HIVECELL EDGE

Reference Architecture

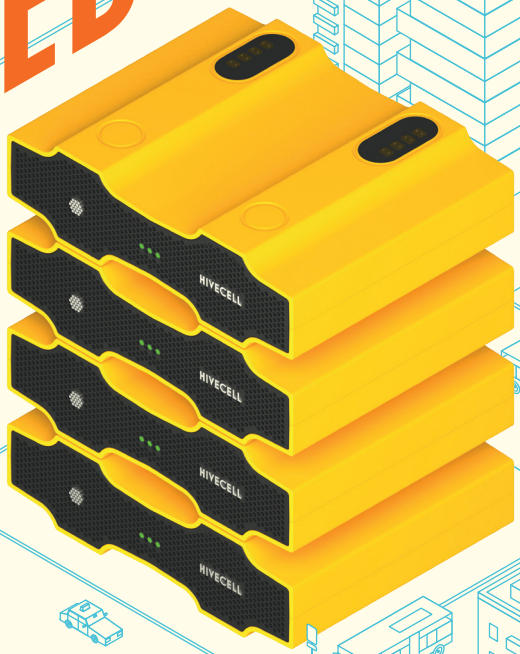

HIVECELL®

HIVECELL EDGE

Reference Architecture

HIVECELL®

INTRODUCTION

Hivecell is a new type of server designed for the Edge.In this paper we will describe how Hivecell addresses the new requirements of edge computing. We will present a reference architecture for deploying a solution of software and hardware at the edge with Hivecells.

SITUATION

It was commonly assumed in the industry that we would process all the data from the internet of things (IOT) in the cloud. Perhaps that occurred because the market began discussing the concepts of cloud and IOT at about the same time. However, as the market began to implement IOT, it became obvious that we had greatly underestimated how much data we were dealing with.

In 2018, the estimated amount of data in the cloud was 12 zettabytes (ZB). At the same time, the estimated amount of data at the edge was 26 ZB. By 2021, the cloud data is expected to nearly double to 20 ZB, but the edge will be three times larger at more than 60 ZB.

Industry has realized that it is impossible to move all the edge data to the cloud for processing. Instead, we must process raw data as close to the source as possible and push only business relevant data to the cloud.

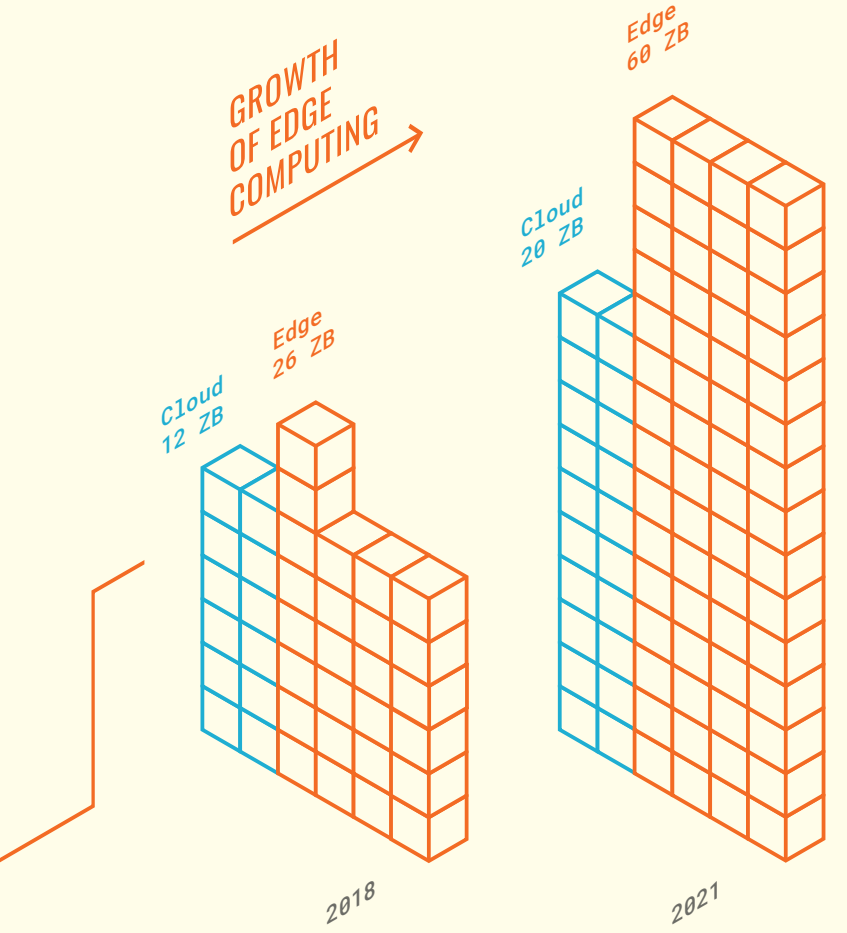
GROWTH OF EDGE COMPUTING
Edge 60 ZB
Cloud 20 ZB
Edge 26 ZB
Cloud 12 ZB
2018
2021

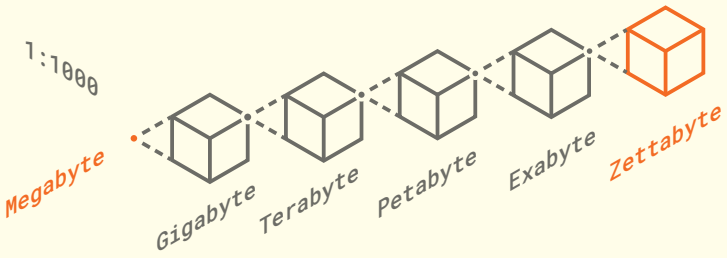
1:1000
Megabyte
Gigabyte
Terabyte
Petabyte
Exabyte
Zettabyte

REASONS FOR EDGE

The sheer size of edge data is not the only reason for edge computing. There are several other reasons for moving compute from the cloud to the edge, most notably security, compliance, bandwidth, latency, reliability and cost.

Security: *some data is considered too sensitive to move across the internet to the cloud, even if encrypted and in a virtual private cloud (VPC)*

Compliance: *there are regulations in many countries that dictate where and how data can be stored which prevents using the cloud*

Bandwidth: *some data is so large that it is not feasible to move it from its source to the cloud*

Latency: *it some situations the network latency of moving data to the cloud and back again is impractical or even dangerous*

Reliability: *system safety or even business continuity may demand local processing of data to avoid possible network outages*

Cost: *moving very large amounts of data across the internet and storing in the cloud can be expensive*

CHALLENGES

For our discussions here, edge computing means deploying outside of the data center or data closet the compute power normally associated with servers. Edge computing deployments usually involve hundreds or even thousands of physical locations. It may be a fleet of a thousand vehicles, or a chain of a thousand retail locations or a windmill farm of a thousand turbines. It may be hundreds of plant floor locations in dozens of plants. The scale of physical distribution brings certain challenges. No one has the army of technicians required to install, service, maintain, upgrade or troubleshoot a thousand locations. The labor cost would quickly outweigh any business advantage sought in the deployment. As such, deploying edge compute must require no special technical know-how. Anyone should be able to install or replace an edge computer.

Deploying edge computers should be able to go anywhere. You cannot rely on special power, networking or cooling to be installed. Edge computers must be small, low power, quiet and reliable so they can go anywhere.

Edge computers must be able to run the new types of distributed software such as Kubernetes, Hadoop, Kafka and Tensorflow. *Distributed software is designed to grow linearly* as the workload increases. The hardware should be able to scale linearly as well. It should require no backplane or additional equipment.

HIVECELL CAPABILITIES

Each Hivecell has a 6 core 64 bit ARM CPU, a 256 core GPU, 16 GB RAM and a 500 GB SSD. It has the compute power for not only running servers and services but also running machine learning models.

Hivecells stack together like toy blocks. You plug in power and ethernet in the back of a Hivecell. Then you stack other Hivecells on top. There is no need for additional wires.

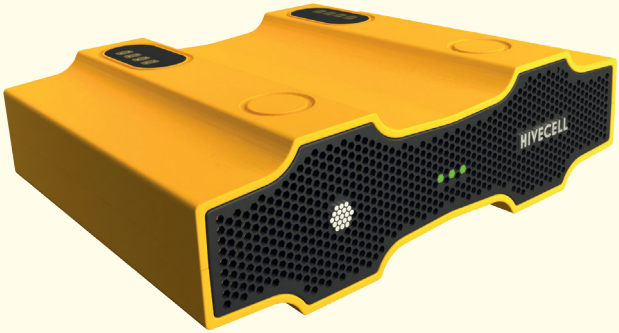

64-bit ARMv8 Processing

6 CPU cores, 2.4 GHz

256 GPU CUDA cores

8GB RAM LPDDR4

500 GB SSD

1G Ethernet

Wifi IEEE 802.11a/b/g/n/ac dual-band 2x2 MIMO

Size 220x175x65 mm

Weight 1.36 kg (3.0 lbs)

Power 15 W (Max 25 W)

The Baranovsky connectors on the top and bottom of the Hivecells pass power and ethernet. You simply add more Hivecells to the stack as you need more compute power.

Hivecells have a built in provision system. That means you can install and update software remotely, without the need to have a technician present. If you remove a Hivecell from the stack, it immediately switches to battery power and wireless communication with the stack so distributed processing can continue uninterrupted. Once you place the Hivecell back on the stack, it immediately switches back to wired power and ethernet.

Hivecells have their own built power supply and routing. There is no need for an uninterrupted power supply (UPS) or router for deployment. They are small and operate at 15 watts, so they do not require special power or cooling to be installed. They can be stacked anywhere.

HIVECELL DEPLOYMENT

PHYSICAL HIVE

A hive is a logical collection of Hivecells that serve a common purpose. A hive may be a cluster for frameworks such as Kubernetes or Hadoop. Alternatively, a hive may be the same application running on two or more Hivecells in order to provide high availability.

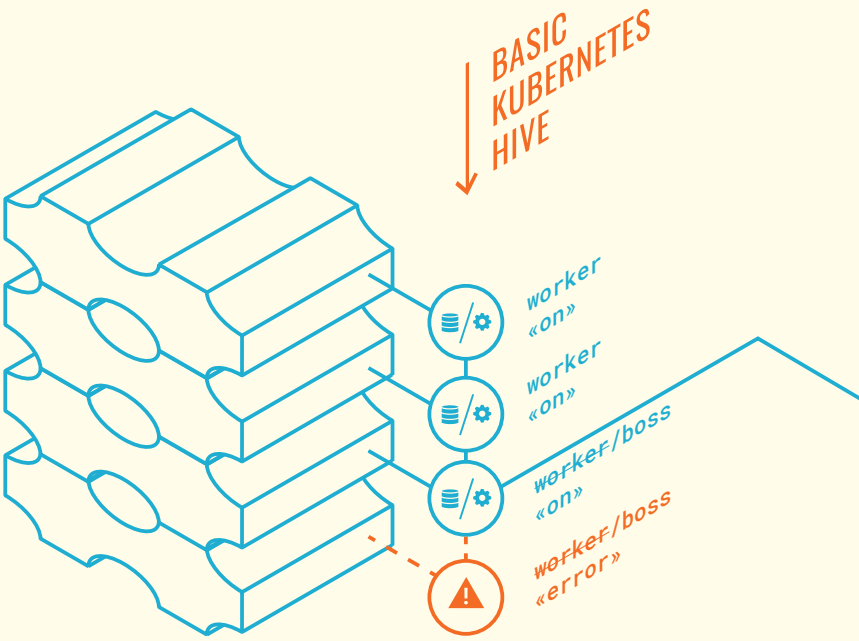

Physically, the basic Kubernetes hive consists of four Hivecells: one boss and three workers. If the workload of the hive increases, then more Hivecell workers can be stacked on top.

The boss is (usually) the bottom Hivecell which has the power and ethernet connection. It is an elected role, so should the boss Hivecell fail, another worker can assume the role of boss.

Each Hivecell can provide not only compute but also storage. A standard hive has encrypted distributed file storage so that if you lose a Hivecell you do not lose data.

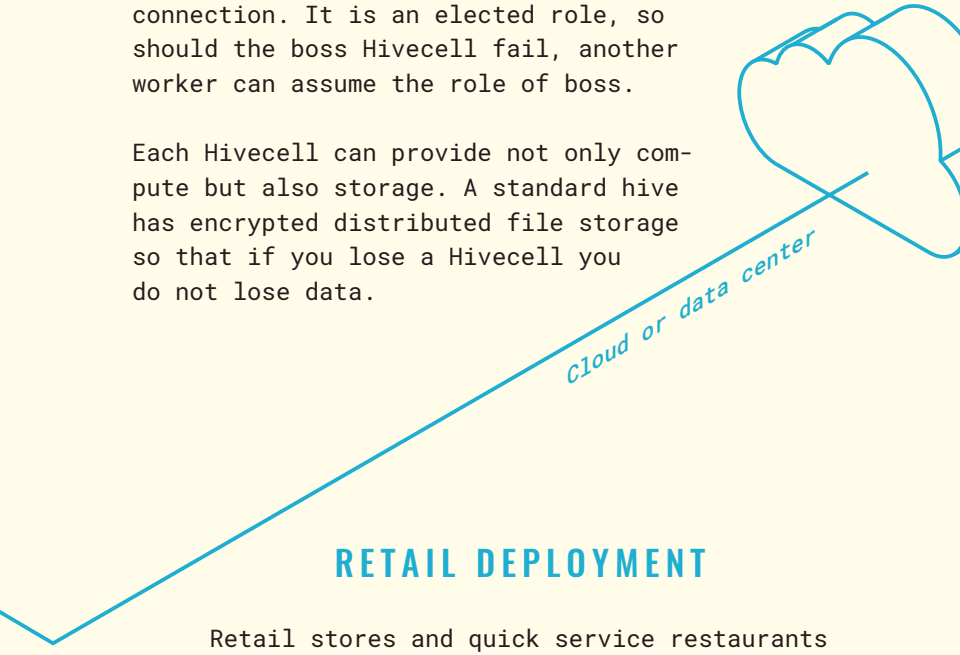

RETAIL DEPLOYMENT

Retail stores and quick service restaurants are now filled with smart equipment that generate and publish all their performance metrics data to the network. The raw data may tell you what is wrong if something does go wrong, but the real value of this data comes from applying predictive analytics and machine learning. Hivecell provides compute power at the edge, at the source of the data. This capability enables

companies to process the raw data on site and ship only the business relevant data to the data center or the cloud. This capability is particularly relevant when applying machine learning models to video.

Why ship all the hay to the cloud when all you really need are the needles?

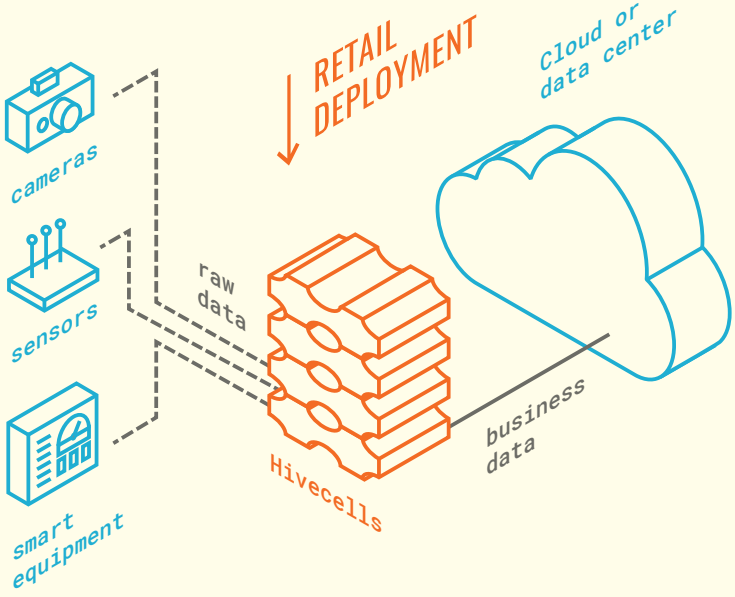

INDUSTRIAL DEPLOYMENT

In an industrial deployment, an IOT gateway usually sits between the things and a hive. The gateway simply normalizes the myriad protocols used by things (CANBUS, MODBUS, Zigbee, etc.) to the standard TCP/IP internet protocols over

Ethernet. From there, applications running on Hivecell can process the data and forward the results “northbound.”

There are usually several hives for raw data processing within a given industrial plant, factory or refinery. The physical distance between sources or the different types of equipment will determine the number and location of the hives. Each zone of things will have its own hive for raw data processing. There will also be other hives for shared workloads. For instance, there will be a manager hive for updating hive software as well as monitoring hives.

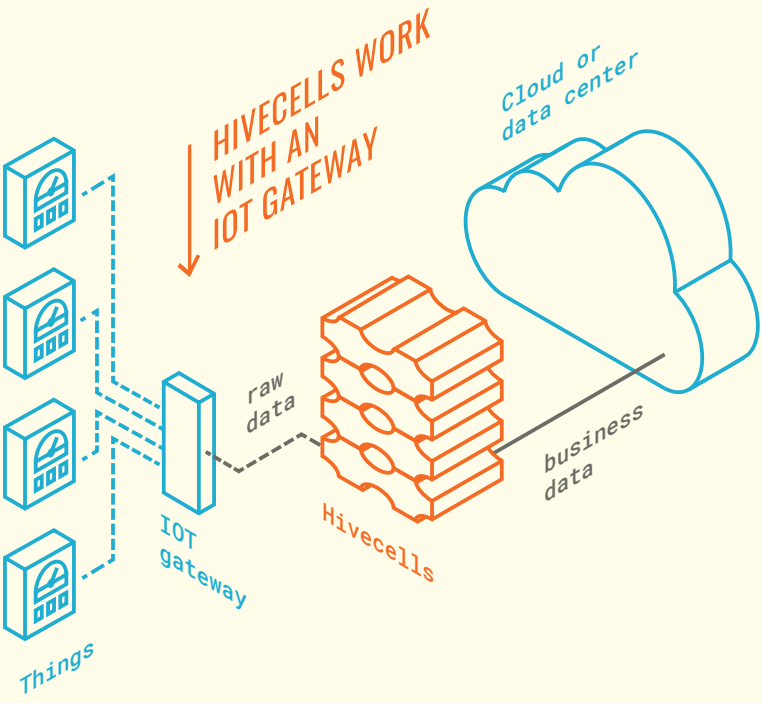

There may be a messenger hive for running Kafka or other shared messaging infrastructure between hives and the datacenter. There may also be one or more analysis hives that enable advanced analytics, development or testing across zones.

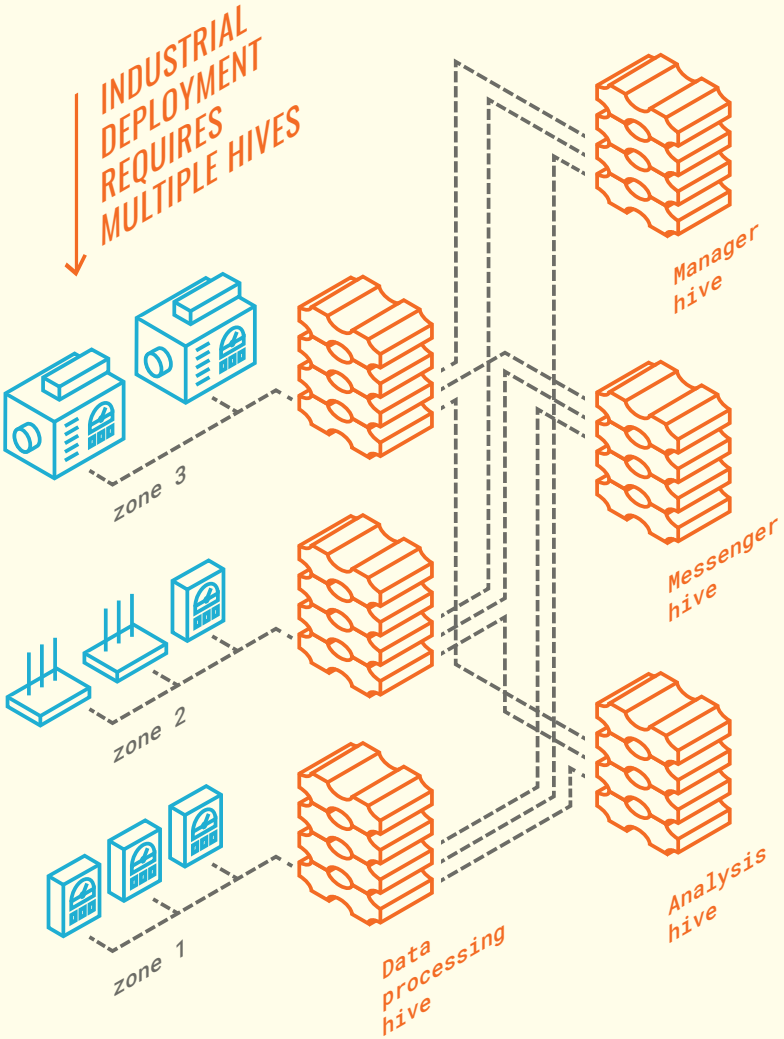

DEVOPS

Deploying and upgrading software is a key challenge of edge computing. Hivecell enables three different modes for software deployment:

Mode 1: Cloud based

Mode 2: Enterprise

Mode 3: Disconnected

We will discuss each of these modes.

MODE 1: CLOUD BASED

The simplest mode for deploying software to Hivecells is a cloud centric approach.

1. Hivecell periodically identifies itself with digital signature to cloud Hive Manager
2. Hivecell receives notification from Hive Manager of a major upgrade
3. Hivecell reboots from worker drive to drone drive
4. Hivecell downloads new image
5. Hivecell confirms new image is installed correctly
6. Hivecell reboots to worker drive

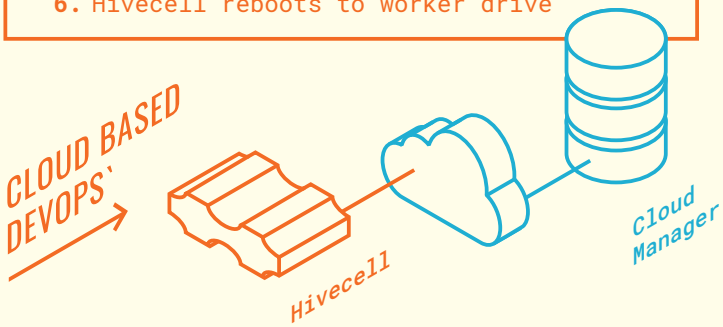

MODE 2: ENTERPRISE

A cloud-centric approach is not acceptable in most large enterprises. An alternative approach is to move the Hive Manager to behind the corporate firewall.

1. Hivecell periodically identifies itself with digital signature to well known Hive Manager behind corporate firewall
2. Remaining steps are the same as mode 1
3. Hive Manager can pull new images from cloud if necessary

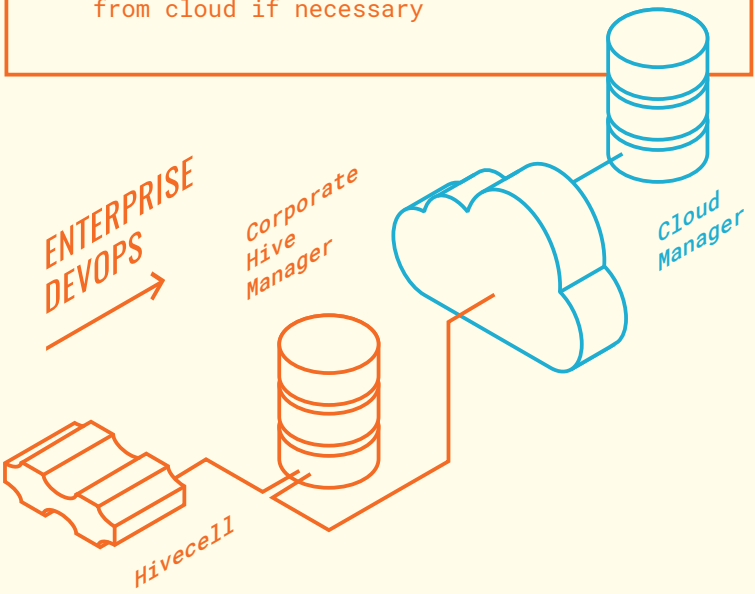

MODE 3: DISCONNECTED

In some edge cases, it may not be possible to have an Internet at all. The location may be too remote or too secure. Hivecell enables a third and particularly novel approach for deploying software.

1. Engineers create upgrade image in lab and write to a Hivecell
2. The upgrading Hivecell is physically carried into a secure environment
3. The upgrading Hivecell is placed on top of a Hive in the secure environment
4. The Hive checks the digital signature of the upgrading Hivecell and vice versa
5. The Hive implements the upgrade instructions
6. The upgrading Hivecell is removed and can be used to upgrade other Hives if necessary

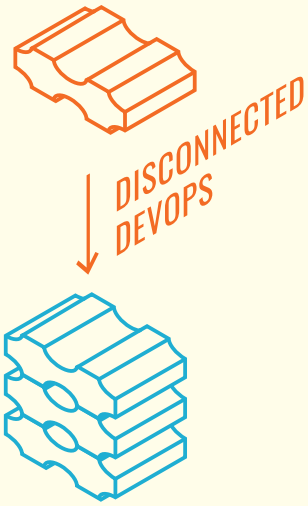

DEVELOPMENT

A key part of removing the complexity of the edge is removing the hardware variable. By only deploying to Hivecells, one greatly simplifies the development, testing and monitoring processes. If it works on Hivecells in the development lab, then one knows it will work at the edge. Any challenges faced after deployment can be isolated to other variables such as network configurations and permissions.

As such, companies should include Hivecells as part of its DevOps infrastructure. We recommend that software engineers have Hivecells for their development and testing environments.

RICKER LYMAN

HIVECELL®

319 Main Street Beacon,
NY 12508 845.216.2248

http://hivecell.io