




Serverless eStatements

Millions of PDFs once a month




The Problem

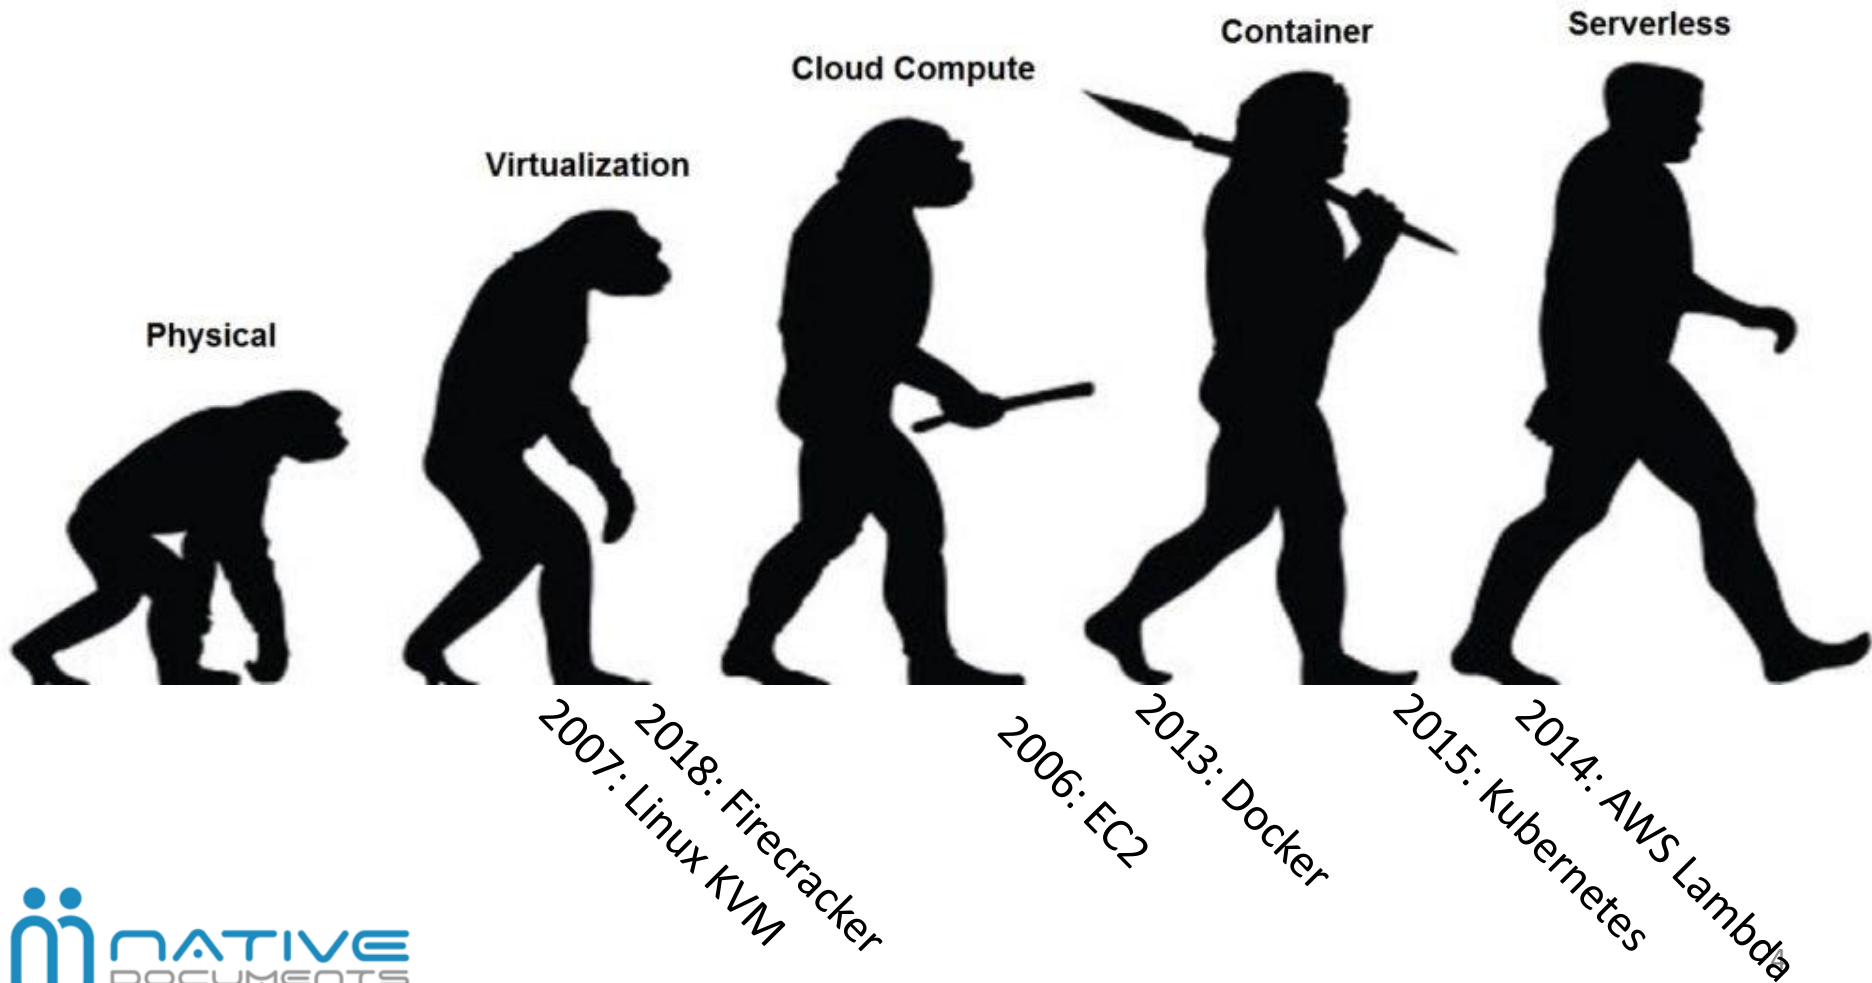
- Docx to PDF conversion is CPU intensive
 - Depends on “complexity” (number of pages, tables, page breaks etc)
- so producing lots has meant provisioning servers
- and that’s expensive: hardware, software and people



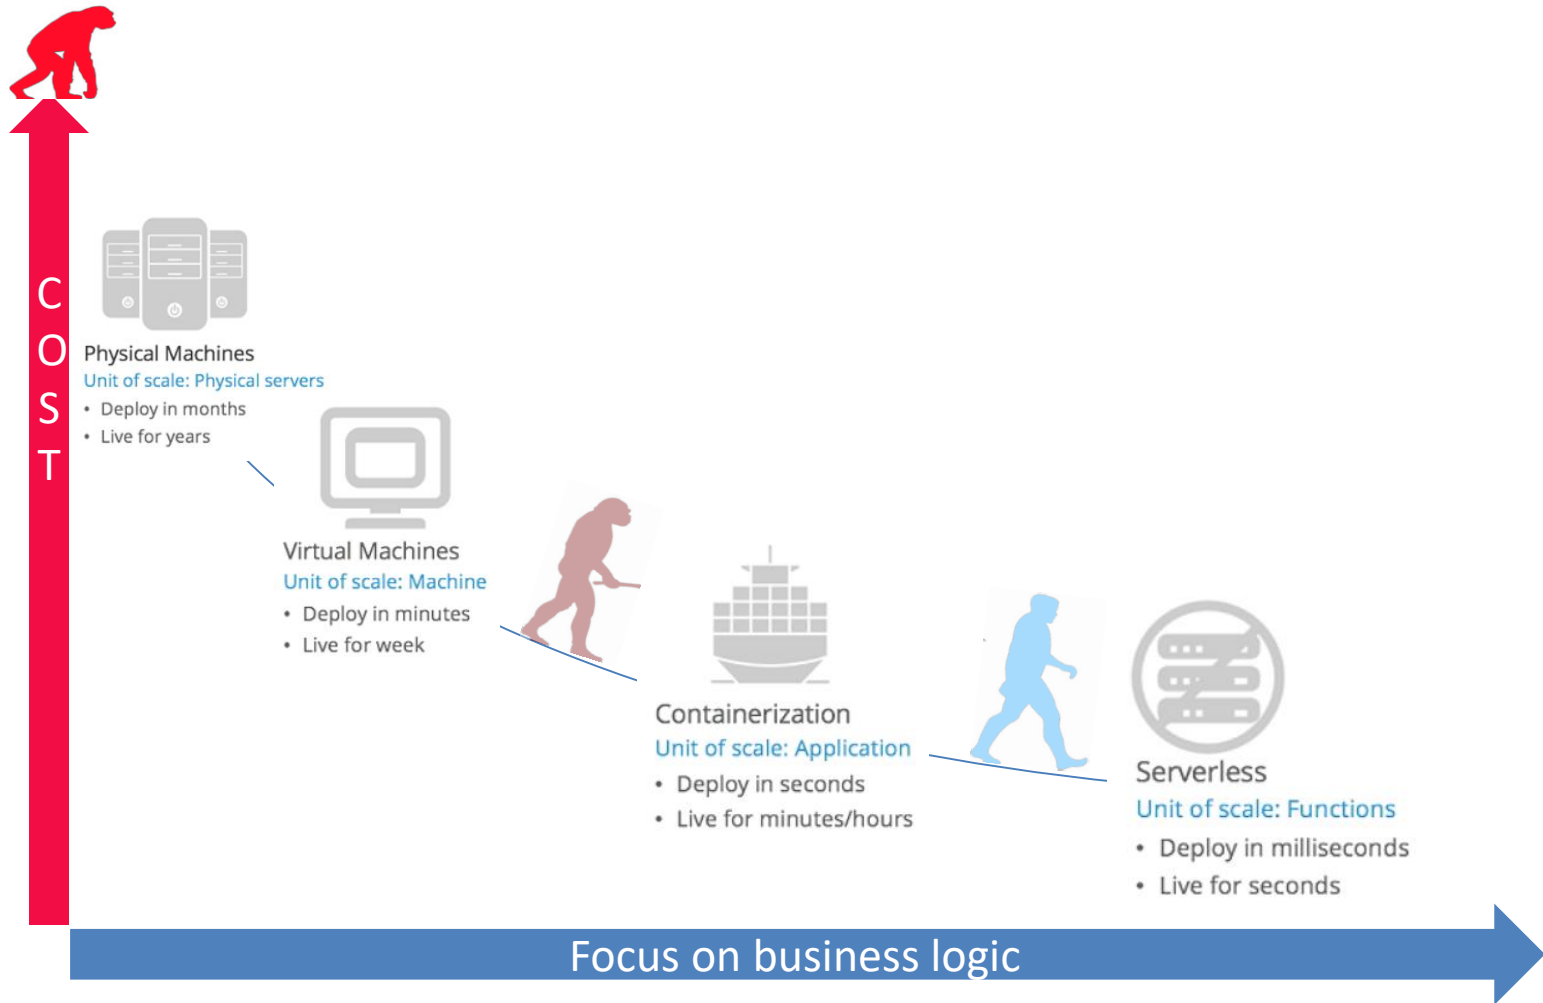
The Solution: “serverless”



Serverless: its the next step in the evolution of computing



Serverless: Focus on *functions* implementing your business logic



All the big providers now offer it

Walled Gardens:

- **AWS** Lambda
- **Azure** Functions
- **Google** Cloud Functions
- IBM, Oracle...

Cross platform

- Serverless Framework

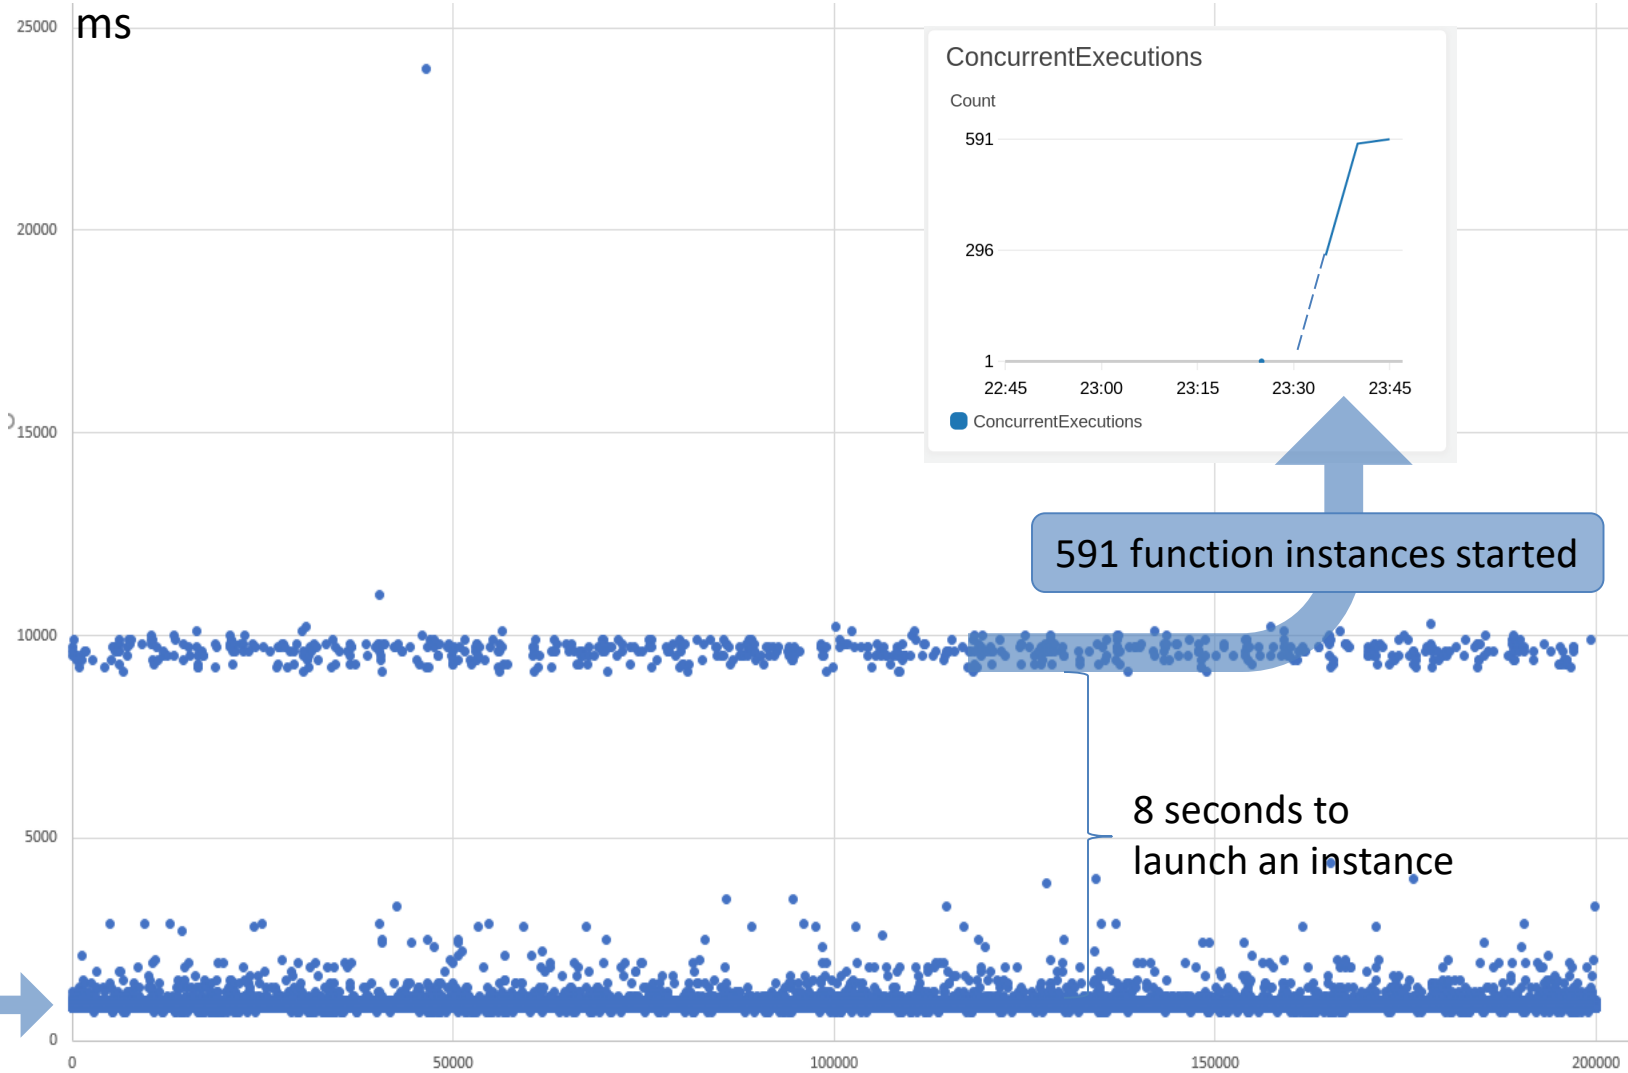
Serverless deployment of Docker containers:

- OpenWhisk (deploy anywhere)
- For Kubernetes:
 - Deploy a Docker Container:
 - in Knative
 - OpenWhisk
 - OpenFaaS


Concepts as implemented in AWS Lambda

- Common concepts:
 - Trigger events
 - Scheduler executes Function
- Languages
 - Javascript (node.js)
 - Others, depending on platform
- Function packaging
- Lambda event triggers include:
 - Your REST API
 - S3
 - SQS (messages)
 - Step functions
 - *etc.*

200,000 PDFs in 10 mins with Native Documents on AWS Lambda




Compared to earlier paradigms



Problems overcome:


- Slow start time if a server instance needs to start
- Effort to configure server auto-scaling



AWS Lambda cost for those 200,000 PDFs:
175,000 seconds
= ~\$5

How? You need “serverless-ready” PDF Conversion code

- Needs to be able to run on the platform:
 - Lowest common denominator is node.js
- Ours, for example works as follows
 - C/C++ code base
 - Proprietary doc/docx layout/editing engine
 - Exports PDF using Skia (Google 2D graphics lib) PDF backend
 - Converted to Web Assembly (wasm) using Emscripten



- <https://www.npmjs.com/package/@natedocuments/docx-wasm>

Web Assembly is also a great foundation for editing Word docs

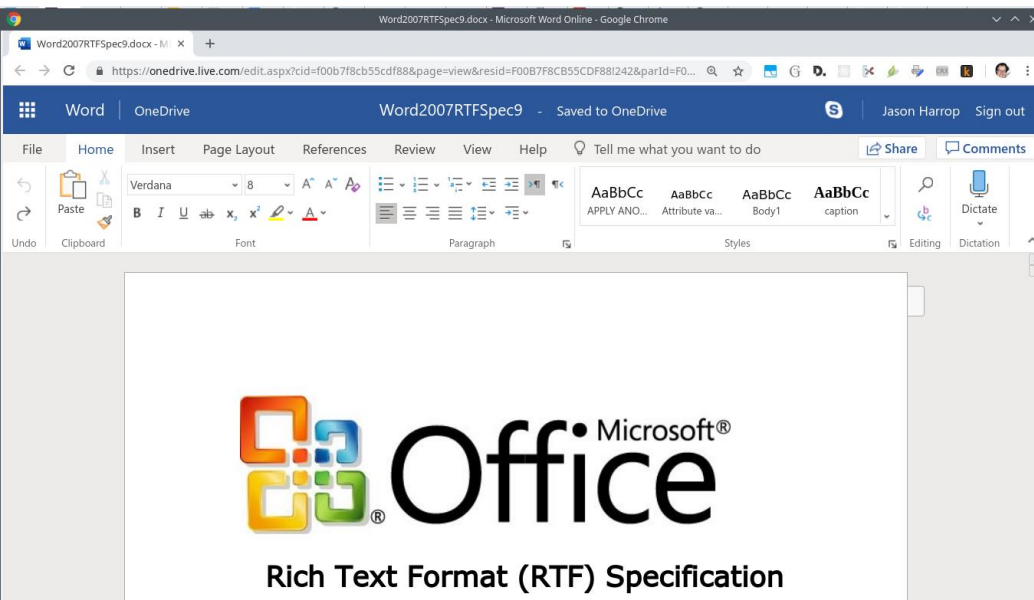
- Docx page layout
 - Must be done (accurately!) to generate a PDF
 - Done by Native Documents rendering engine
 - Rendering engine also used in Word File Editor
- Thanks to wasm, we run that rendering code in the browser
- Compared to Word Online (and Google Docs):
 - wasm approach uses under half the resources
 - User perceives better performance on long documents
- For example, the RTF spec:

In Word Online:-

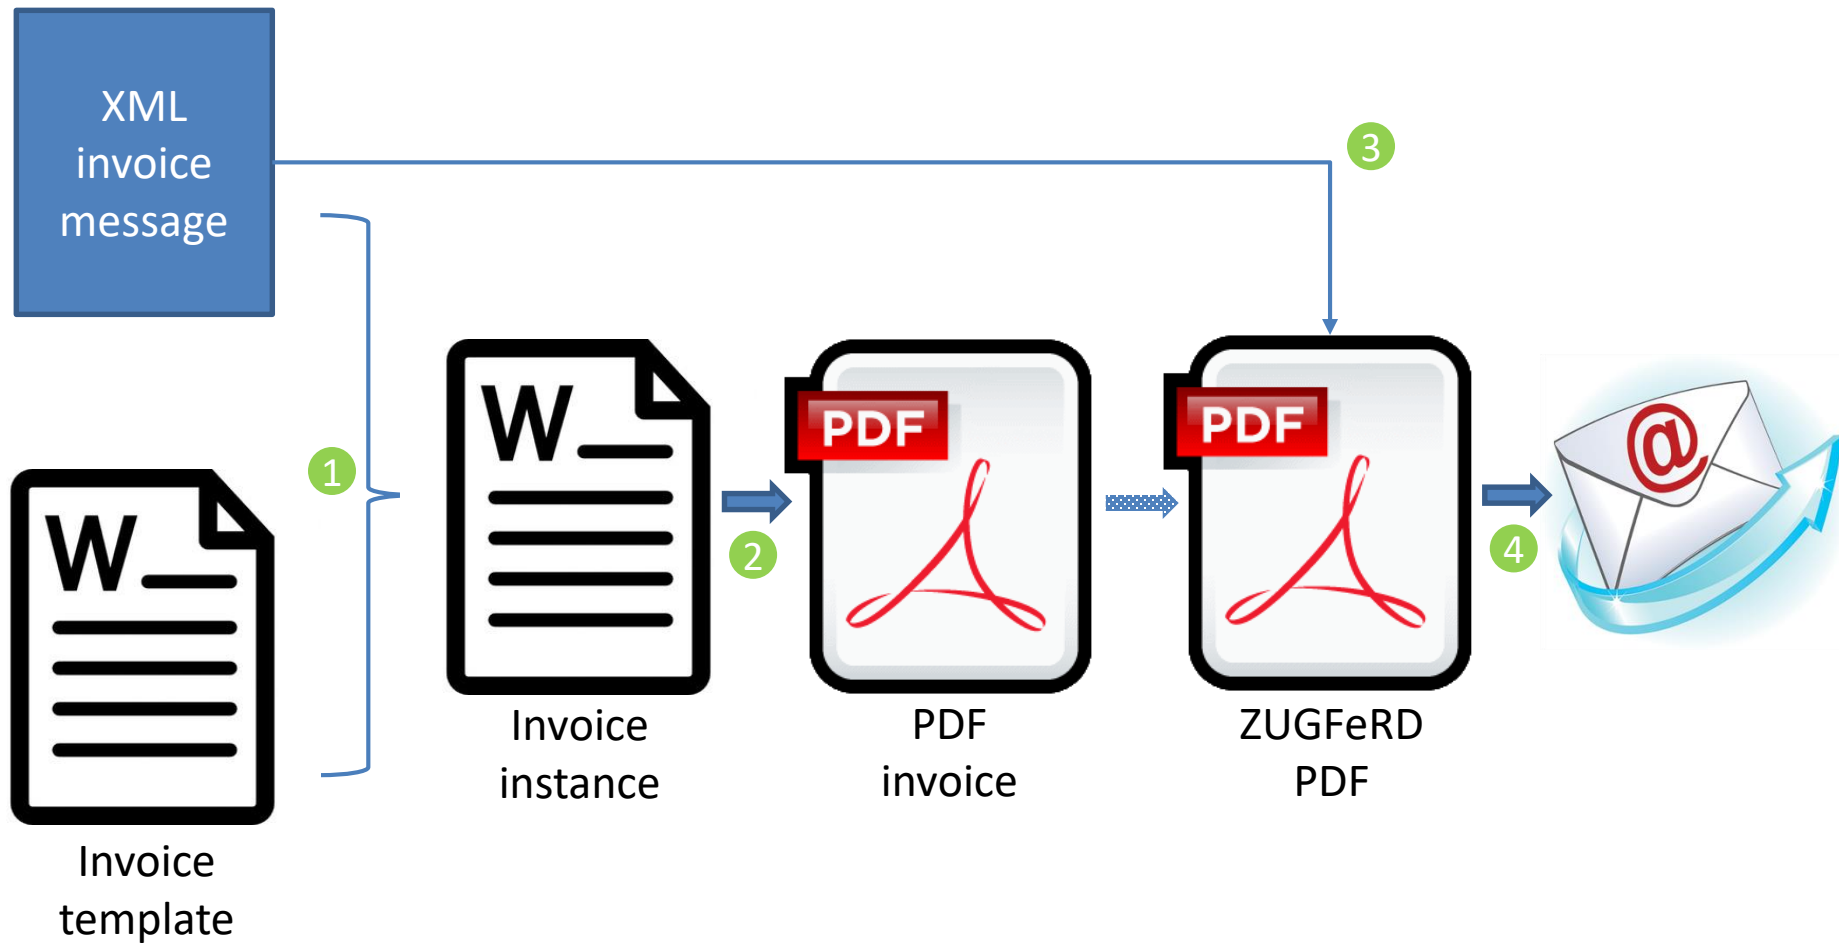
Memory footprint	GPU memory	JavaScript memory	CPU
711,008K	11,202K	536,164K (413,380K live)	4.0
<i>Scroll to the end:-</i>			
1,232,996K	11,202K	684,644K (585,490K live)	14.0

compare Native Documents:-

283,684K	48,148K (37,860K live)
<i>Scroll to the end:-</i>	
425,428K	92,180K (80,576K live)



Case study: High volume serverless ZUGFeRD PDF invoices



4 functions to be orchestrated here

But is eInvoicing really high-volume?

- Generally eInvoices are required in government procurement
 - Low volume for most sellers
 - Except intermediary hubs/services (billing service providers)
- (Compare eStatements)
- Increasingly corporate customers demand an eInvoice
 - Facilitates reconciliation/analysis of electricity bills
- Easy then for utilities to provide consumers with an eInvoice
 - Consumed by:
 - Personal finance software
 - Online banking
 - Government tax authority
 - A FinTech opportunity? Maybe..

How best to wire micro-services together?


State machine definition

Define your state machine using the Amazon States Language (ASL), and review the visual representation of your workflow. [Learn more](#)

Generate code snippet

[Learn more](#)

```
1 {
2   "StartAt": "PreMergeValidation",
3   "States": {
4     "PreMergeValidation": {
5       "Type": "Task",
6       "Resource": "arn:aws:lambda:us-east-1:123456789012:function:PreMergeValidation",
7       "Next": "FetchFrontMatter"
8     },
9     "FetchFrontMatter": {
10      "Type": "Task",
11      "Resource": "arn:aws:lambda:us-east-1:123456789012:function:FetchFrontMatter",
12      "Next": "MergeData"
13    },
14    "MergeData": {
15      "Type": "Task",
16      "Resource": "arn:aws:lambda:us-east-1:123456789012:function:MergeData",
17      "Next": "PDFConvert"
18    },
19    "PDFConvert": {
20      "Type": "Task",
21      "Resource": "arn:aws:lambda:us-east-1:123456789012:function:PDFConvert",
22      "Next": "Deliver"
23    },
24    "Deliver": {
25      "Type": "Task",
26      "Resource": "arn:aws:lambda:us-east-1:123456789012:function:Deliver",
27      "End": true
28    }
29  }
30 }
```



One state machine per document processed

Case study 2: doc gen in Salesforce.com

The screenshot displays the iDialogue Admin interface for configuring a document room. The top navigation bar includes 'iDialogue Admin', 'Home', 'Quick Start', 'Files', 'Room Builder', 'Guides', 'Document Rooms', and 'Room Events'. The main header shows 'Marketing Materials' with 'Publish' and 'Help' buttons. A 'Desktop' view selector and a 'Save' button are also present.

The main content area is titled 'eStatement Portal' and 'Document Room'. It contains a welcome message with placeholders for member and owner names, a list of documents including 'samplePDFDoc.pdf', and a file upload section with an 'Upload Document' button and a 'Drop File' area. A 'Document Plugin' button is also visible.

The right sidebar contains configuration options for the document plugin, including 'Content Library' (set to 'Document Room'), 'Documents' (set to 'Select an Option'), 'Display Name' (set to 'Document Plugin'), 'Share With Room Members' (checked for Administrators and Collaborators), 'Object Source' (set to 'ContentDocument'), 'Export/Download Format' (radio buttons for Word and PDF), 'Document Journey' (checked for 'Enable Status Button'), and 'Completed Event' (radio buttons for 'Opened/Viewed', 'Approved', 'Submitted', and 'eSignatures').

Serverless architecture scales from 0 to N documents.

Compare the previous architecture, which required up to 20 servers to execute full doc gen batches.

Lessons learnt

- Faster can sometimes be cheaper
- Web assembly:
 - Production-ready
 - Great if you have C/C++ code
 - Great fit for serverless
- Concern that cloud APIs and sensitive documents don't mix
 - Serverless makes DIY easy
- Choose serverless-ready tech
- Choose serverless-ready business models
 - Per-core/socket/CPU pricing doesn't fit
- Where is your bottleneck now?
- Be aware of cloud vendor lock-in

Thank you!

We appreciate your participation.

jason.harrop@natedocuments.com