



NOMAD CUSTOMER STORY

How Nomad powers a Google-backed indoor farming startup to disrupt AgTech

// Infrastructure Enables Innovation



2 Weeks

to deploy Nomad from PoC into production



1 SRE

to manage Nomad



From a Week to 2 Hours

to deploy a new application in production



7x

improvement in resource utilization

About Bowery Farming

Bowery Farming, the modern farming company, was founded in 2015 with the belief that technology and human ingenuity can grow better food for a better future. Bowery's proprietary software system, BoweryOS, uses vision systems, automation technology, and machine learning to monitor plants and all the variables that drive their growth 24/7. Because Bowery controls the entire process from seed to store, farms grow produce year-round, ensuring a safer supply of food that's reliable and consistent. The farms are 100+ times more productive on the same footprint of land than traditional agriculture and grow traceable produce with zero pesticides. Nomad powers their data engineers to optimize the production process from germination to harvest.

The Challenge

In the early days of Bowery, there was no dedicated DevOps resource, and each team had adopted their own methodologies and tools. The variety of workflows across development teams were causing challenges with reliability and velocity at scale. The data science team in particular was struggling with deployment processes and runtime failures. This team, which builds production-facing data products that monitor the agricultural and operational functions in Bowery's network of indoor vertical farms, is central to Bowery's vision as well as its day-to-day operation.

Moreover, Bowery is built around hybrid infrastructure—public cloud and on-prem. Each Bowery farm is equipped with physical bare metal servers, which help facilitate the daily high-compute workloads needed to cultivate the produce. The reliability and velocity at scale challenges were further

exacerbated with a lack of visibility and complexity of deploying applications to two entirely different environments. Some of the issues Bowery had with running services in production were:

- A fragile deployment pipeline caused runtime failures of data services in production
 - Inconsistent, non-standardized application deployment workflows
 - Missing dependencies with improper packaging
- Took up to a week to deploy a new service into production
- Overprovisioning of resources to deploy applications
 - Running a single, dedicated EC2 server for each service

John Spencer joined Bowery as the company's first Site Reliability Engineer (SRE) in early 2019. John immediately recognized that he needed an orchestrator that would be lightweight enough for the on-prem edge deployments at each Bowery farm, unify the data science team and later broader software team with a simple application deployment workflow, provide a monitoring plane into their hybrid infrastructure, and still be easy enough to be operated by one person – himself.

“ Each team had their own method of deploying services and containers. Some did it on Beanstalk or other AWS managed services, some just cloned their repo onto an EC2 instance and ran the service right on the VM. For on-prem, we had Docker running on a server. We'd have to SSH into the server and run the Docker command by hand, which created lots of opportunity for operator error.”

Why Nomad

Bowery evaluated Kubernetes and HashiCorp Nomad. Nomad won on the following criteria:

Operational Simplicity

As the only SRE at Bowery juggling multiple responsibilities, John needed an orchestrator that he could easily understand and operate by himself. His strong positive experience using other HashiCorp products, such as Terraform and Packer, led to a strong consideration of Nomad. Nomad's simple architecture and operational ease-of-use enabled John to set up a Nomad cluster in one day and move to production in just two short weeks.

“From the research I’ve done and the stories I had heard about operating Kubernetes, I was worried it was too much for a single SRE to use efficiently and effectively. I would rather start with a simpler service that solved our immediate needs (running batch jobs) while laying the groundwork for a more advanced, feature rich system down the road. Nomad was perfectly suited for this task, especially because it’s easy to extend the feature set by adding on other HashiCorp products, like Consul and Vault, which all integrate seamlessly. What’s nice about Nomad is that even though it’s powerful, operating Nomad is not an insurmountable task and it did not take me too long to get it up and running. Nomad makes it really easy to get started with a dev mode and bootstrap quickly to get one’s feet wet and then build out a production configuration from there.”

// Within one day, I had a Nomad cluster running in AWS with a pretty barebones configuration. The fact that I was able to do that in a single day was pretty awesome. From there, I built out the infrastructure automation with Packer and Terraform, creating immutable AMIs, AutoScaling groups, load balancers, IAM profiles, and security groups, and within 2-3 weeks I was ready to schedule the first job in production. To be able to accomplish all of that with an SRE team of one was a testament to the easy operability of Nomad.”

Lightweight Edge Deployment

Bowery has a hub-and-spoke distribution model. Internal workloads run centrally on the public cloud while farms run critical services on bare metal at the edge. Nomad’s lightweight single binary (~30MB) and agent made it possible for John to deploy a Nomad client at each edge location and build one single cluster across their distributed infrastructures. Unlike other orchestrators whose complexity requires separate clusters in each edge location with more than 10X capacity needed, Nomad won on enabling a simple topology with less management and resource overhead.

“We build farms, which are our edge locations. Because of this, we cannot run services exclusively in the cloud. That’s where a tool like Nomad came in perfectly. We have a Nomad cluster in AWS with a bunch of client nodes. Every one of our farms is a separate Nomad datacenter, and depending on how many servers or VMs we have at each farm, we run one or more Nomad clients. It’s a big win to be able to schedule jobs through the same Nomad API, regardless of the datacenter.”

Efficient Hybrid Deployment

In addition to Bowery's distributed infrastructure, the company has a hybrid environment, with central services running on public cloud and edge deployments on-prem. The installation and configuration process for Nomad is the same across Bowery's hybrid environments. Nomad treats both cloud instances and on-prem bare metal servers identically as fingerprinted pools of resources. John can enable a unified workflow to deploy services to each environment without friction.

The Result

Today, Nomad is deployed in production at Bowery, enabling data engineers to optimize the production process from germination to harvest 365 days a year. The initial success with Nomad led to broader adoption by migrating software team services into Nomad as well.

Improved Productivity

Nomad unified Bowery's teams with a single application deployment workflow, which has enabled developers to self-serve and deploy services to the cloud or a specific farm at the edge. Consequently, the time to deploy applications has been significantly improved. Some of the improvements Bowery's seen include:

- Self-service deployment workflow for developers without need for SRE intervention
- Reduced time to deploy an application in production from one week to two hours
- Less than two hours to onboard a new engineer and start deploying applications

" I helped the team lift and shift a couple of jobs into Nomad. And then I said, okay, here's the [Nomad] cluster, here are the tokens, here's the API endpoint, and here are some docs. Deploy away! And from there, the team was able to deploy dozens of services, using an easy, unified workflow, with very little intervention from me. This was a huge productivity win compared to the previous deployment workflow, which was much more time intensive and error prone."

Single Monitoring and Control Plane with Fault Tolerance

With Nomad, John was able to achieve a single cluster topology for Bowery's hybrid infrastructure. This enables John to have a single monitoring and control plane for all of Bowery's applications with the ability to deploy services freely to the cloud or to a specific farm at the edge at any time. Nomad's visibility also extended into the application level, providing John a picture of resource usage of each service.

"We were experiencing memory issues where an EC2 instance would get stuck because the service running on it used up all the memory. Nomad and Docker were able to help us better profile the memory usage of each of our applications. With Nomad, there's a configuration that lets us just reschedule the same job if it runs out of memory. So while we occasionally may lose an application to a memory issue now, it takes no intervention now with Nomad to get that application up and running again."

Nomad's built-in fault tolerance has also brought resilience to Bowery at the application and node level. Before Nomad, deployments would occasionally fail and require a manual rebuild of infrastructure and redeploy of the service. With Nomad, Bowery's infrastructure is now self-healing, which has freed up John's capacity to focus on more pressing responsibilities and less on crisis management and firefighting duties as the single SRE.

// We've seen EC2 instances fall over and disappear from the console, which results in losing a Nomad client or server node, but I've never had to intervene. I've never been woken up in the middle of the night to patch our Nomad cluster. AWS Autoscaling and Nomad bring up a new instance and it gets brought into the cluster within a matter of minutes, or even seconds. That's a huge win."

Cost Savings

Nomad enables Bowery to achieve high levels of resource utilization and efficiency, resulting in significant cost savings. Prior to Nomad, developers would overprovision resources to deploy their applications by running a single dedicated EC2 instance for each service. Nomad's bin-packing enabled Bowery to densely schedule multiple applications on a single machine. The company went from one application running on an EC2 instance to 30-40 jobs across 5 Nomad clients, achieving 7x in resource utilization and cost savings.

"We certainly have fewer resources, even though we've deployed many more services. Before, each resource was living on its own EC2 instance even if it was just a single Docker container. Now those have all been shrunk down to a couple of Nomad clients that are running somewhere around 30 to 40 services, definitely a big reduction there."

The Conclusion

With Nomad, John helped Bowery improve deployment reliability, reduce the time-to-deployment, and achieve significant capacity savings. Moreover, as engineers have rapidly adopted Nomad, John was able to free up his time to support more critical initiatives. Bowery recently finished its second farm, which is 3 times bigger than its first farm. Nomad will remain as the orchestrator of choice to help Bowery Farms scale in the future.

// We're confident in our choice of Nomad as our application scheduler, and have since integrated Consul into our stack as well. Bowery will continue to build more and more farms which will each have a number of Nomad clients, ready to accept and run any service we dream up. Nomad has proven itself to be highly scalable, and we're excited to scale our business alongside it."

Technology Stack

- Infrastructure: AWS + on premise
- Workload type: EC2 / Bare Metal
- Container Runtime: Docker
- Orchestrator: Nomad
- CI/CD: CircleCI
- Version Control: Git / GitHub
- Networking: Consul / Meraki / Unifi
- Provisioning: Terraform, Packer, Ansible

Bowery Farming Partner



John Spencer was born and raised in Oakland, CA. After moving to NYC to pursue a music degree in jazz drum set, John took an intro to computer science course and fell in love. He now spends his time automating the software systems which operate the world's most advanced indoor farms, while still finding time to play drums in punk rock bands.

John Spencer, Senior Site Reliability Engineer

