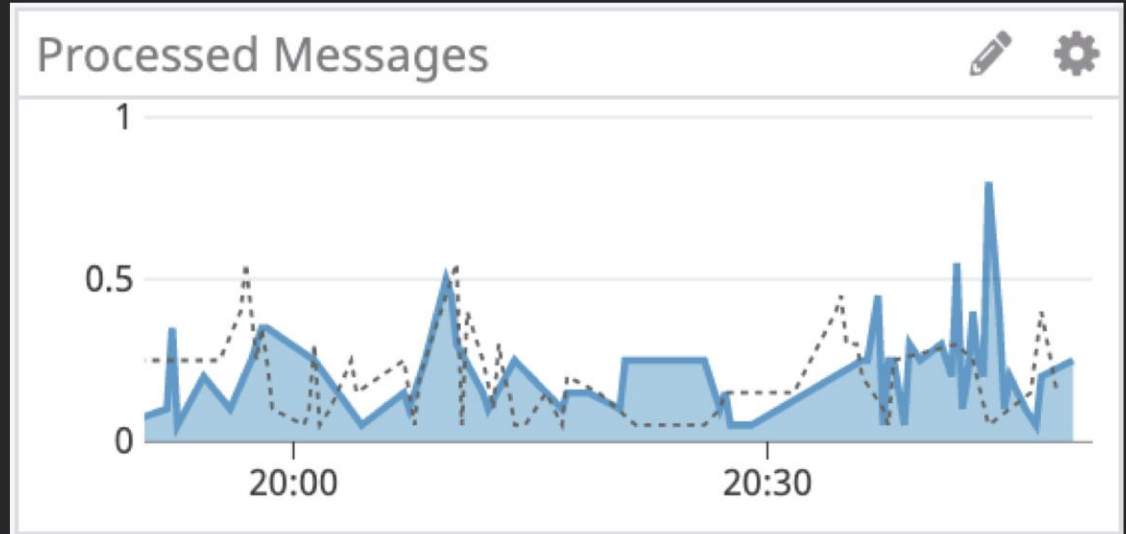
Destination Worker






<1 event/second

Destination Y



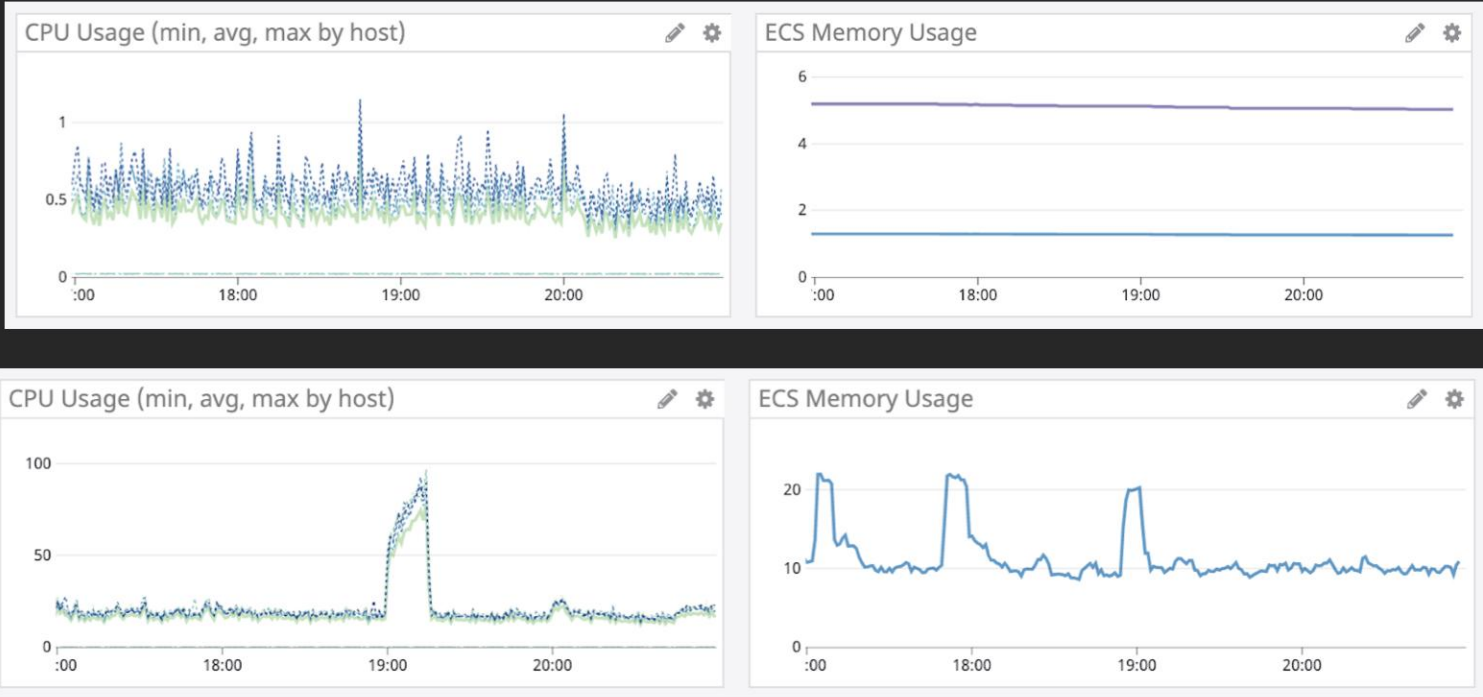


Destination Y





Auto Scaling Issues






The Breaking Point




Number of destinations



Operational Overhead



Developer Productivity



Operational Overhead



New feature development

Overview **Debugger** Schema Settings

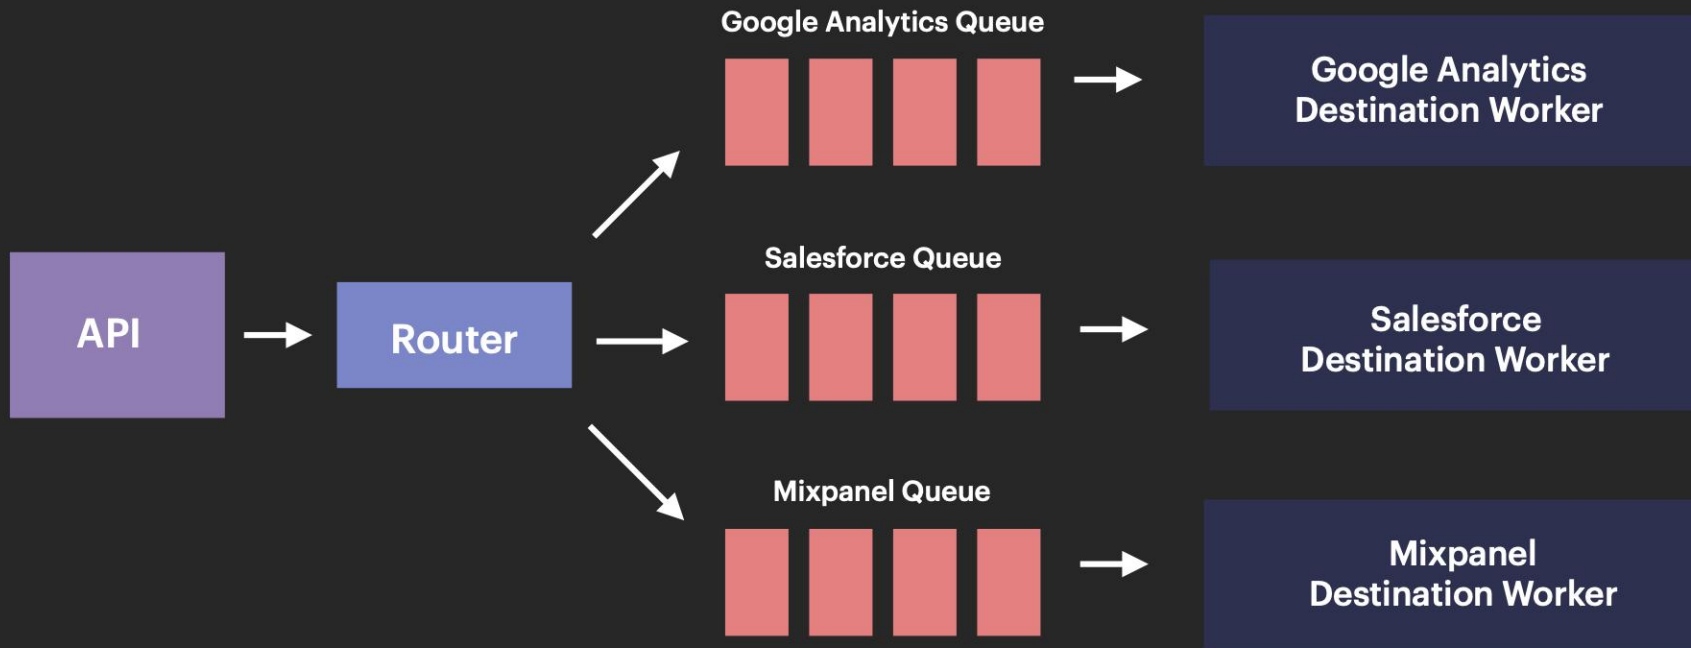
Live Pause

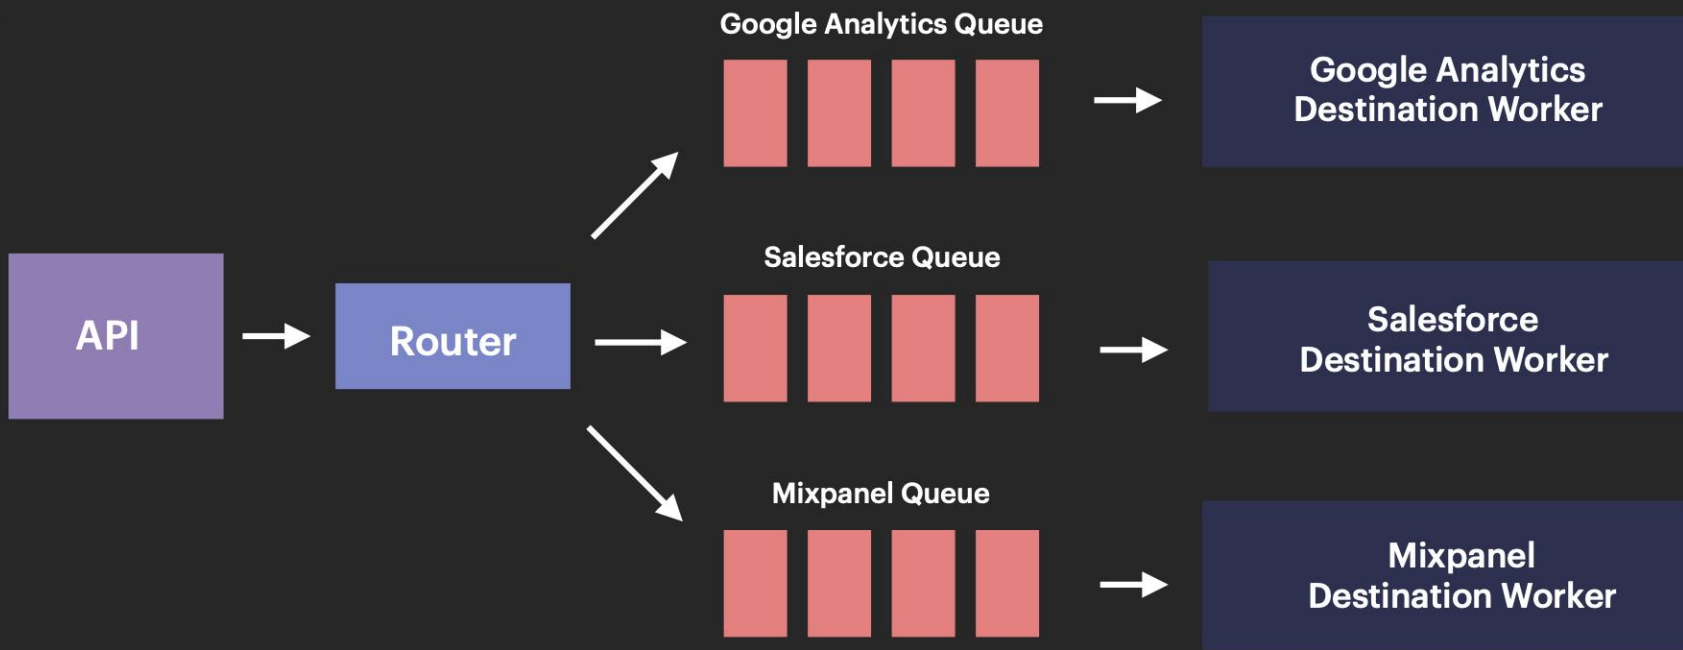
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04
<input checked="" type="checkbox"/>	TRACK	QA Event	2020/02/18 20:38:04

QA Event Allowed

Pretty **Raw** Violations

```
1  {
2    "anonymousId": "1Y0WF0dpV8RUA5WySp3LLcldzgk",
3    "timestamp": "2020-02-19T02:38:04Z",
4    "type": "track",
5    "messageId": "1Y0WEzMMTqha6NPXhCsnR2Ex0wK",
6    "integrations": {},
7    "__centrifuge_qa_id": "1Y0WF3PEnFkV7goMHZ1ISZu3iNX",
8    "properties": {
9      "pipeline": "centrifuge",
10     "qa-proper-1": 123
11   },
12   "event": "QA Event",
13   "receivedAt": "2020-02-19T02:38:04Z",
14   "sentAt": "2020-02-19T02:38:04Z",
15   "originalTimestamp": "2020-02-19T02:38:04Z",
16   "context": {
17     "library": {}
18   }
19 }
```



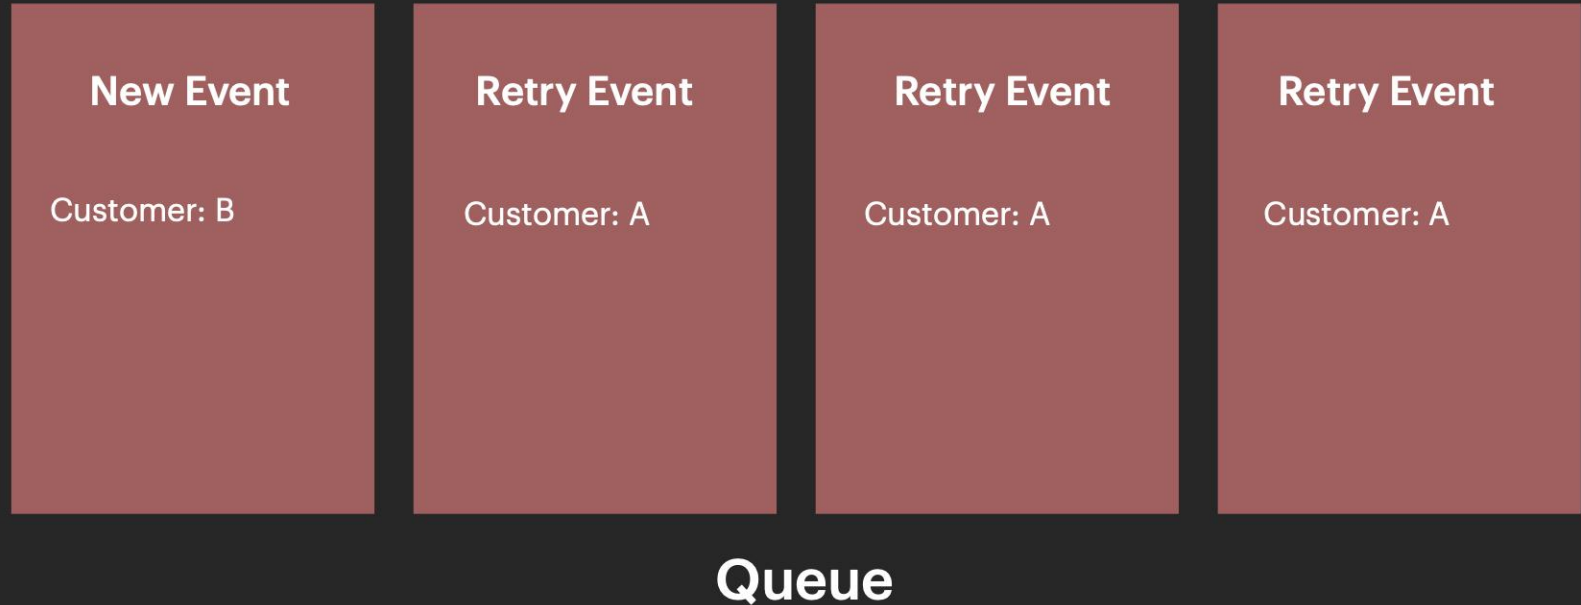


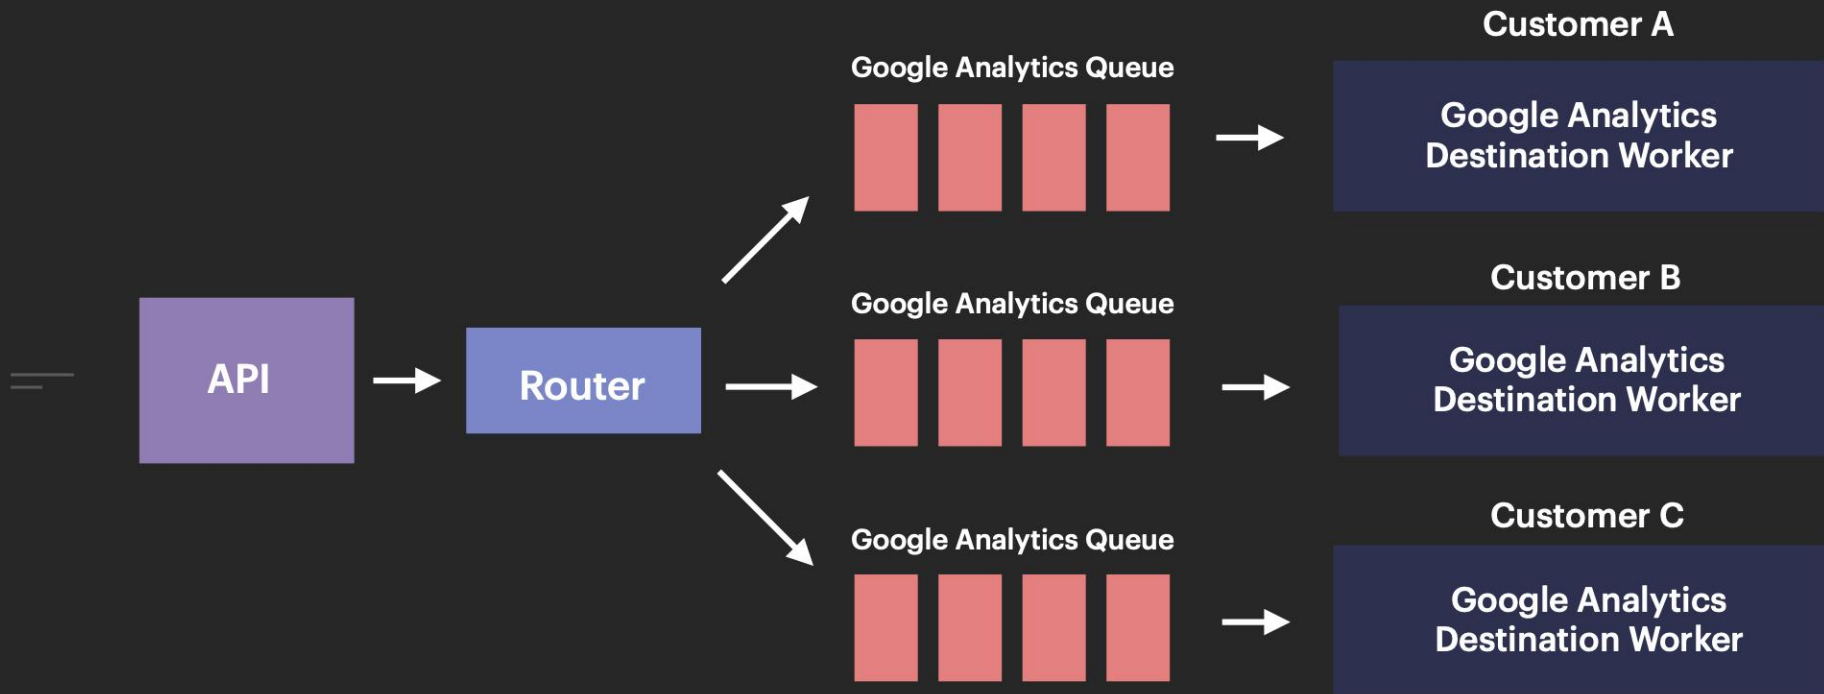
TRADE-OFFS


-  **Operational overhead**
-  **Environmental isolation**
-  **Improved modularity**
-  **Visibility**




Head-of-line blocking





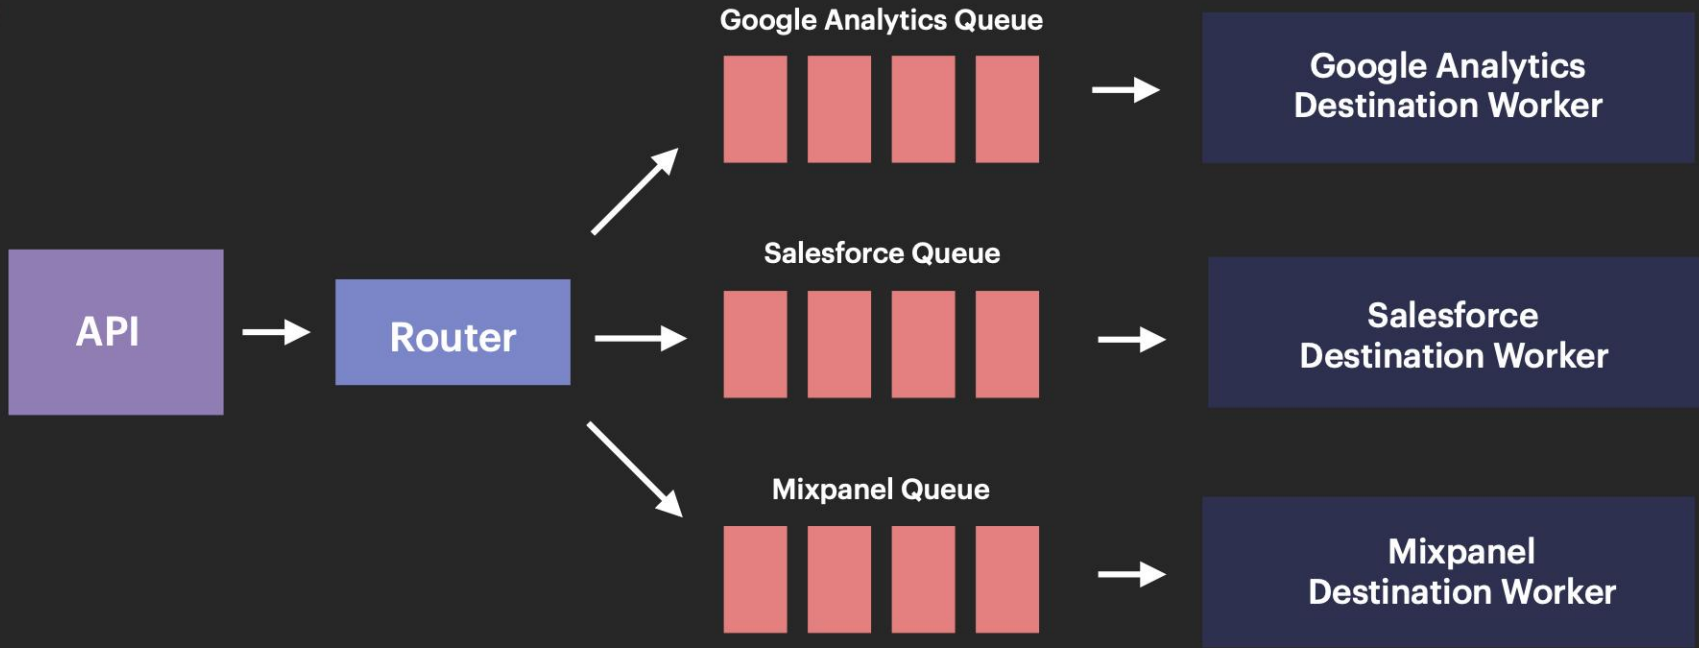




(ARTIST RENDERING)

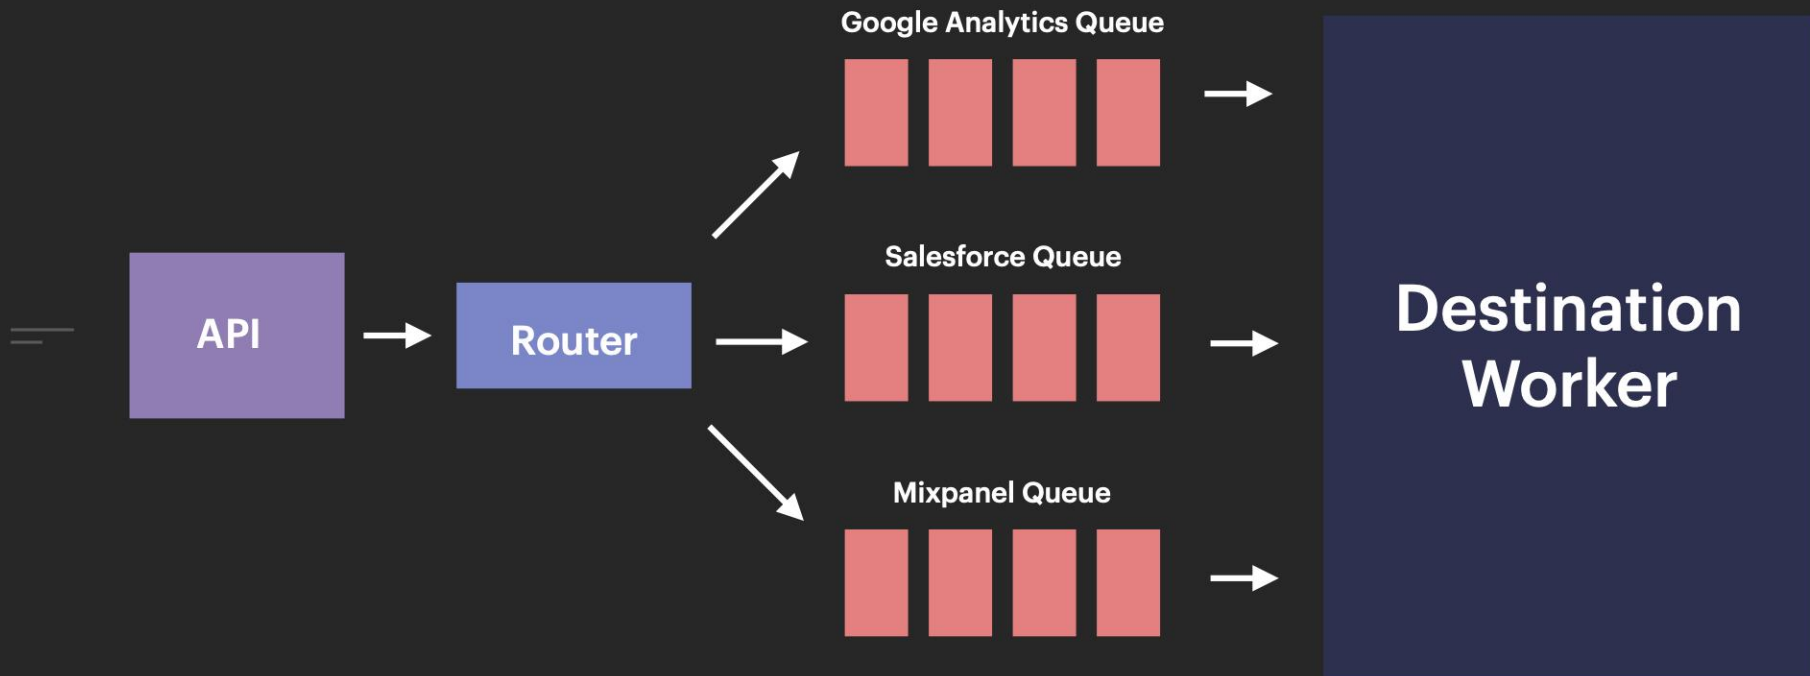


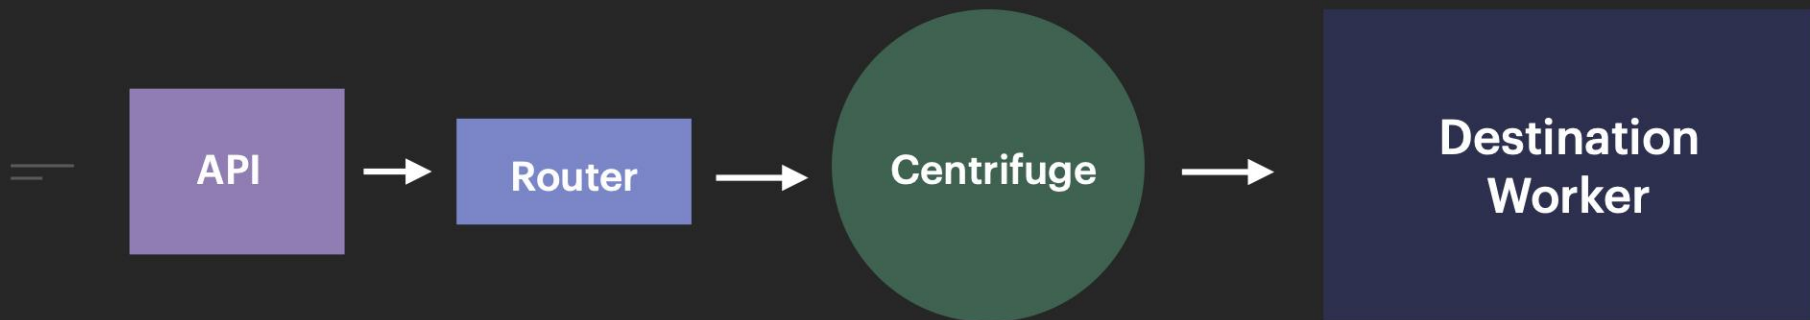
— Moving back to a Monolith



TRADE-OFFS

-  **Operational overhead**
-  **Environmental isolation**
-  **Improved modularity**
-  **Visibility**







Back to a monorepo

[/ destinations]

[google-analytics]

[marketo]

[drift]

[mixpanel]

[salesforce]

[webhooks]



Failing tests


[/ destinations]

[google-analytics]	✓
[marketo]	✗
[drift]	✓
[mixpanel]	✓
[salesforce]	✗
[webhooks]	✓



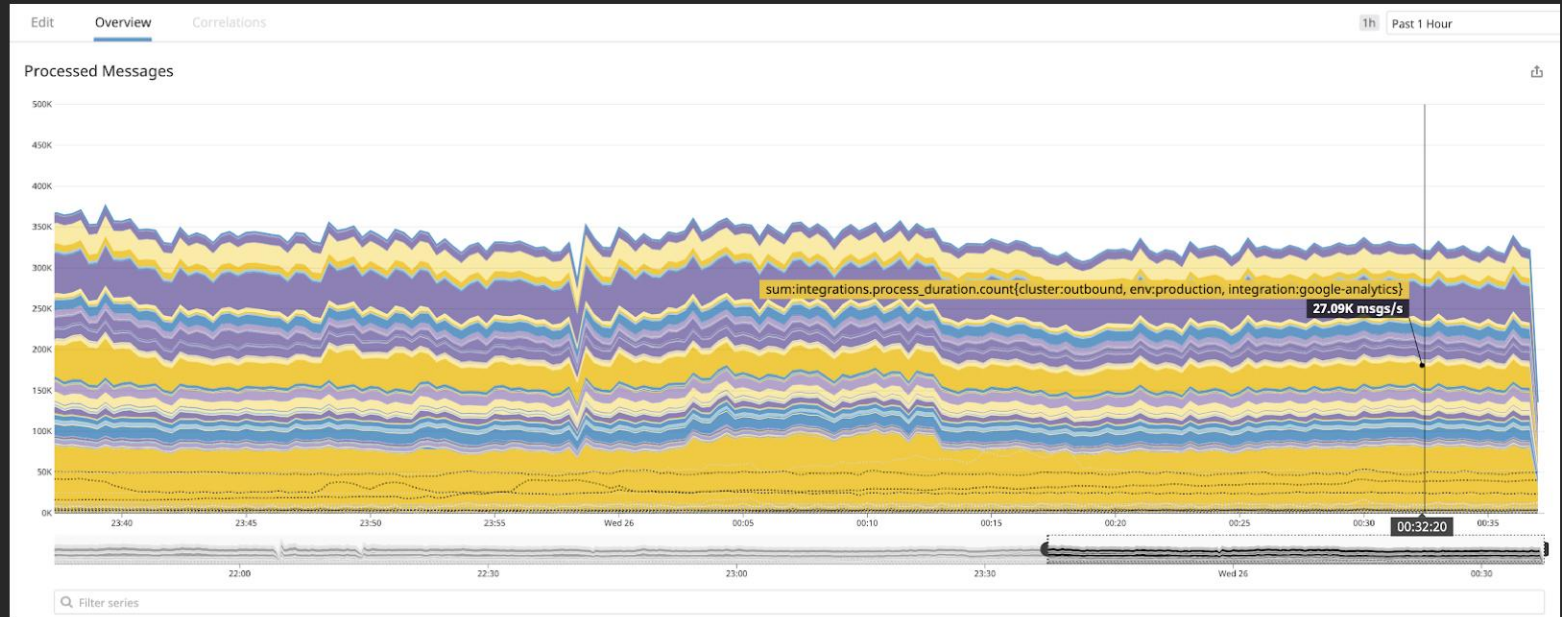
Traffic recorder

```
/*  
 * GET /pagead/conversion/1012091361/?rdid=159358&idtype=idfa&lat=1&bur  
 *  
 * host: www.googleadservices.com  
 * accept-encoding: identity  
 * user-agent: Segment.io/1.0  
 * connection: close  
 */  
  
module.exports = function (req, res) {  
  res.statusCode = 200;  
  
  res.setHeader("p3p", "policyref=\"https://www.googleadservices.com/pa  
  res.setHeader("timing-allow-origin", "*");  
  res.setHeader("content-type", "text/html; charset=ISO-8859-1");  
  res.setHeader("x-content-type-options", "nosniff");  
  res.setHeader("date", "Tue, 10 Apr 2018 16:38:22 GMT");  
  res.setHeader("server", "cafe");  
  res.setHeader("content-length", "0");  
  res.setHeader("x-xss-protection", "1; mode=block");  
  res.setHeader("alt-svc", "hq=\"googleads.g.doubleclick.net:443\"; ma  
  res.setHeader("connection", "close");  
  
  res.setHeader("x-yakbak-tape", path.basename(__filename, ".js"));  
  
  res.end();  
  
  return __filename;  
};
```






Auto-scaling issues





Developer Productivity





New product development

EVENTS DELIVERED ⓘ		EVENTS NOT DELIVERED ⓘ		LATENCY ⓘ
3,203	20	0	0	2s
Delivered on first try	Delivered on retry	Failed events	Unsupported events	P95 Latency

Delivery Trends

Event Delivery

Delivered on first try Delivered on retry Failed events



Delivery Issues ⓘ

Description	Count	Last Seen
504 Gateway Timeout	22	3 minutes ago



Dealing with the trade-offs



TRADE-OFFS



Operational overhead



Improved modularity





Environmental isolation



Visibility



Visibility





OOM debugging with sysdig

```
681370 00:17:03.650041317 3 node (22966) < write res=27 data=144 connections still open.
709526 00:17:04.649247094 1 node (22966) < write res=27 data=144 connections still open.
745016 00:17:05.650249890 5 node (22966) < write res=27 data=144 connections still open.
798750 00:17:06.650611831 3 node (22966) < write res=27 data=120 connections still open.
852165 00:17:07.650266570 1 node (22966) < write res=27 data=112 connections still open.
888505 00:17:08.650220339 5 node (22966) < write res=27 data=112 connections still open.
927071 00:17:09.659049434 6 node (22966) < write res=27 data=112 connections still open.
965320 00:17:10.658550092 6 node (22966) < write res=27 data=112 connections still open.
1002589 00:17:11.658417243 3 node (22966) < write res=27 data=112 connections still open.
1036850 00:17:12.658187124 5 node (22966) < write res=27 data=112 connections still open.
1075760 00:17:13.658146286 7 node (22966) < write res=27 data=111 connections still open.
1177228 00:17:14.716056029 2 node (22966) < write res=27 data=111 connections still open.
1299199 00:17:15.745282564 3 node (22966) < write res=27 data=111 connections still open.
1385124 00:17:16.768392181 7 node (22966) < write res=27 data=109 connections still open.
1469912 00:17:17.789551431 6 node (22966) < write res=27 data=109 connections still open.
1508855 00:17:18.789833892 0 node (22966) < write res=27 data=109 connections still open.
1536147 00:17:19.789544971 6 node (22966) < write res=27 data=108 connections still open.
1569629 00:17:20.789905914 1 node (22966) < write res=27 data=108 connections still open.
1604373 00:17:21.789517756 6 node (22966) < write res=26 data=94 connections still open.
1639905 00:17:22.790208306 2 node (22966) < write res=26 data=94 connections still open.
1749353 00:17:23.820944468 0 node (22966) < write res=26 data=94 connections still open.
1852364 00:17:24.821539565 7 node (22966) < write res=26 data=94 connections still open.
1936605 00:17:25.821801924 1 node (22966) < write res=26 data=94 connections still open.
1969927 00:17:26.821837262 4 node (22966) < write res=26 data=94 connections still open.
2012231 00:17:27.821369508 2 node (22966) < write res=26 data=94 connections still open.
2039840 00:17:28.821914211 3 node (22966) < write res=26 data=66 connections still open.
2071314 00:17:29.822263594 0 node (22966) < write res=26 data=38 connections still open.
2104017 00:17:30.823036049 3 node (22966) < write res=26 data=38 connections still open.
```



OOM debugging

▼ Request

▼ Request @1065083

- ▶ `_agent` :: Agent @109483
- ▶ `req` :: ClientRequest @826011
- ▶ `_data` :: @1066519
- ▶ `__proto__` :: Stream @240359
- ▶ `map` :: system / Map @1038221
- ▶ `_events` :: @1066515
- ▶ `_header` :: @862915
- ▶ `header` :: @862921
- ▶ `_callback` :: () @1066525
- ▶ `_timer` :: Timeout @692175
- ▶ `_formData` :: system / Oddball @51
- ▶ `_maxListeners` :: system / Oddball @47
- ▶ `domain` :: system / Oddball @51
- ▶ `url` :: "https://api-segment.intercom.io/events" @1066511
- ▶ `properties` :: (object properties)[] @1066527



Environmental isolation



Destination Worker

Destination Worker

Destination Worker

Destination Worker

Destination Worker

Destination Worker

Destination Worker

Destination Worker



Uncaught exceptions

```
try {  
    return await Promise.all(requests)  
} catch (e) {  
    throw e.response.error  
}
```



Environmental isolation





Environmental isolation






Back pressure





Wrapper





TRADE-OFFS



Operational overhead



Improved modularity

Traffic recorder




Environmental isolation

Centrifuge, Wrapper



Visibility

Heapdump on hosts






TRADE-OFFS



Environmental isolation



Operational overhead





TRADE-OFFS



Improved Modularity



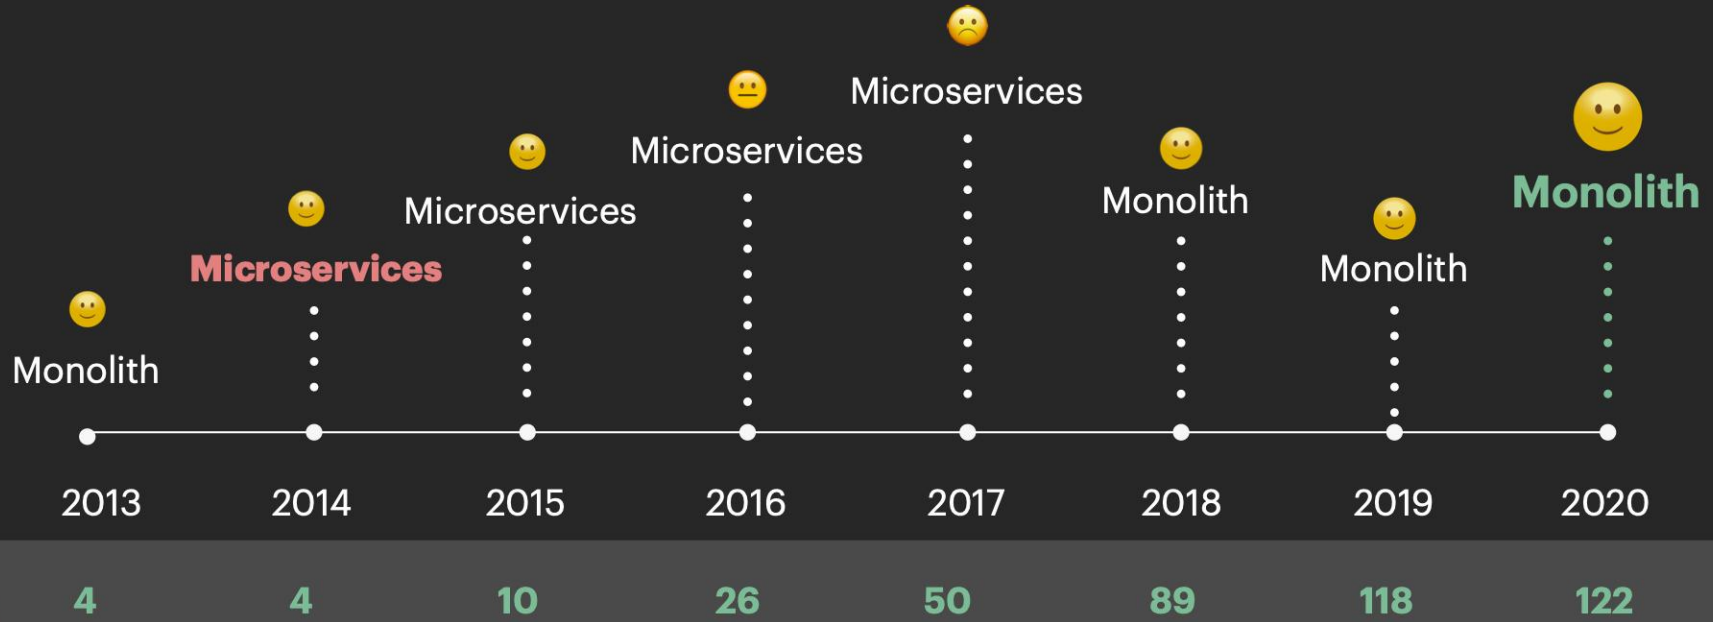
Visibility



Environmental isolation




Operational overhead





TRADE-OFFS

- ✓ Improved Modularity
- ✓ Environmental isolation
- ✓ Visibility
- ✗ Operational overhead






TRADE-OFFS

Improved Modularity

Visibility

Environmental isolation

Operational overhead






TRADE-OFFS

— Improved Modularity

✓ Visibility

✓ Environmental isolation

✓ Operational overhead





Takeaways

No silver bullet

Have a good
story around
trade-offs

Solve the root issue



— Thank you.



To Microservices... and back again

Alexandra Noonan

alexandra@segment.com

 anoonan