Building a Large-Scale Wagtail Hosting Platform

Vince Salvino – @vincesalvino

CodeRed

The Aviary

What you want your multi-client hosting environment to look like...

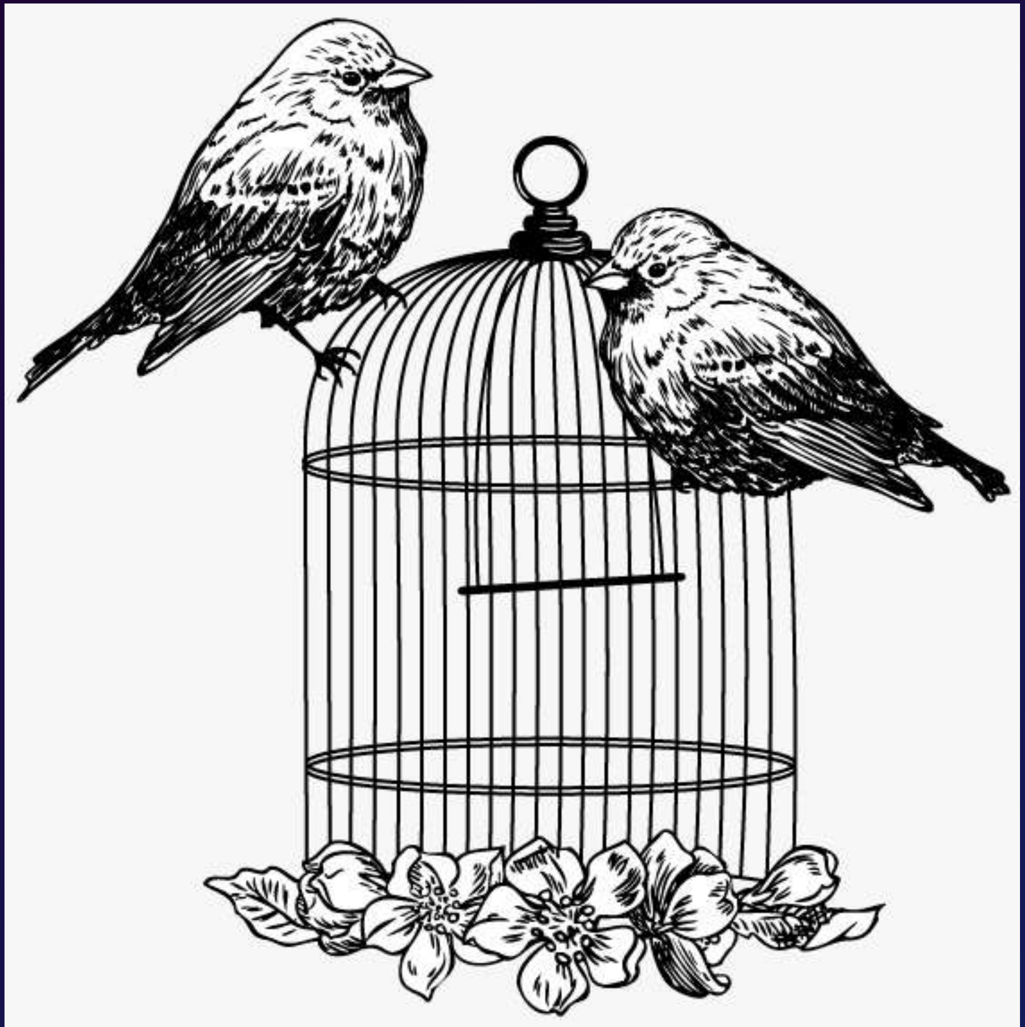

The Aviary

...versus what it will quickly descend into.

Common Wagtail Hosting Setups

“The Old School”

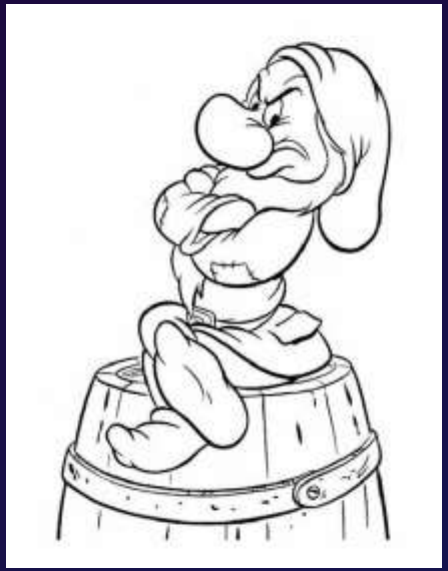

*You’ve been around the block. You’ve seen some s**t.*
If it ain’t broke, don’t fix it!

- Linux machine
- Apache + mod_wsgi
- MySQL or Postgres (probably running on localhost)
- Python installed from the system package manager, like any god-fearing true-blue American. It’s all gone downhill since 2.7 anyways.

“The Captain”

“For there is no folly of the beast of the earth which is not infinitely outdone by the madness of men.” –from Moby-Dick

- Everything is Docker.
- Python container running a wsgi server.
- Front-end web server that also talks wsgi.
- Database container.
- 3rd party object storage (static/media).
- A few more containers for the heck of it.

“The Pioneer”

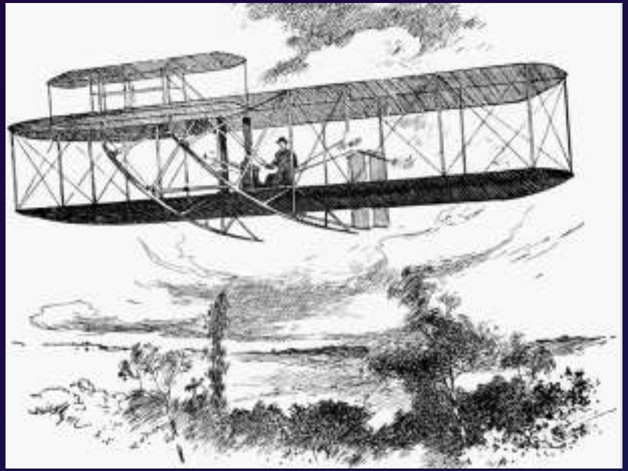

“With 12 horsepower at our command, we considered that we could permit the weight of the machine with operator to rise to 750 or 800 pounds.” –Orville Wright

- Total cloud-based infrastructure.
- Git → CI → Container Registry → Deployment → Network gateway → Reverse proxy → Load balancer → Certificate issuer → ...
- Database service, Redis/Memcached service, Object storage, probably Kubernetes in there somewhere...
- Everything but an actual Linux VM.

Now let’s scale it up…

100 x “The Old School”

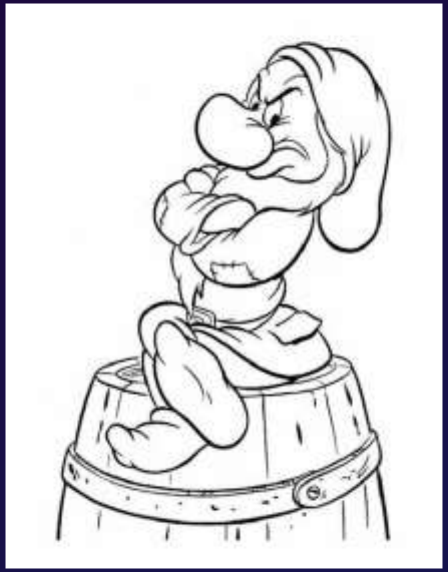

- You now must manage 100 virtual machines.
  - Not too bad… this is a problem with very mature solutions (Ansible, Chef, Puppet, etc.)
  - Slow – bring a VM up or down can take many minutes.
- No economies of scale… the price of a VM is the same whether you buy 1 VM or 100 VMs. Reserved pricing becomes even harder unless you have all your customers on yearly contracts.

100 x “The Captain”

- You now must manage 100 (or possibly 200, 300, 400) containers.
  - Containers are very manageable, robust tooling available.
  - But you are still incurring a 3-4x number of management subjects compared to number of websites.
- Media files are still a problem, you now need to manage 100 separate object stores.
- Solutions such as Kubernetes or Swarm are of no use to you… they are design around the idea of running a fleet of a single app, not a fleet of 400 different apps.
- Economies of scale… maybe. If you run containers of a VM, you benefit from buying bigger VMs. If you run containers on a cloud service, cost goes through the roof.

100 x “The Pioneer”

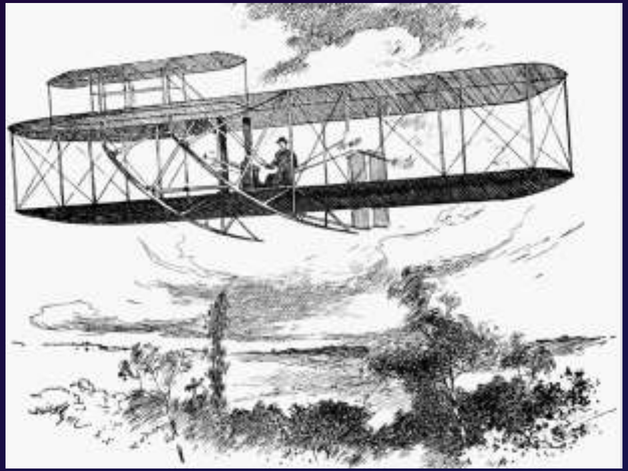

- You now need to manage 100 x however many cloud services you had before. If you’ve got 12 services per app, you now have 1,200 subjects to manage!
  - But… some services are designed to scale up and are easy to manage. Others not so much.
  - You now need to start writing YAML full time in order to manage these cloud services.
- Economies of scale… no. Like the VM approach, pricing is flat and metered based on usage. But costs can be kept low depending on which services are used.

So how to host 100 Wagtail sites?

…nobody wins 😭

Anatomy of a Wagtail

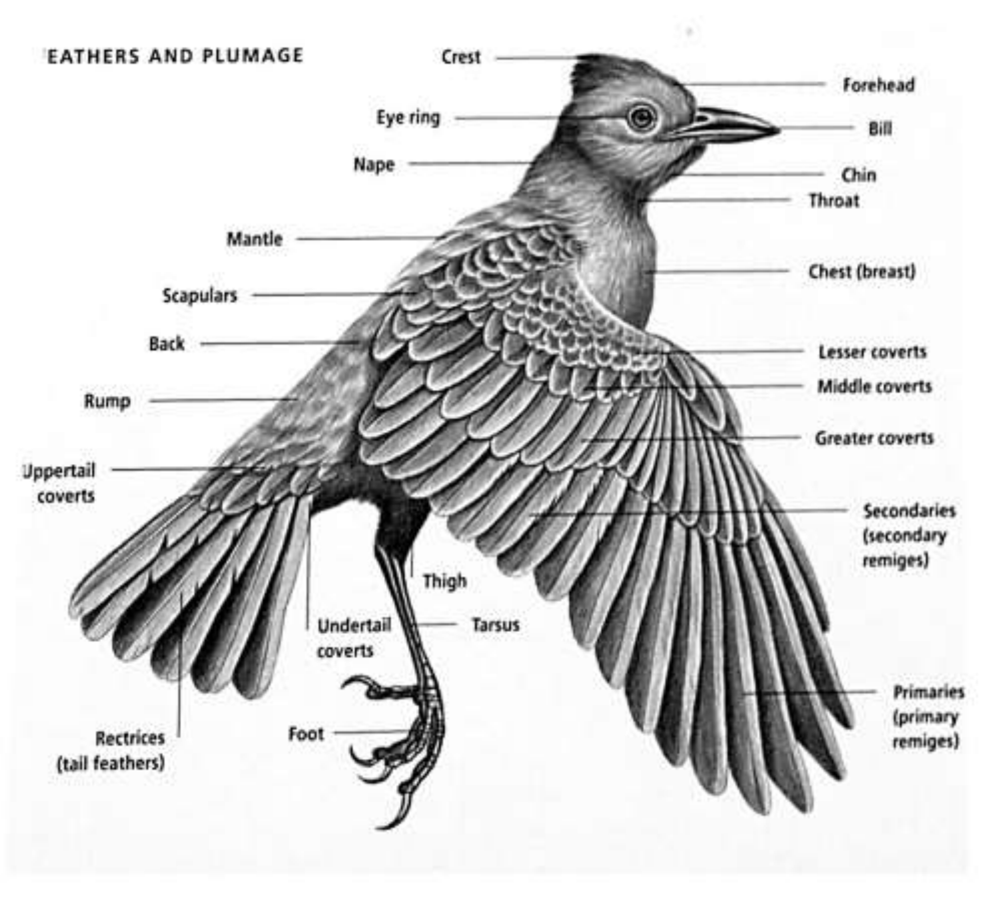

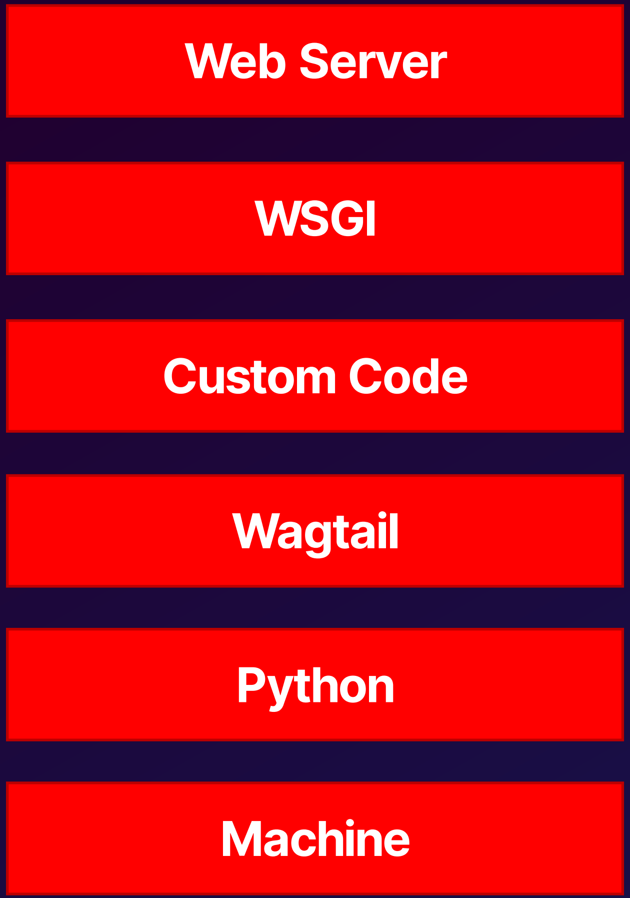

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

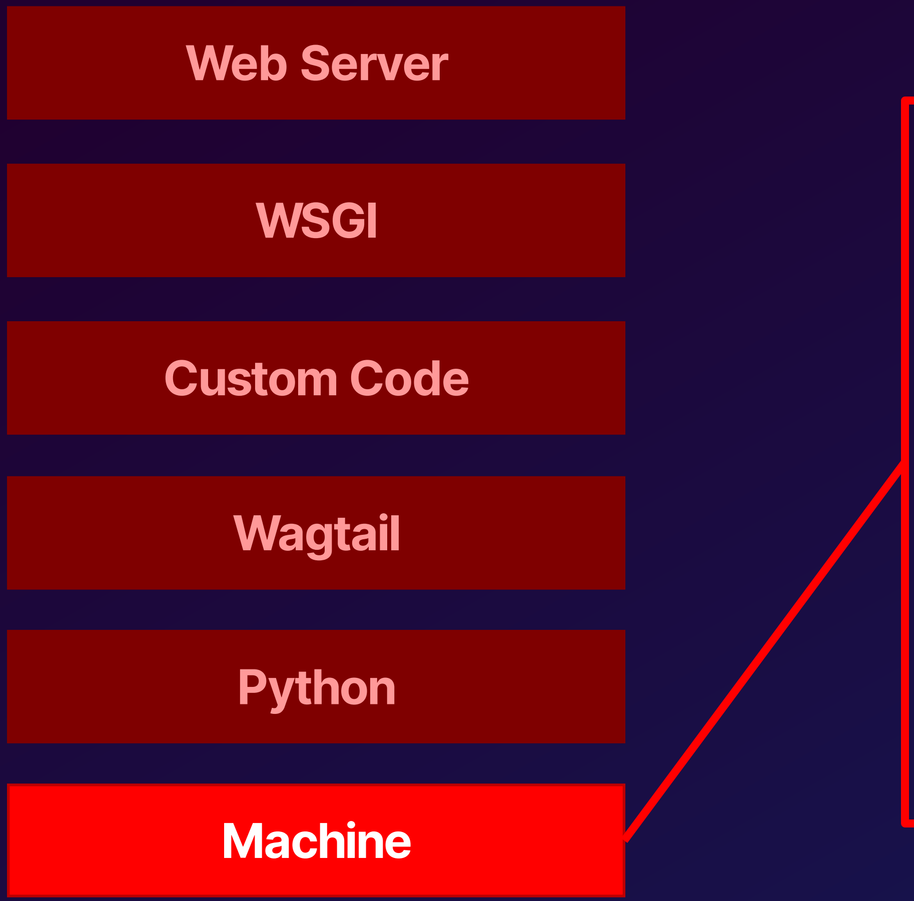

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

- Convenient tools
  - Cron jobs
  - Fast persistent filesystem
  - Strong user/file permissions
- Almost zero overhead (VM)
  - Management overhead

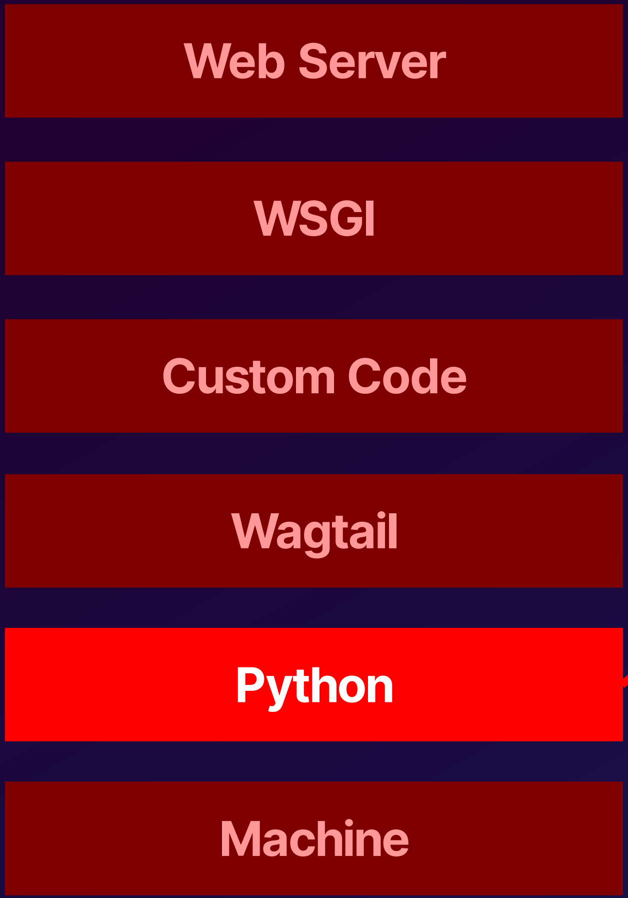

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

- Python’s downfall is (kind of) our benefit
  - GIL = predictable resource usage
- Very low memory usage
  - < 10 MB

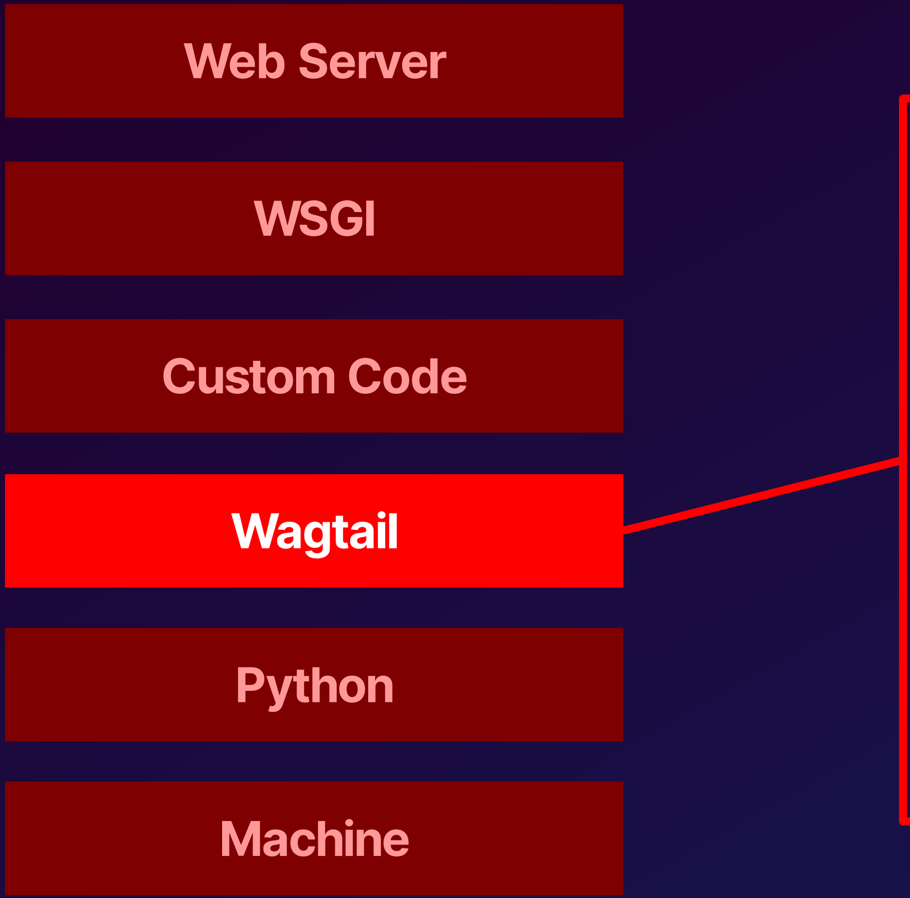

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

- ~100 MB mem usage
- 1 request per process, ~100ms response time.
- Media files
  - Requires them but cannot serve.
- Well behaved and mostly bug-free

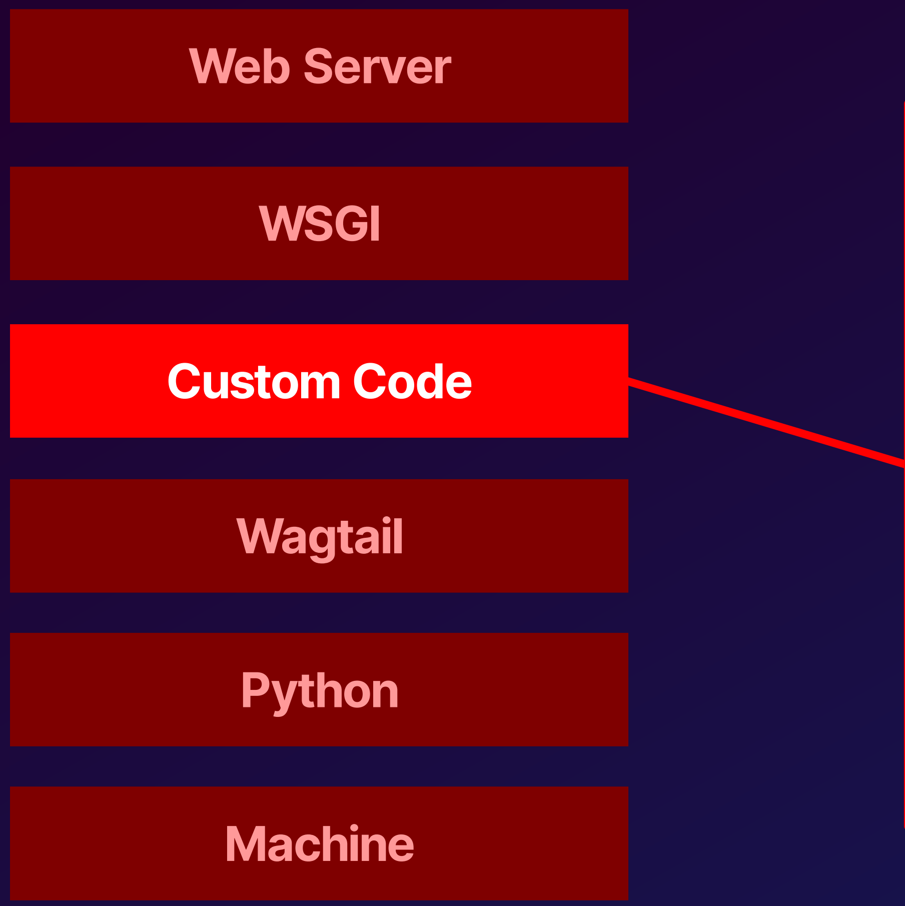

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

- Total wildcard
- Accesses:
  - RAM
  - CPU
  - Filesystem
  - Network
- If hacked, could wreak havoc!

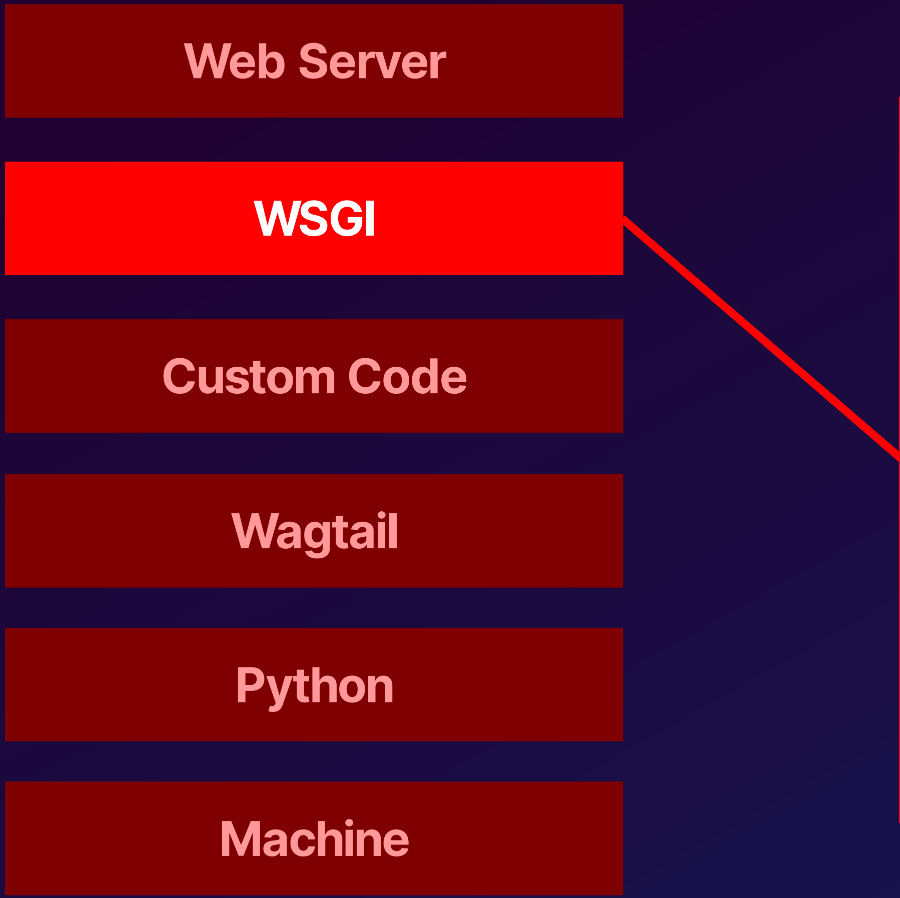

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

- Turns Python into a website.
- Boilerplate – devs don’t care!
- Proctors the resources of the entire stack below.
- Can be a source of massive resource savings!
    - Requires difficult fine-tuning.

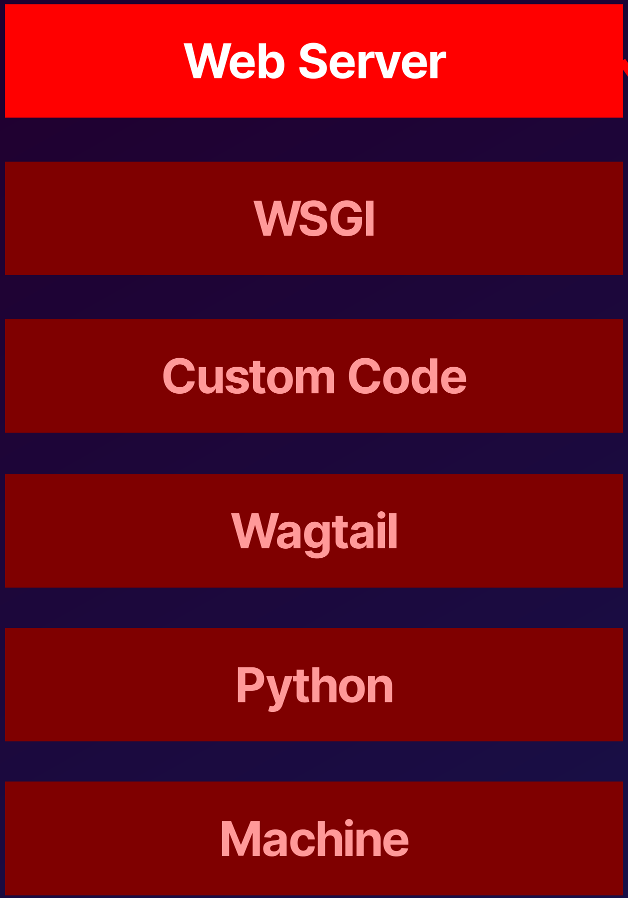

Web Server

WSGI

Custom Code

Wagtail

Python

Machine

- Apache / NGINX
- Boilerplate – devs don’t care!
- Highly customizable
  - Caching
  - Proxy
  - Rewrite
  - SSL

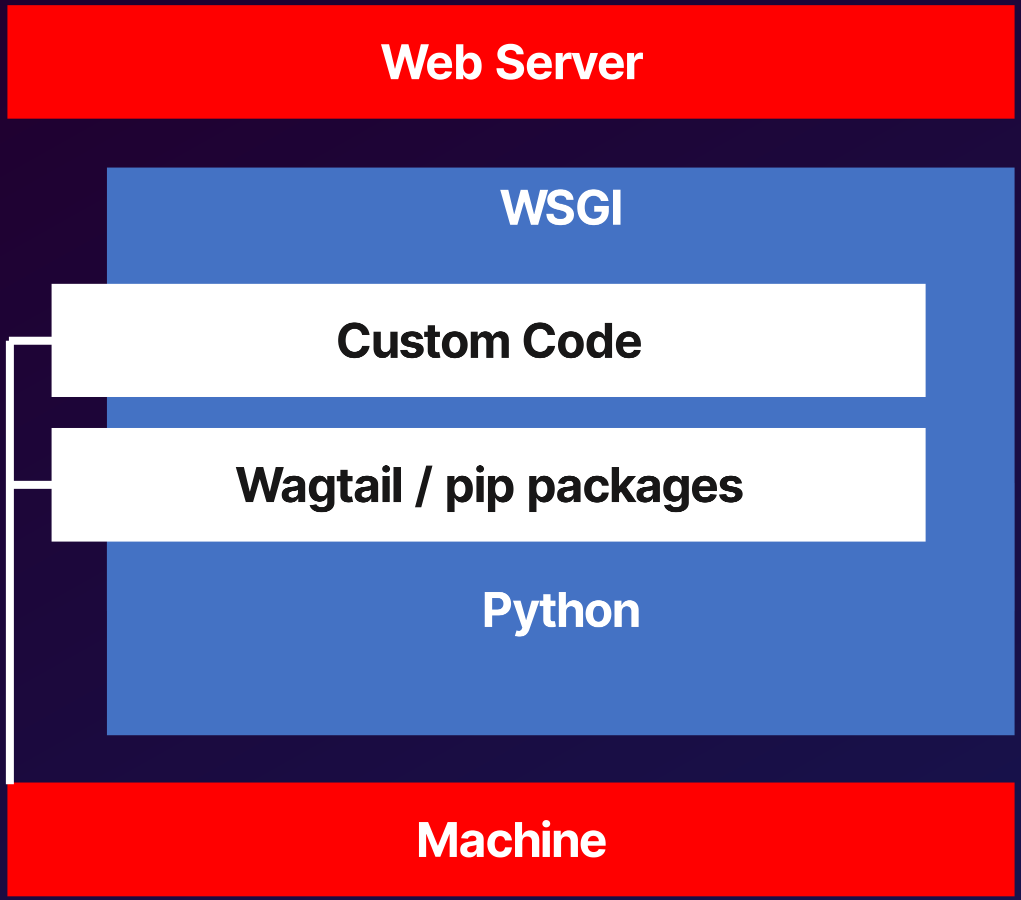

One native high-performance web server / reverse proxy.

Condense WSGI and Python into one single re-usable container image.

Client code and dependencies live on filesystem.

Operating system manages filesystem, resources, etc.

And now we scale...

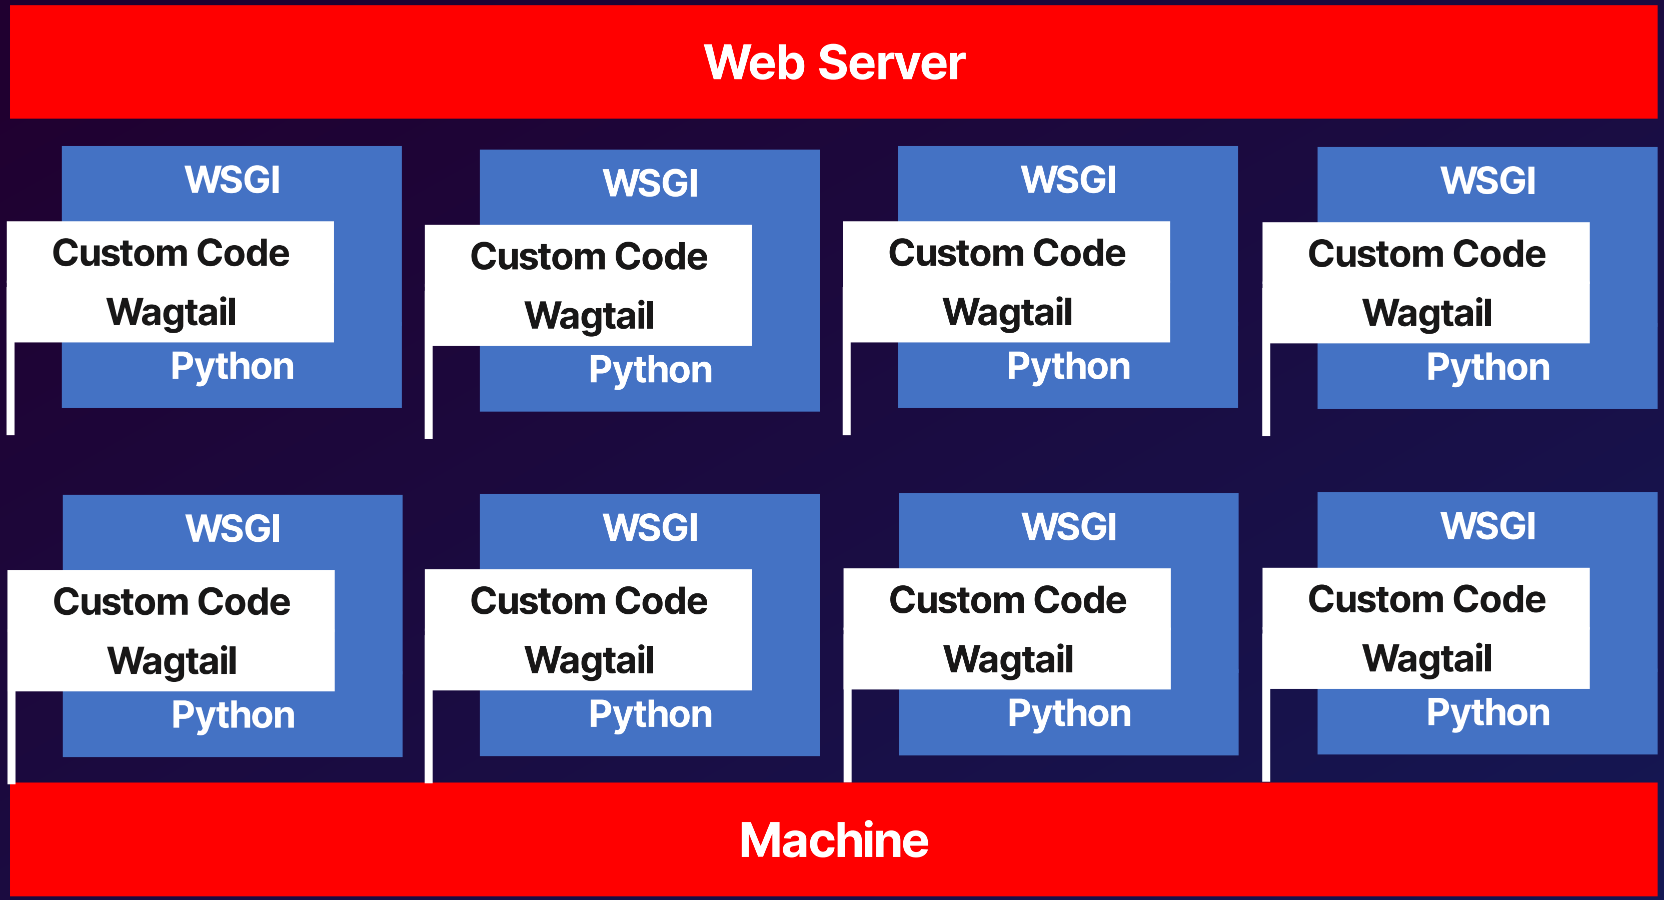

…and we scale…

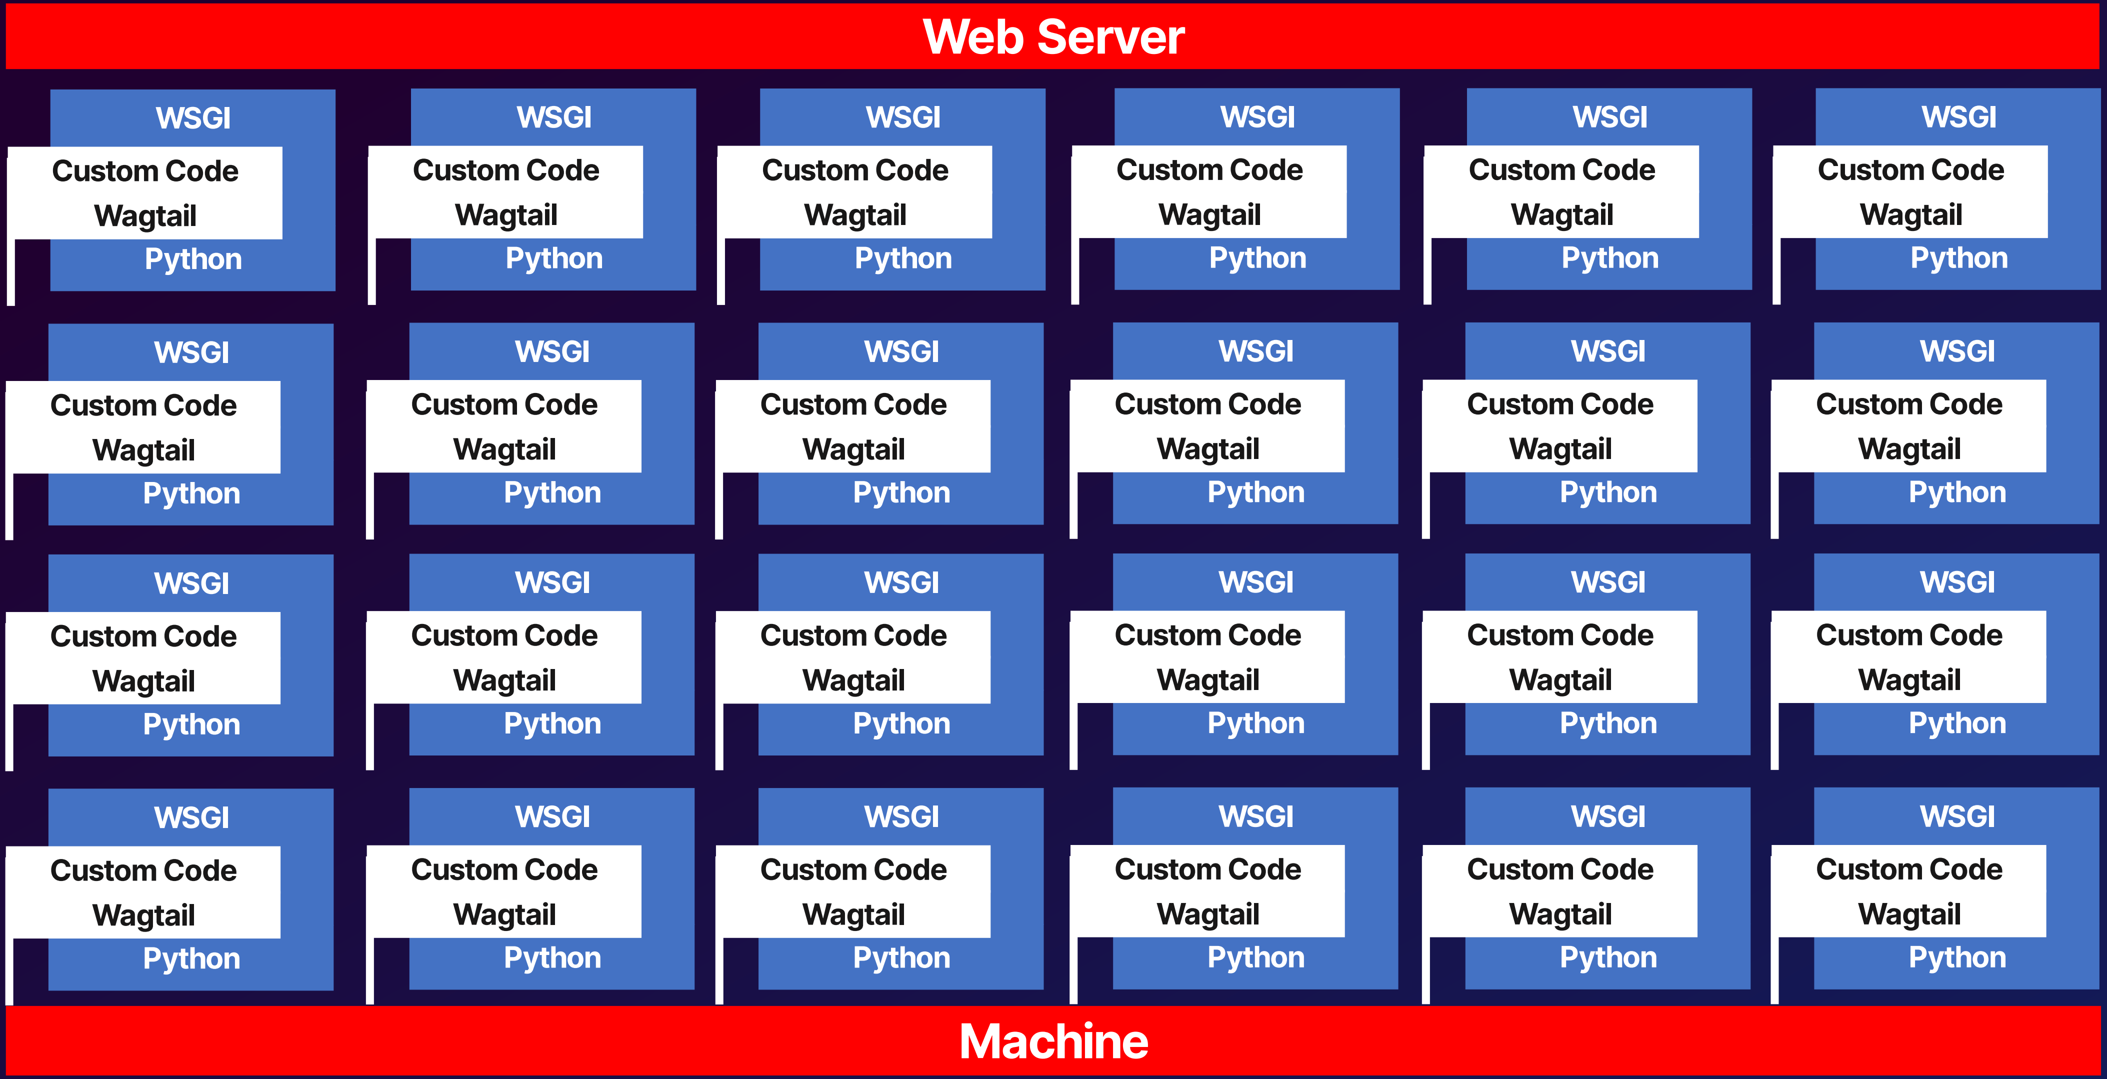

Web Server

WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python
WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python
WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python
WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python	WSGI Custom Code Wagtail Python

Machine

...and scale...

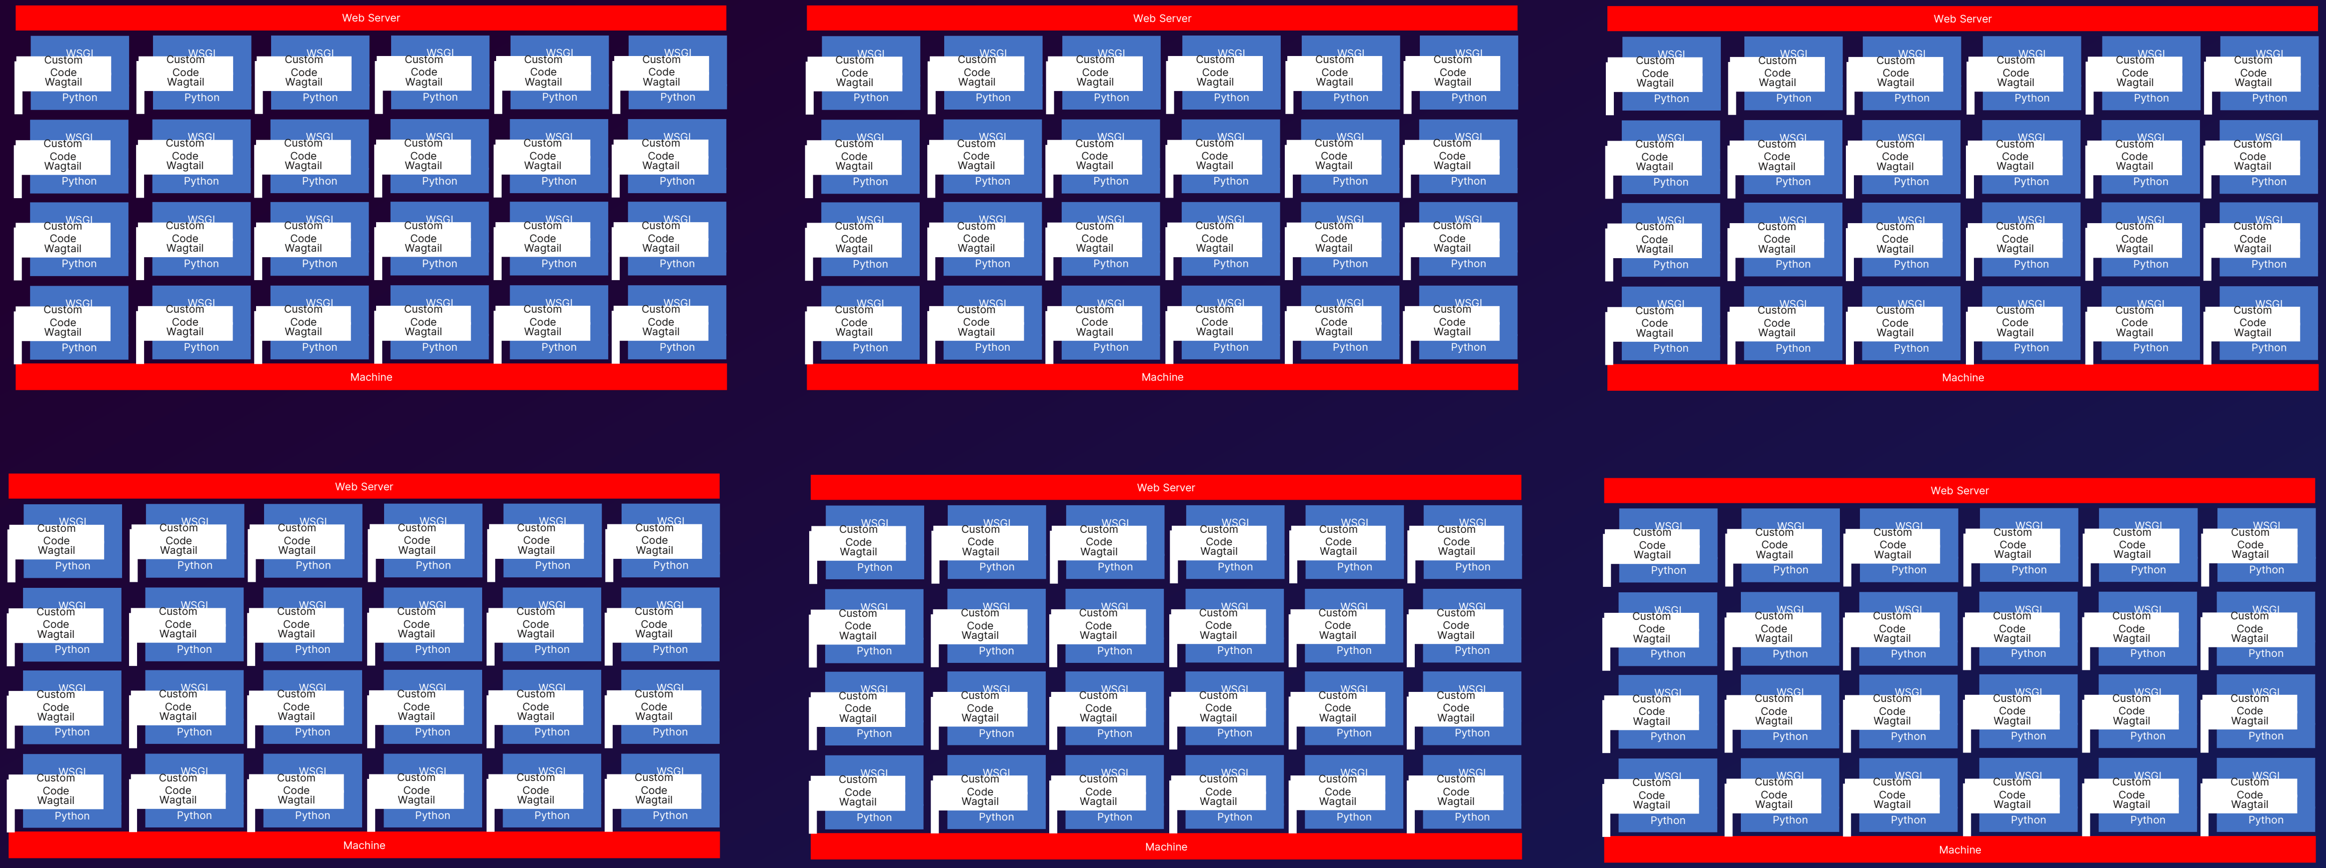

...and scale...

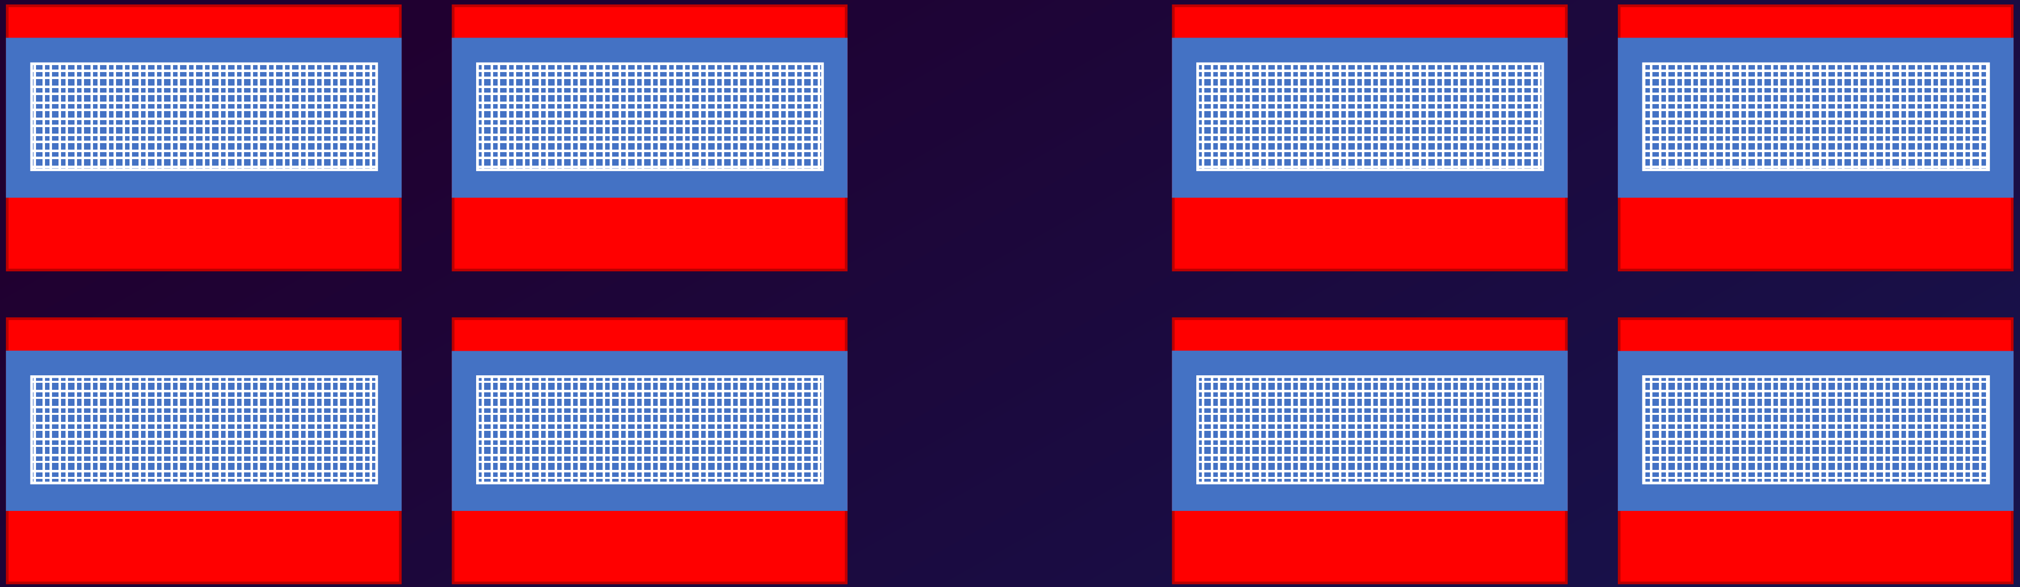

...add geographies and cloud services...

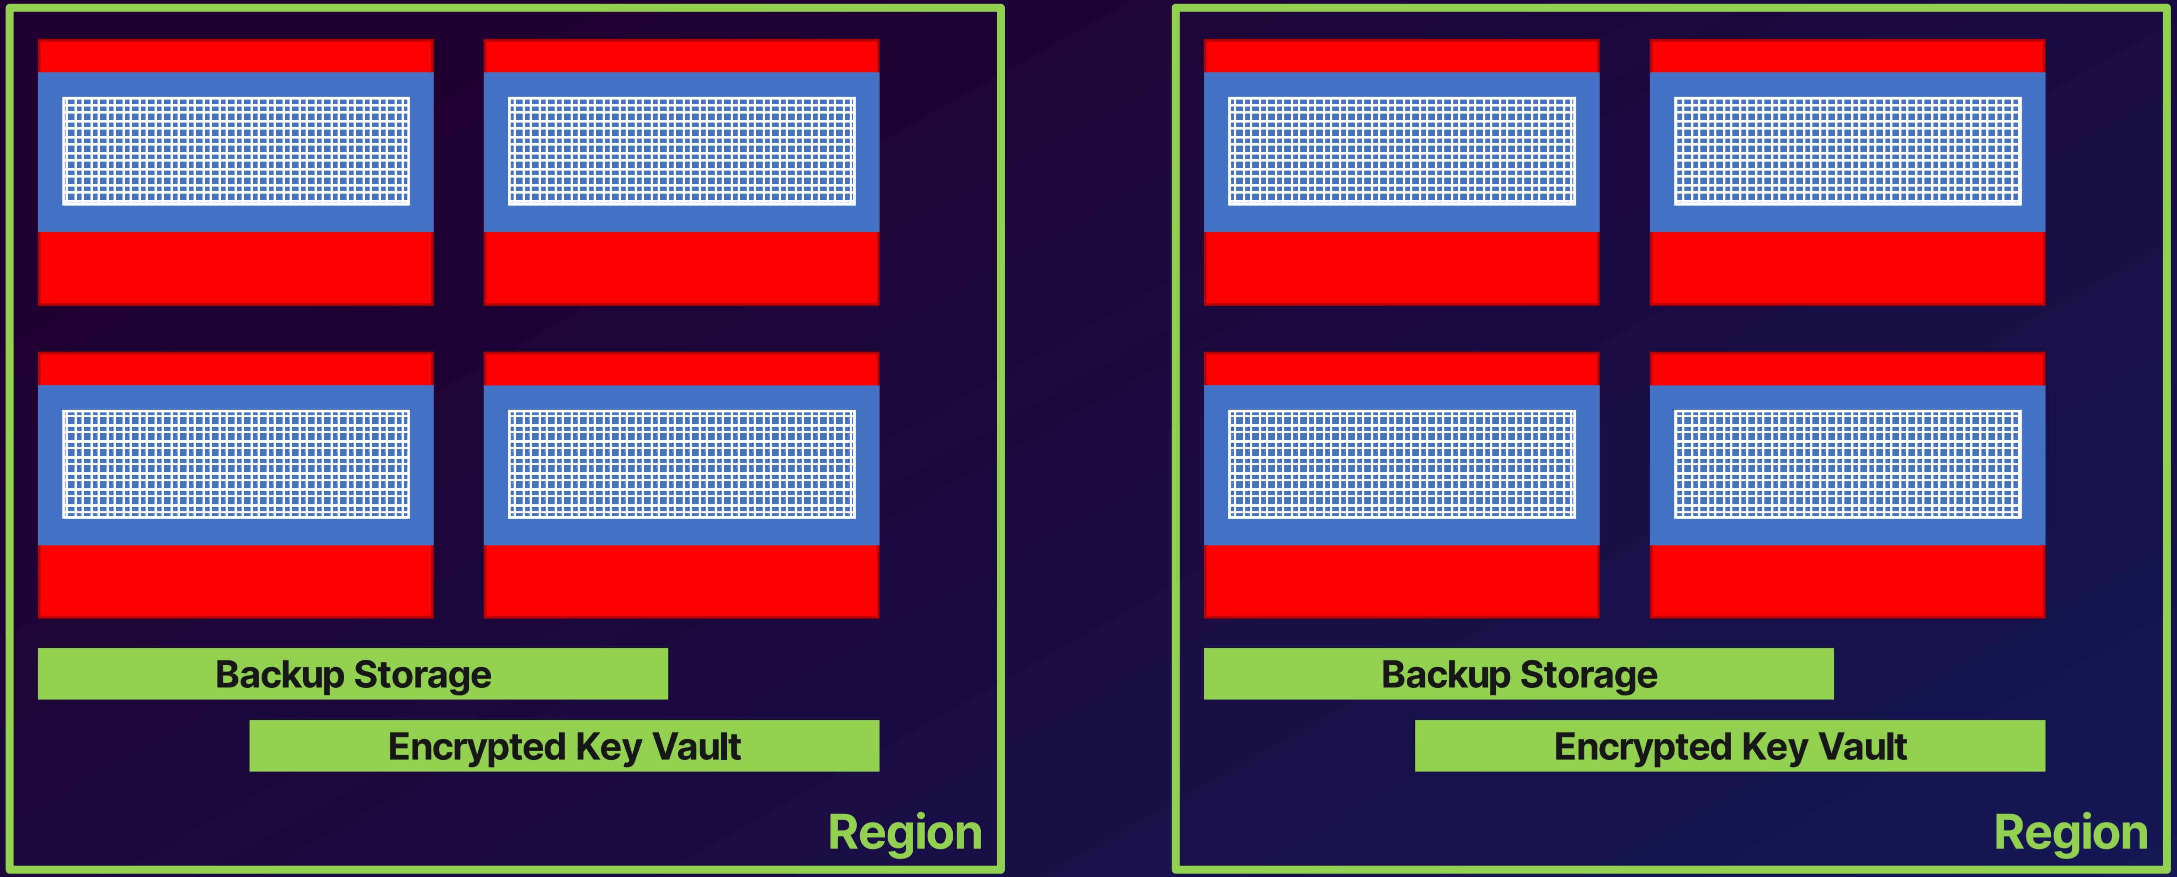

...and we create management infrastructure...

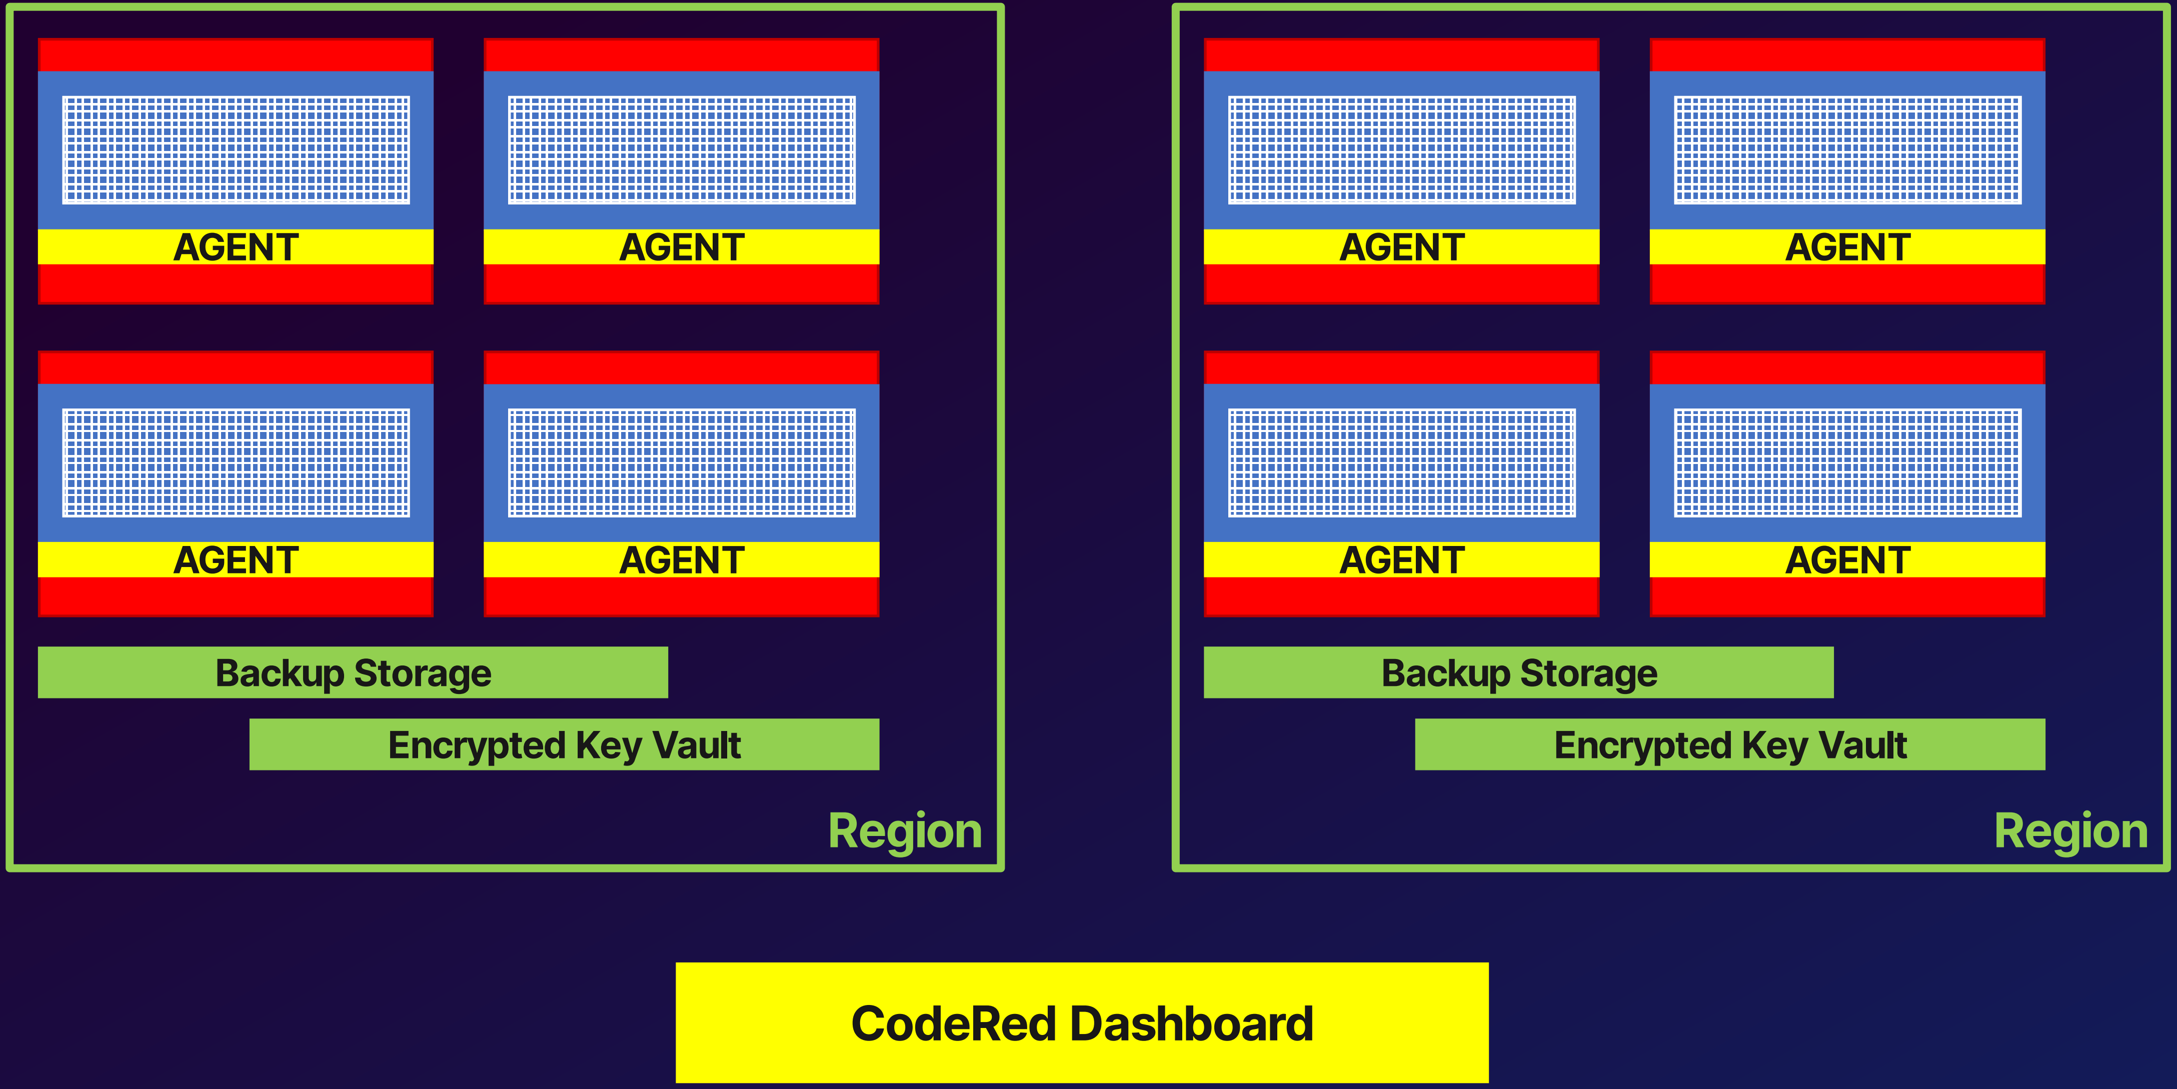

...and we create management infrastructure...

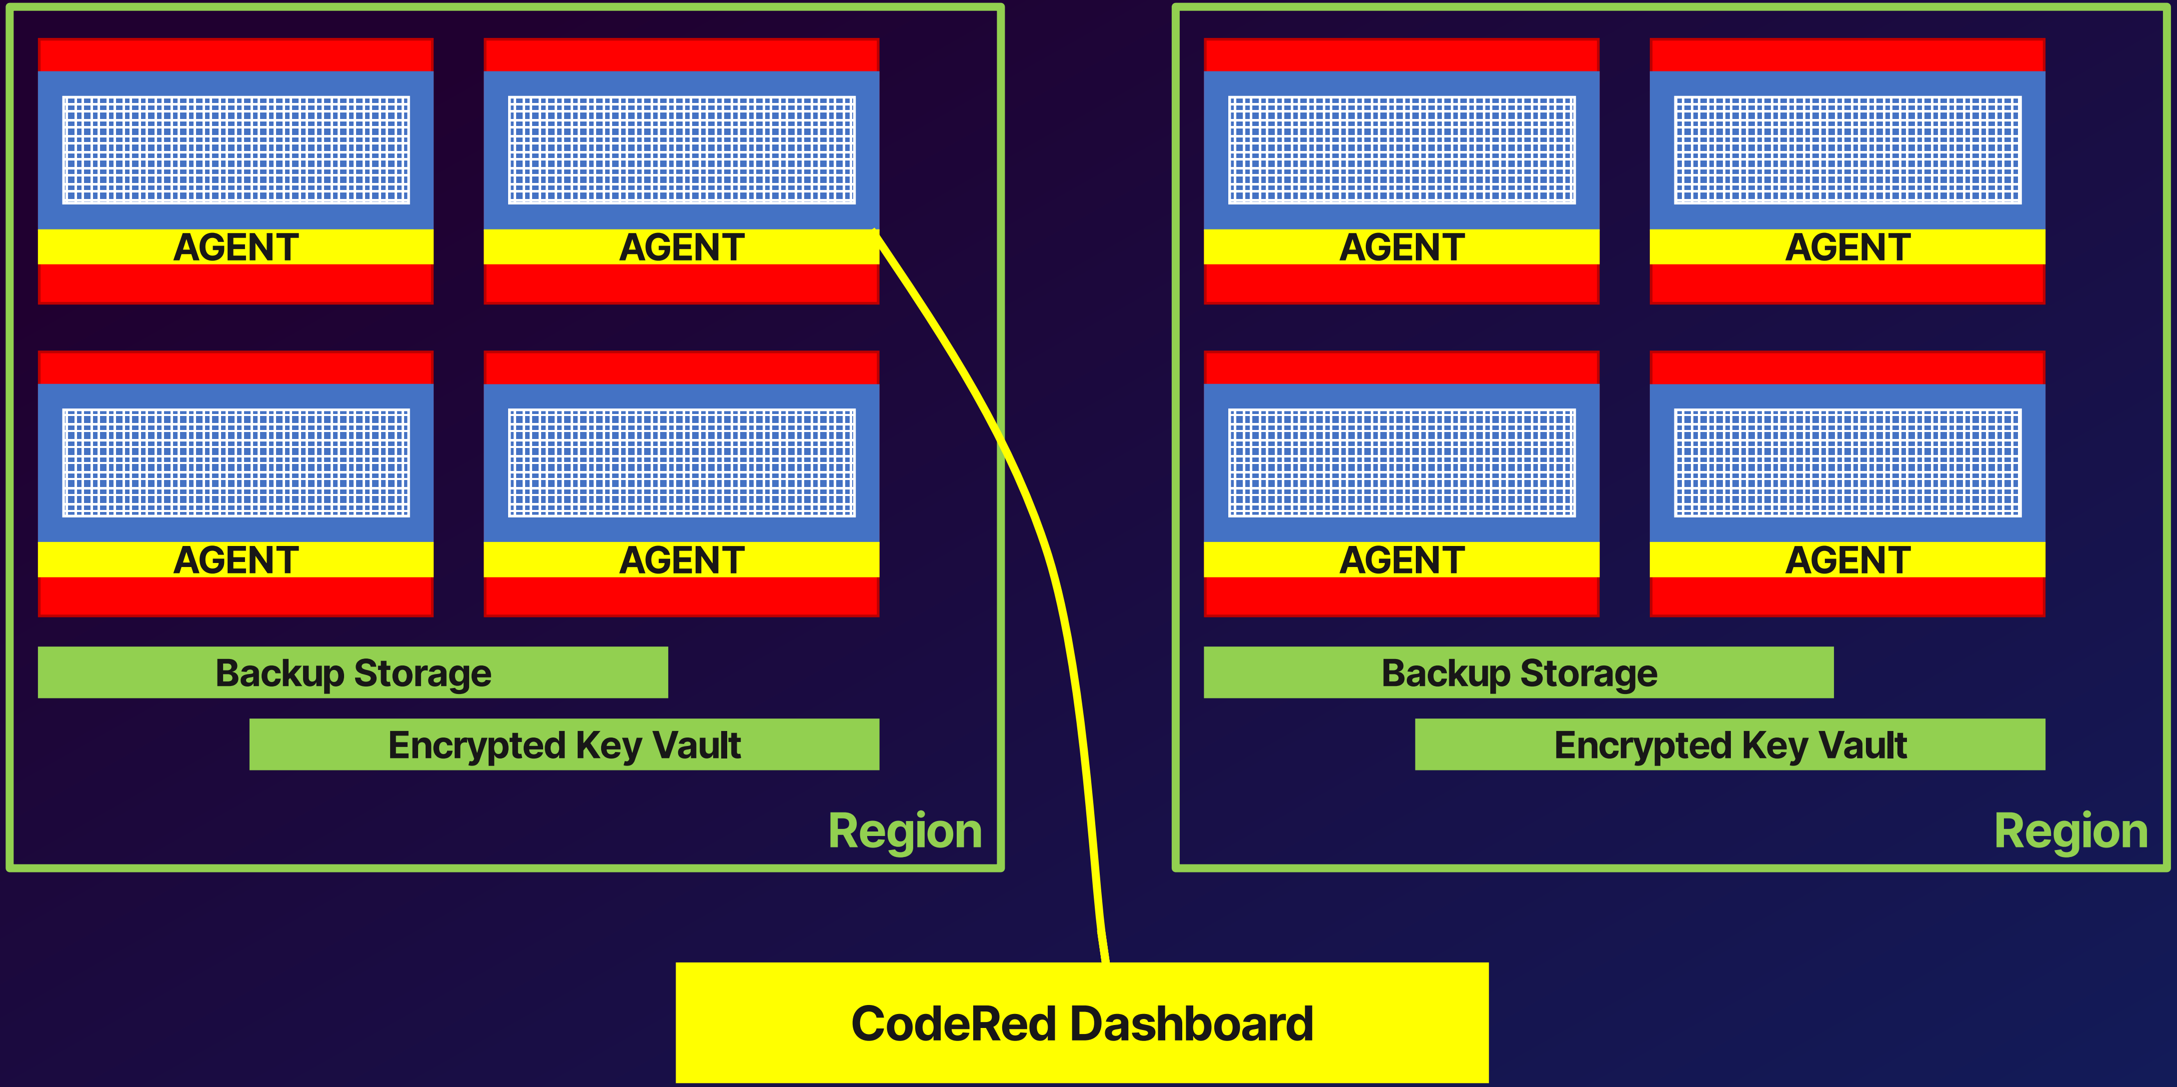

AGENT

AGENT

AGENT

AGENT

Backup Storage

Encrypted Key Vault

Region

AGENT

AGENT

AGENT

AGENT

Backup Storage

Encrypted Key Vault

Region

CodeRed Dashboard

...and we create management infrastructure...

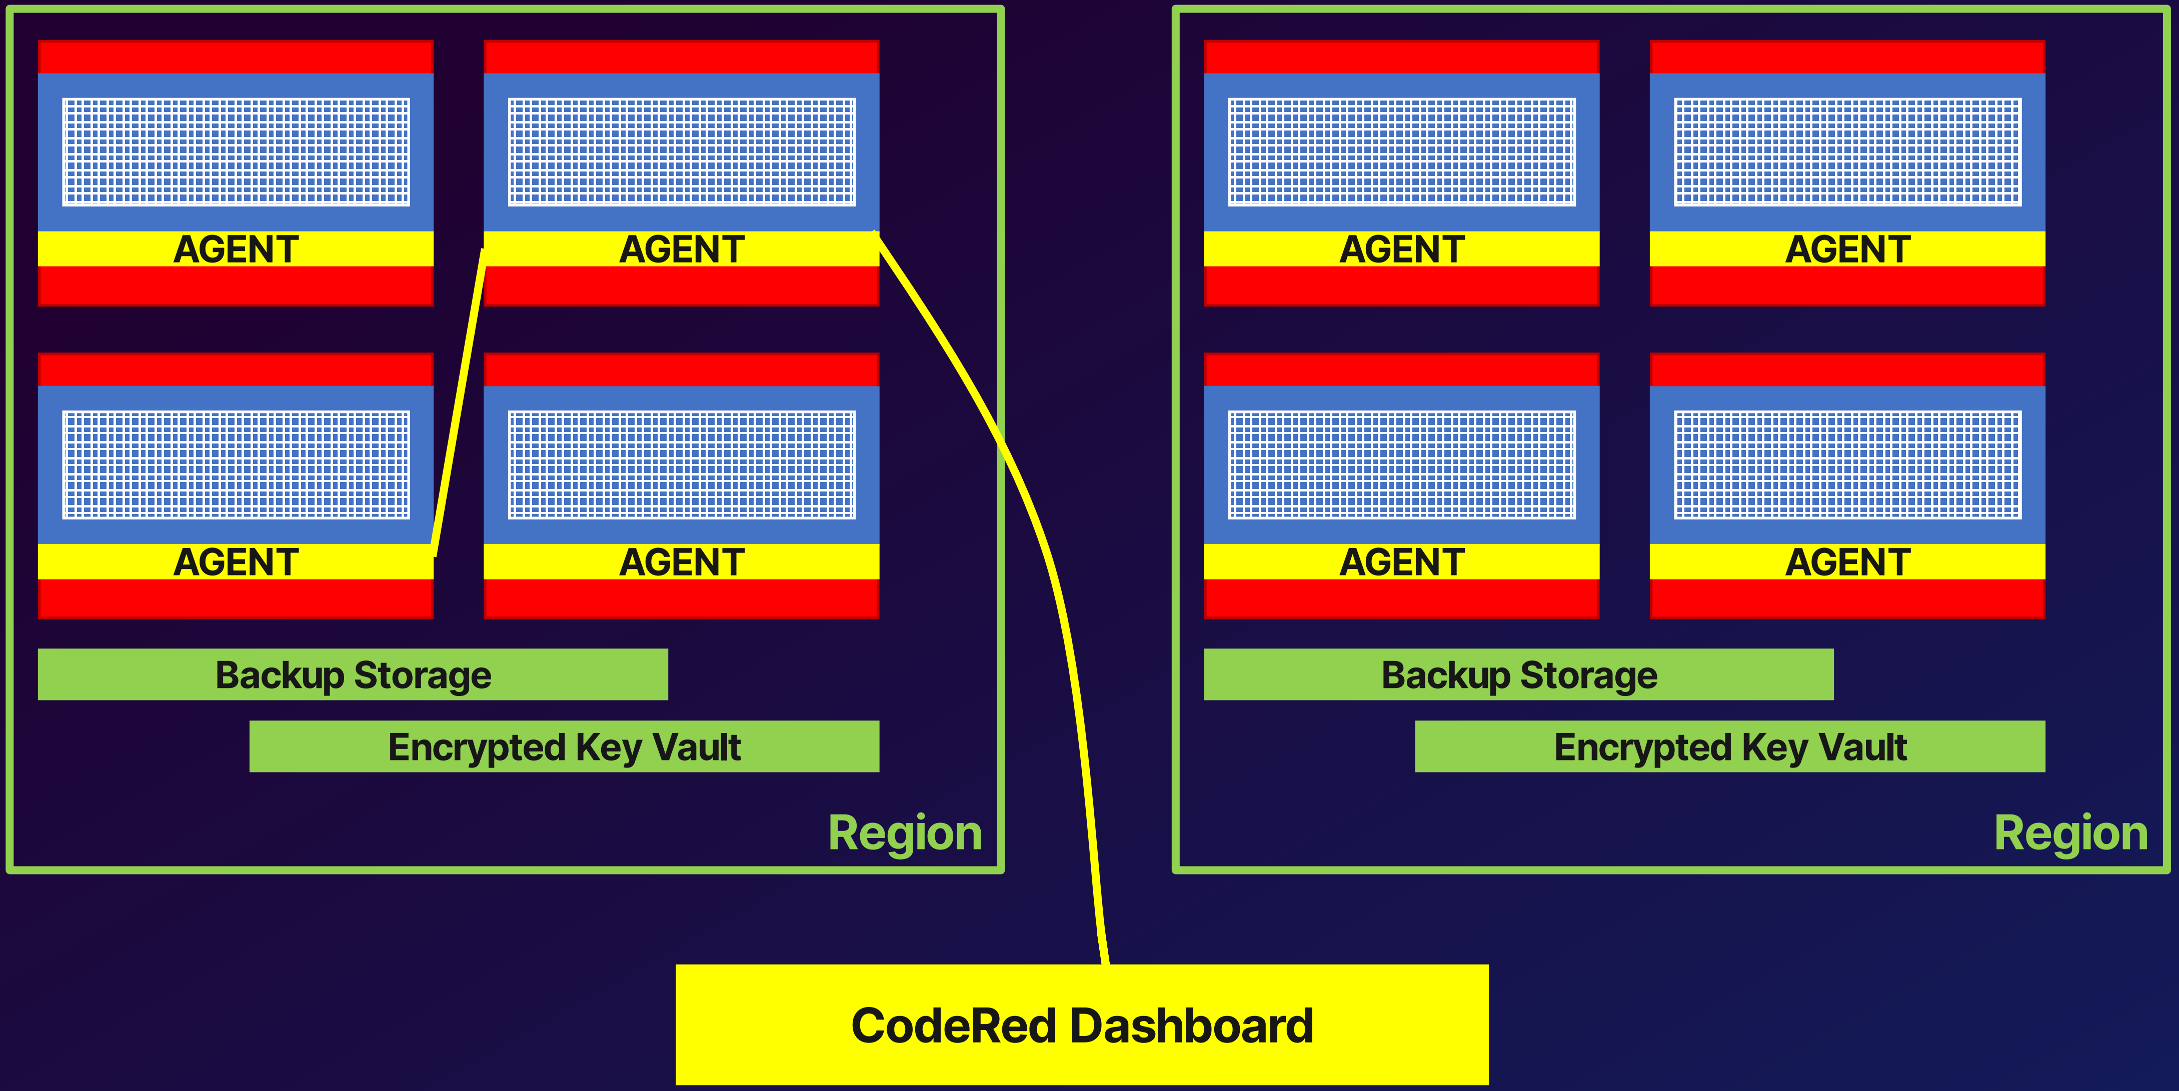

How was this all built?

How was this all built?

How was this all built?

🧙 + 🐍

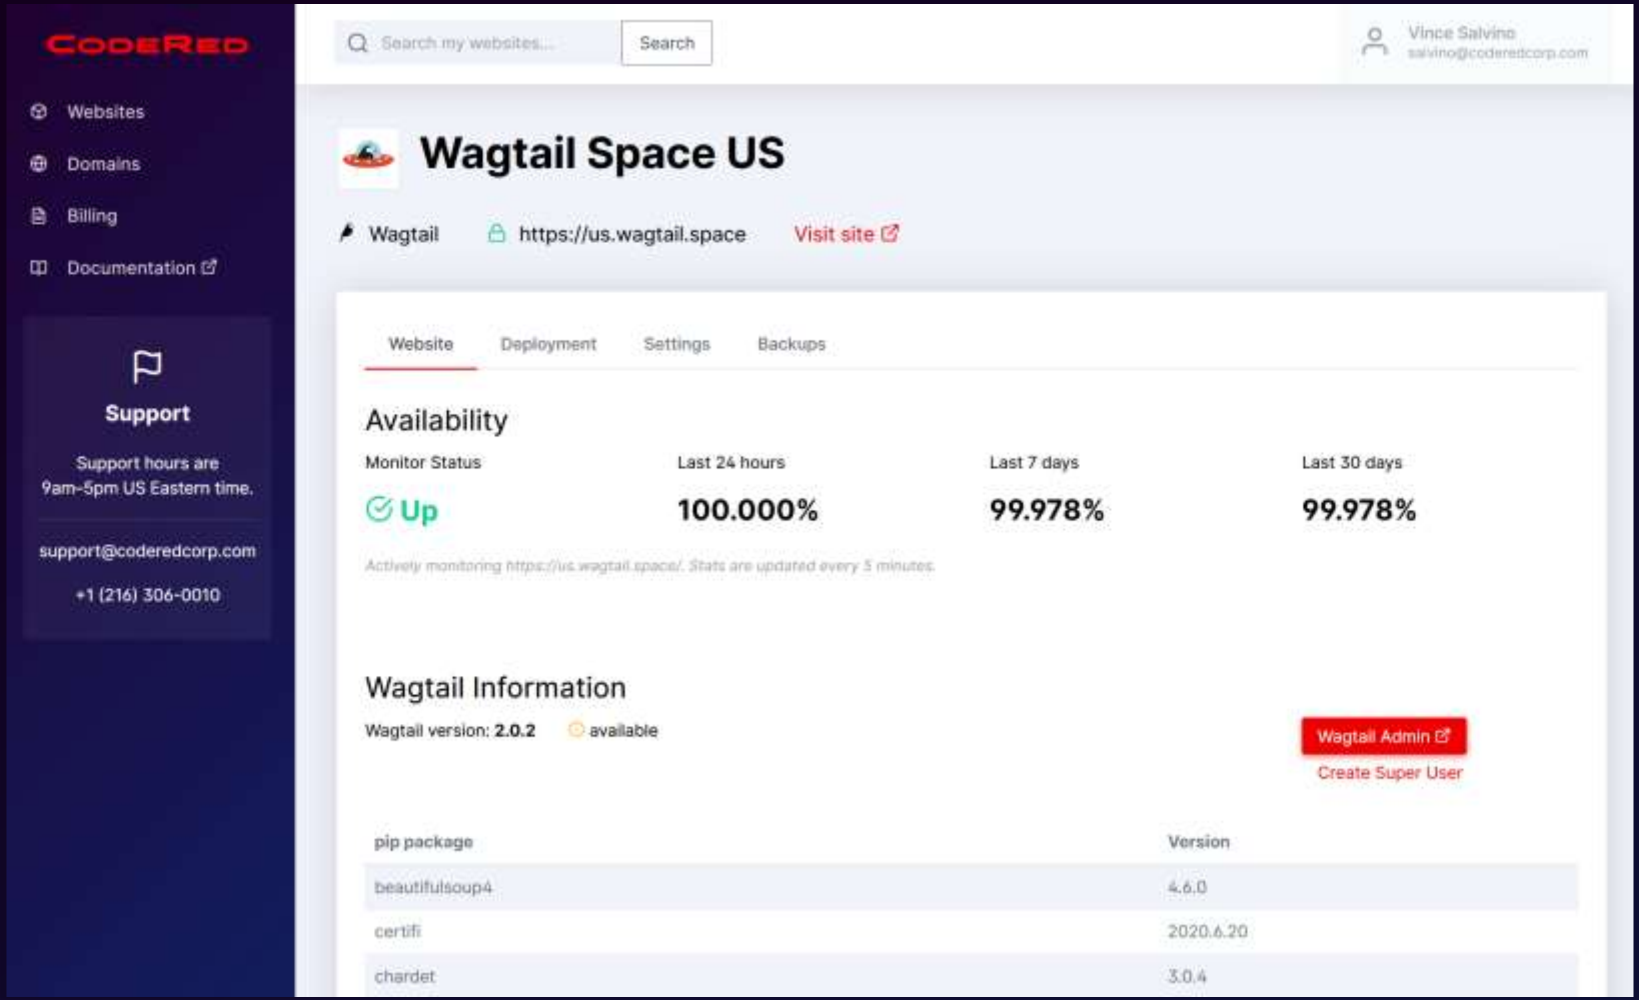

Support

Support hours are 9am-5pm US Eastern time.

support@coderedcorp.com

+1 (216) 306-0010

Vince Salvino
salvino@coderedcorp.com

Wagtail Space US

Wagtail https://us.wagtail.space

Availability

Monitor Status	Last 24 hours	Last 7 days	Last 30 days
✓ Up	100.000%	99.978%	99.978%

Actively monitoring https://us.wagtail.space/. Stats are updated every 5 minutes.

Wagtail Information

Wagtail version: **2.0.2** ○ available

pip package	Version
beautifulsoup4	4.6.0
certifi	2020.6.20
chardet	3.0.4