mimik

Cloud Decentralization & the Trustless Revolution

Fay Arjomandi

What trends drive the need for cloud decentralization? What is the mimik technology platform? How does mimik relate to the consensus-based trustless transactional revolution? What is the mimik exchange?

Software development trend: from monolithic to ephemeral serverless microservice software development

The complexity of managing software solutions made of many components in the past led to monolithic solutions. However, with the consumerization of on-demand IT, the advent of rich communication APIs, and light-weight virtual machines such a Docker & CoreOS, this complexity has been drastically reduced. As a result, today, most software is developed as a collection of many instances of single purpose, well-defined system components called microservices. The consequences of designing systems this way are:

- more granular utilization of infrastructure resources to align with the demand curve
- simplification of the design of complex attributes (session, tenancy)
- better distribution and utilization of computing resources within or between data centers

Furthermore, to achieve even more efficiency for software solutions, a current trend is to use ephemeral microservices within a serverless architecture. Microservices are instantiated (launched and run) based on API calls received by API gateways.

Traffic pattern trend: from consumer devices that mostly consume content (hosted on the cloud) and generate mostly downlink (server to client) traffic to powerful devices that produce content and generate significant uplink (client to server) traffic.

The ever-increasing computing power, storage capacity, memory, graphic processing, etc. on edge devices (mobile phones, game consoles, Smart TVs, IoT devices, etc.) that are usually connected to high-speed networks, is changing the distribution of activities in a "client-server" environment. As a result, the hyper-connected world is generating massive uplink and downlink traffic compared to the past where most of the traffic was generated and transmitted on the downlink. In the past, fewer devices (hereby referred to as nodes) were connected directly to the internet. This shift in traffic pattern has stressed both the network and central cloud infrastructure. Last year millennials posted more than 35 Billion selfies on social media, 27 Billion photos (13.7 petabytes) were uploaded to google photo app alone. The communication model amongst people has shifted from text-based communication to rich media. The up and coming augmented and virtual reality (AR/VR) technologies will further increase the amount of data exponentially. According to the Intel CEO, one autonomous vehicle generates 4TB of data per day equivalent to the traffic generated by 3000 people in one day. As a result, both network and cloud providers are ceasing the opportunity to use edge nodes to offload computing from servers in the central cloud to edge nodes. This is obviously more economical and efficient.

There are enormous unused computing resources on the edge that sit idle most of the time. For instance, 50 million Sony PlayStation 4 game consoles have roughly four times the computing, 2.5 times the storage, 2.5x computing and five times the memory of the Amazon data centers globally. There are billions of laptops, mobile phones, WiFi access points, IOT gateways, routers, set-top-boxes, enterprise servers, connected appliances, autonomous vehicles, AR/VR equipment, etc. These edge devices collectively have millions of times the computing resources of central cloud.

Many have recognized the opportunity to leverage computing resources at the edge. For instance, network infrastructure vendors are experimenting with adding new functions to cellular base stations such as content cashing, content delivery, and edge computing.

Similarly, cloud providers have introduced home gadgets such as Amazon Echo and Google Home to offload processing from the cloud to edge nodes. Running small tasks on the central cloud has similar overhead costs as running large tasks. Therefore, cloud providers have much lower profit margins or even lose money running smaller tasks on central cloud resources. Moving some of the processing to the edge provides the cloud providers increased profitability and better performance since:

- the traffic is localized and bandwidth usage is minimized
- computing, storage and electricity is provided by the consumer, and
- there is much more visibility to consumer behavior that can improve personalization and increase opportunities to further monetize consumer data.

We are also witnessing a similar strategy from large software application companies such as Facebook. Their transition from HTML5 to native client application is mostly aimed at reducing their cost of cloud hosting by utilizing the processing power and storage of consumer devices. This move has also helped improve consumers' perceived quality of experience with faster response times and some offline availability of content.

Microservice-based software development and offloading of computing and communication to edge resources reduces latency, increases network capacity utilization and reduces the cost of cloud hosting services.

These converging trends of 1) ephemeral microservices software development, 2) massive underutilized edge resources and 3) traffic pattern change and offloading to the edge, creates new opportunities to optimally decentralize the cloud and the Internet. To achieve this, we need to:

- treat all available computing devices from servers in large data centers to edge mobile nodes as potential cloud server nodes
- allow edge nodes to dynamically discover each other and form their own ad-hoc clusters based on predefined scopes such as network, proximity and account

- turn each edge node into a cloud node where one can remotely download, install, start and stop ephemeral microservices that are dynamically instantiated and are serverless
- ensure that nodes are two-way connected (directly or indirectly through other nodes) to the internet
- there is transparent information on the availability of resources on edge nodes and/or cluster of nodes (edge cloud)
- enable peer-to-peer communication between the nodes without the necessity for a central cloud element
- minimize signaling and bearer communication to ensure scalability at the system-level

This approach, expands the boundaries of “the cloud” beyond data centers to all computing resources from large data centers all the way to consumer devices on the edge of the network. This in effect creates a cloud fabric that is potentially millions of times larger than the existing central cloud infrastructure. Most computing resources on the edge are under-utilized. We can discover and expose these under-utilized computing resources and utilize them in an opportunistic way for the direct benefit of developers and consumers.

Moreover, we can add data analytics to ephemeral microservices; not only based on availability, policy and context but also including social and other application-level events. This way, software solutions can optimally deploy themselves where it makes the most sense whether on a big computing farm (central cloud) or on edge resources (edge cloud). The result is an evolving continuum of swarm of microservices using computing resources as needed from edge nodes in clusters all the way up to large server nodes in data centers.

mimik decentralized edge cloud provides a platform that expands the central cloud boundaries to include edge resources. The evolution of central cloud started by the likes of Amazon offering IT-as-a-Service platform to the developer community. They achieved this by building large datacenter facilities equipped with server blades (resources) exposed to developers for ease of development and deployment. Application developers can now leverage cloud resources in various levels of abstraction from Infrastructure as a service (IaaS) to Platform-as-a-Service (PaaS) or even Software-as-a-Service (SaaS).

The mimik decentralized cloud platform turns any computing resource into a cloud node by networking all the nodes (from the center to the edge) and enables development and dynamic deployment of serverless microservices on these nodes. It in effect provides a new IaaS platform that is potentially millions of times larger than the central cloud. This new platform not only is more efficient by harnessing unused computing resources on the edge, but it is essential for many applications that require low latencies and hyper-localization. In addition to cost efficiencies and improved profit margins, solutions like artificial intelligence (AI) often require system-level learning and collaboration across clusters of nodes in a subsystem. In many instances, this is best done using a decentralized and trustless architecture.

The central cloud has helped transition monolithic application development to a much more scalable and efficient API centric and serverless microservice software development. This was made possible by exposing on-demand resources through IaaS, PaaS and SaaS. The mimik decentralized cloud platform will help expand the cloud to include underutilized edge nodes. It will help all computing resources to join and participate in building a community cloud environment. It enables developers to apply their standard API centric, serverless microservice development practices and dynamically deploy microservices on new classes of edge resources as needed. Developers can dynamically use all types of computing resources from powerful servers in data centers (public or private) to mobile phones or even IoT devices. In fact, with mimik, the meaning of the term "cloud' will shift from computing resources in data centers to include generic computing resources that may reside anywhere on the internet. We can harness the collective power of edge devices by enabling communication and collaboration amongst these resources in an ad-hoc fashion. This is the most effective way for the industry to scale global computing resources to enable new applications across all industry sectors; IoT, autonomous vehicles, collaborative drones, cloud robotics, automation, IoT, AR/VR, gaming.

In short, mimik technology:

- turns any computing node into a server node to take part in an expanded cloud infrastructure
- enables the nodes to form clusters dynamically
- enables the nodes to discover, connect and communicate within or across clusters (either directly or through dynamic and opportunistic instantiation of signaling and data bearer resources)
- enables the nodes to dynamically deploy (load and run) microservices.

Central cloud has also created a new business model where application providers pay for computing resources on an as–need model charged by the number of transactions, events, I/O, computational resource usage, etc. This business model has reduced the upfront investment in hardware resources and ancillary software tools. However, it is very difficult to predict the platform costs over time as the solution scales. Edge cloud computing can bring more predictability by expanding the resource base, localization of services, leveraging internal resources for enterprises, and minimizing the overhead of central elements.

Distributed applications and blockchain:

Recently, we are witnessing a movement towards decentralization of financial transactions. Decentralized cryptocurrencies still represent a small fraction of the global financial activity. Yet, its exponential growth has fueled a general movement towards consensus based trustless transactions using blockchain technology. The blockchain ecosystem has started generating revenues from computing assets by putting the resource to work. Block chain technology provides the proof of work that drives monetization in the form digital currencies. There is now a community of miners that gain bitcoin (or other digital currencies) in return for participating with computational power to confirm transactions in a trustful manner.

Although the blockchain technology is decentralized in nature, it currently runs on a centralized cloud infrastructure. Certain blockchain algorithms require specialized computing resources. For instance, bitcoin mining today requires specialized hardware. However, in general, it is possible to run trustless applications on heterogeneous computing devices that pool their resources together and collaborate.

We believe the mimik platform is the missing link to help create a decentralized network of computing resources. We have spent seven years of R&D and created a downloadable software development kit (called edgeSDK) to expand IaaS to edge resources. Many trustless transactional protocols (blockchain, *IOTA*, etc.) can leverage this new expanded cloud. The developer community can develop scalable, portable and sustainable software solutions that are easily maintainable leveraging a pool of nodes that is collectively millions of time larger than the existing “cloud” infrastructure. Each node exposes its capability via well-defined RESTful API(s) to others. Many applications can be developed as easily as with the central cloud but in a consensus-based and trustless fashion. Additionally, mimik enables the owners of computing resources to participate in a large pooled network and share the profits of a potentially large network exchange. From the technology point of view, this enables the community to use a serverless microservice architecture expanded to edge devices and as a result expand its benefits for ease of upgradability, maintenance, and improved fault isolation. On the business side, this allows the owners and consumers of computing resources the freedom to trade directly and participate in the share of a network exchange without any middlemen or trust elements.