Tactical Cloudlets: Moving Cloud Computing to the Edge

Grace Lewis (glewis@sei.cmu.edu)
SEI Webinar
December 10, 2014

Software Engineering Institute | Carnegie Mellon University

Motivation

Soldiers, first responders and field personnel operating in tactical environments increasingly make use of handheld devices to help with tasks such as face recognition, language translation, decision-making, and mission planning

Edge environments are characterized by dynamic context, limited computing resources, high levels of stress, and intermittent network connectivity

Tactical cloudlets provide cloud capabilities at the edge that can lead to enhanced situational awareness and decision making, even if disconnected from the enterprise

Tactical Cloudlets $_1$

Goal: Provide cloud computing capabilities at the edge for computation offload, data staging, and increased survivability of mobile systems

R&D Goals

- Discoverable resources
- Operation in DIL environments (disconnected, intermittent, limited)
- Systems perspective on cyber-foraging that includes survivability, trust, and ease of development and deployment
- Flexible architecture to support research and experimentation

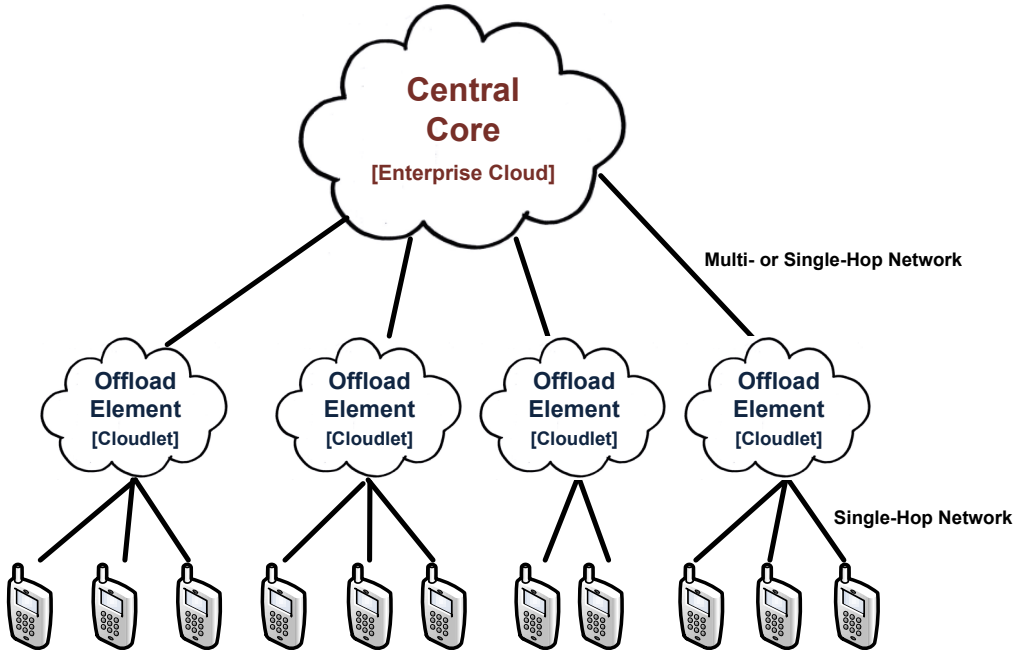

Tactical Cloudlets $_2$

Capabilities: Forward-deployed, discoverable, virtual machine (VM) cloudlets that can be hosted on vehicles or other platforms and provide

- infrastructure to offload computation
- forward data-staging for a mission
- data filtering to remove unnecessary data from streams intended for mobile users
- collection points for data heading for enterprise repositories

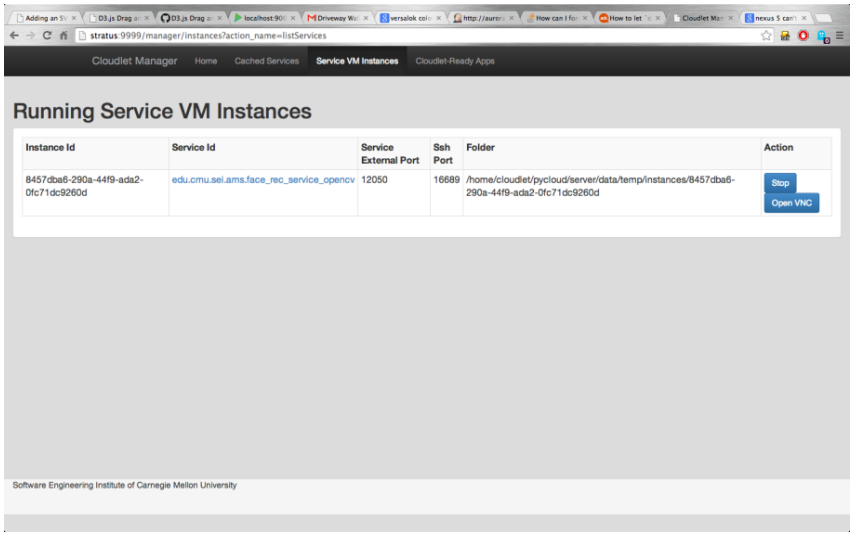

Cyber-Foraging

Cyber-foraging* is the leverage of external resource-rich surrogates to augment the capabilities of resource-limited devices

Two main forms of cyber-foraging

- Code/Computation Offload
  - Offload of expensive computation in order to extend battery life and increase computational capability
- Data Staging
  - Improve data transfers between mobile computers and the cloud by temporarily staging data in transit

* Satyanarayanan, Mahadev (2001). "Pervasive Computing: Vision and Challenges". IEEE Personal Communications (IEEE)

Cloudlet-Based Cyber-Foraging

Discoverable, virtual-machine based, forward-deployed servers located in single-hop proximity of mobile devices

- Can operate in disconnected mode
- Communication with the central core is only needed for provisioning

Applications are statically partitioned into a client and server

- Very thin client runs on mobile device
- Computation-intensive server runs on cloudlet

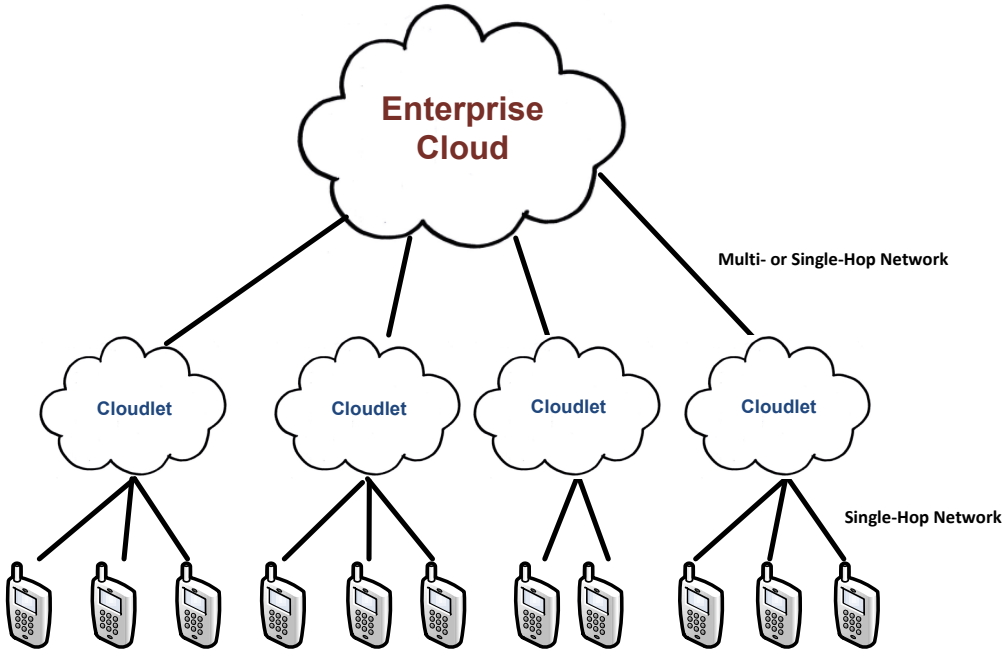

Reference Architecture for Cloudlet-Based Cyber-Foraging

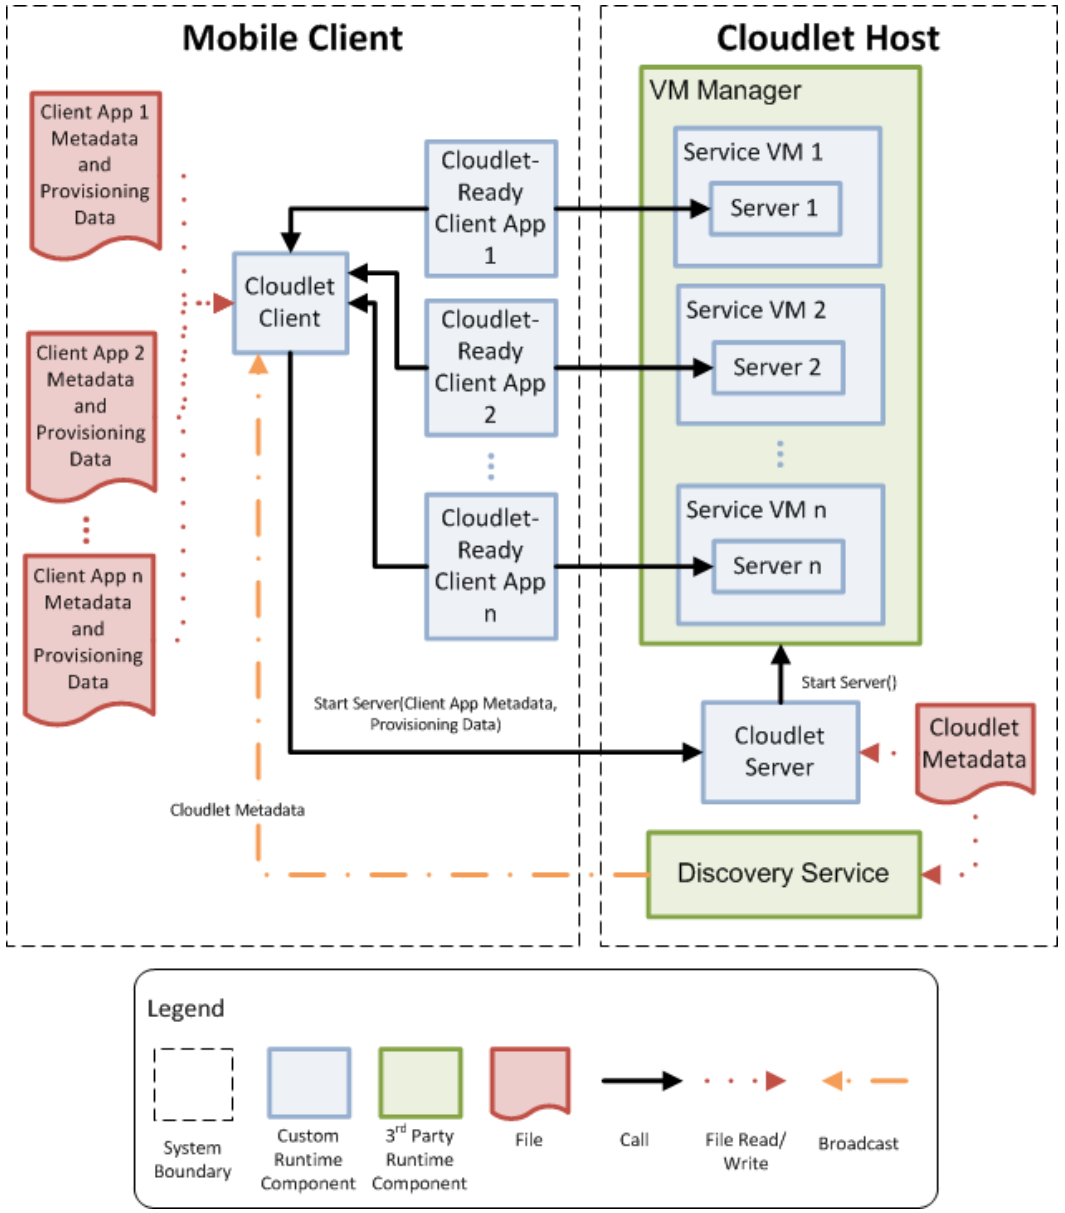

Cloudlet-Based Cyber-Foraging: Operations

1. Cloudlet Discovery

2. Cloudlet Provisioning and Setup

3. Application Execution

Cloudlet Discovery

Discovery Service implementation based on zeroconf

- Uses DNS Service Discovery (DNS-SD) along with Multicast DNS
- Multicast addresses are used to allow the client to request the service without knowing the IPs of the servers

Enables mobile devices to locate available cloudlets

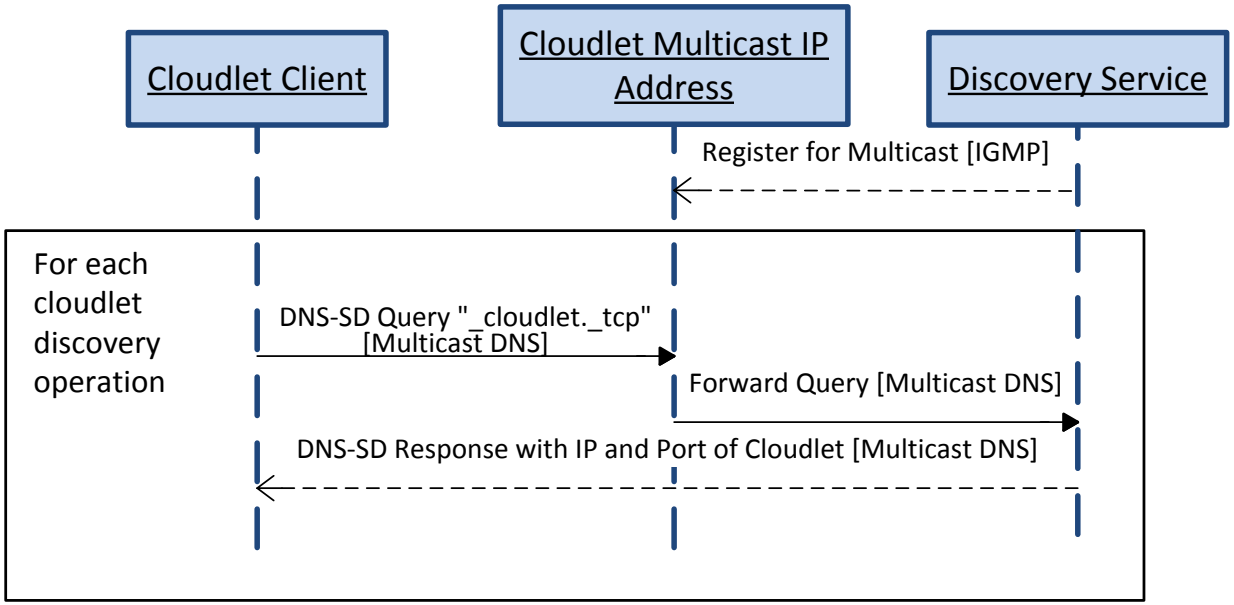

Cloudlet Provisioning

Configuring and deploying the Service VM that contains the server code on the cloudlet so that it is ready to use by the client running on the mobile device

Working prototypes for five different cloudlet provisioning mechanisms

- Runtime — provisioning from the mobile device
  - VM Synthesis
  - Application Virtualization
- Deployment time — cloudlets pre-provisioned based on mission needs
  - Cached VM
  - Cloudlet Push
- On-Demand — capabilities assembled at runtime
  - On-Demand VM Provisioning

VM Synthesis 1

Cloudlet is provisioned by sending an application overlay from the mobile device to the cloudlet at runtime

Application overlays are created in advance for server portions of applications — they represent the diff between a baseline VM and that same VM with the application installed

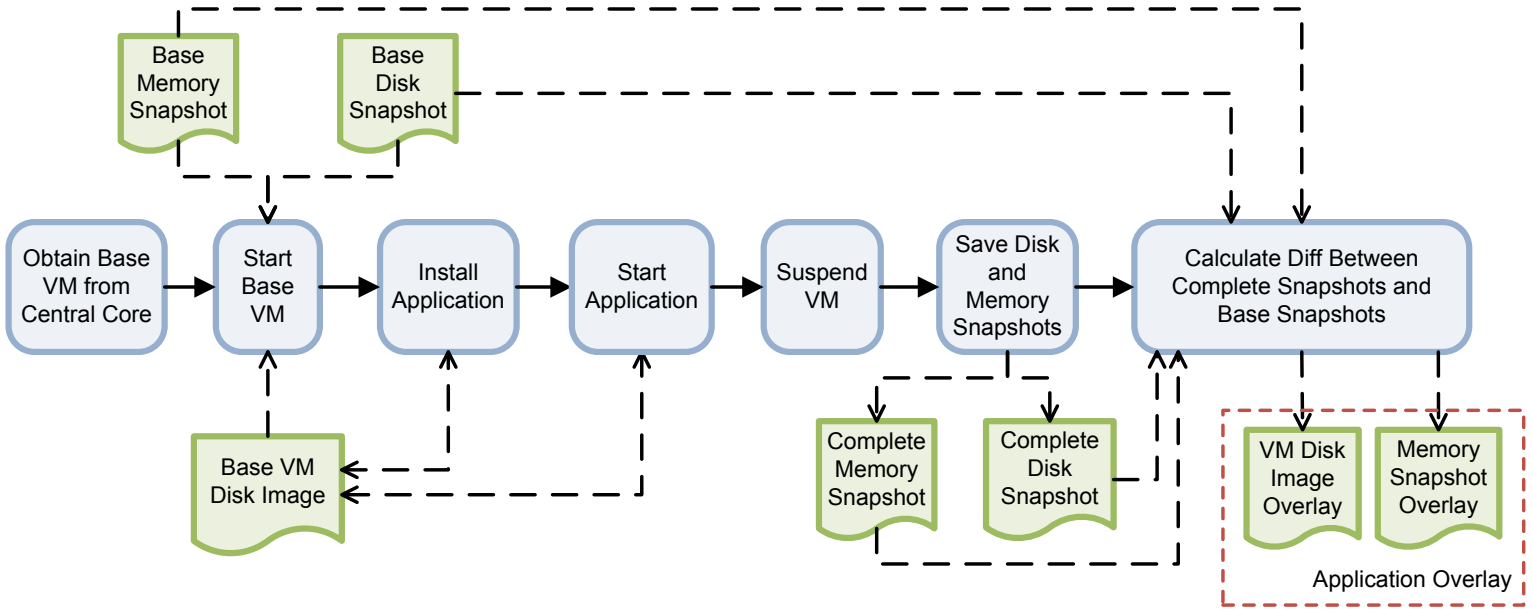

VM Synthesis 2

The reverse of the application overlay creation process is done at runtime to create the Service VM

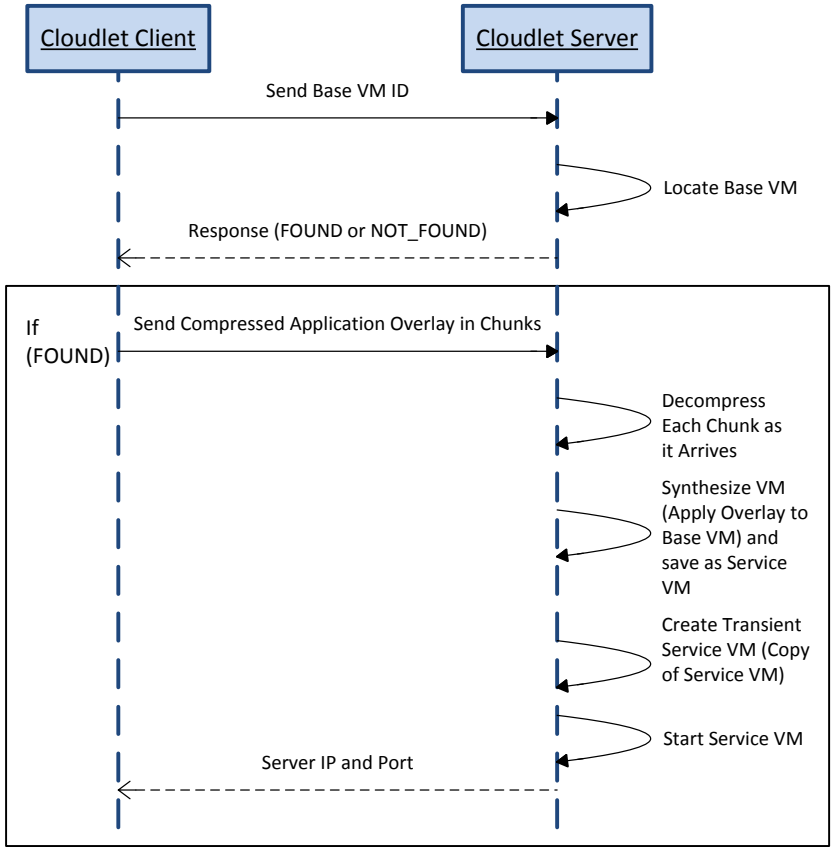

Applications	Payload Size (MB)	Application -Ready Time (s)	Client Energy (J)
FACE (Windows)	55	53.4	57.8
OBJECT (Linux)	332	175.7	333.3
SPEECH (Windows)	194	85.9	175.5
SPEECH (Linux)	147	99.0	172.5

Cloudlet Content	Exact Base VMs
Mobile Device Content	• Application Overlays • Client Apps + Metadata
Payload	Application Overlay
Advantages	Cloudlet can run any server code that can be installed on a Base VM
Constraints	Requires exact Base VM which limits distributions and patches

Application Virtualization

Cloudlet is provisioned by sending a virtualized application from the mobile device to the cloudlet at runtime

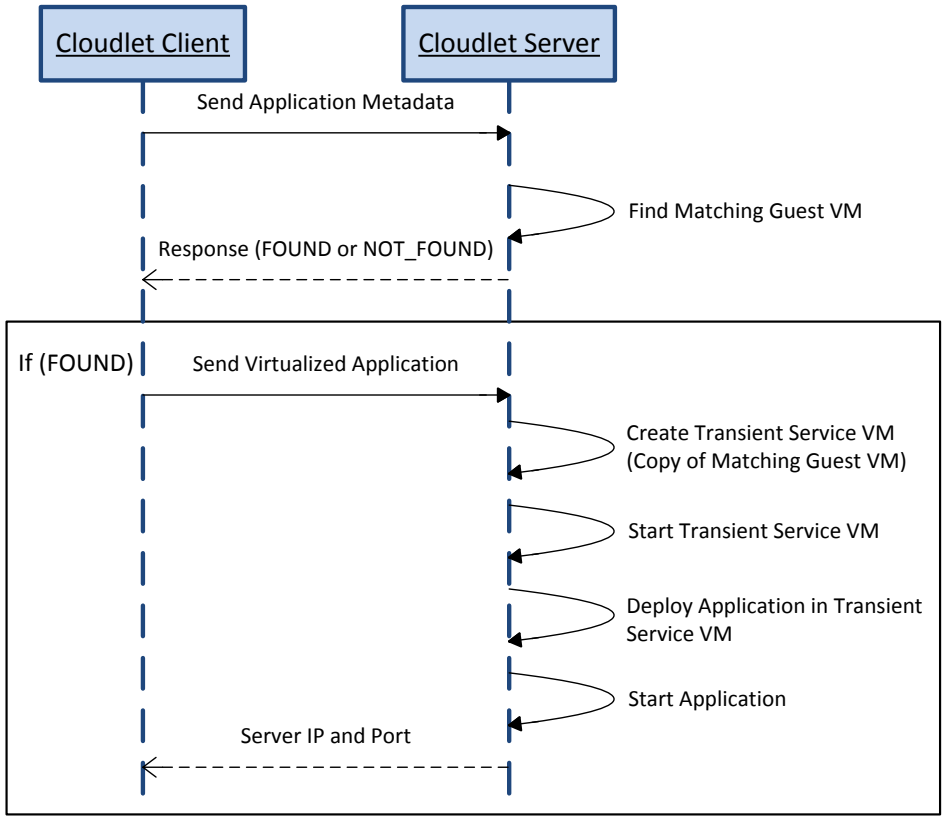

Applications	Payload Size (MB)	Application-Ready Time (s)	Client Energy (J)
FACE (Windows)	14	14.3	10.5
OBJECT (Linux)	29	21.9	24.5
SPEECH (Windows)	66	62.5	66.6
SPEECH (Linux)	68	38.3	54.2

Cloudlet Content	VM compatible with Server Code
Mobile Device Content	• Virtualized server code • Client Apps + Metadata
Payload	Virtualized Server Code
Advantages	Portability across OS distribution boundaries
Constraints	All server code dependencies have to be captured at packaging time

Cached VM

Cloudlet is pre-provisioned based on mission needs

Repository of VMs that contain capabilities

Each VM is treated as a service

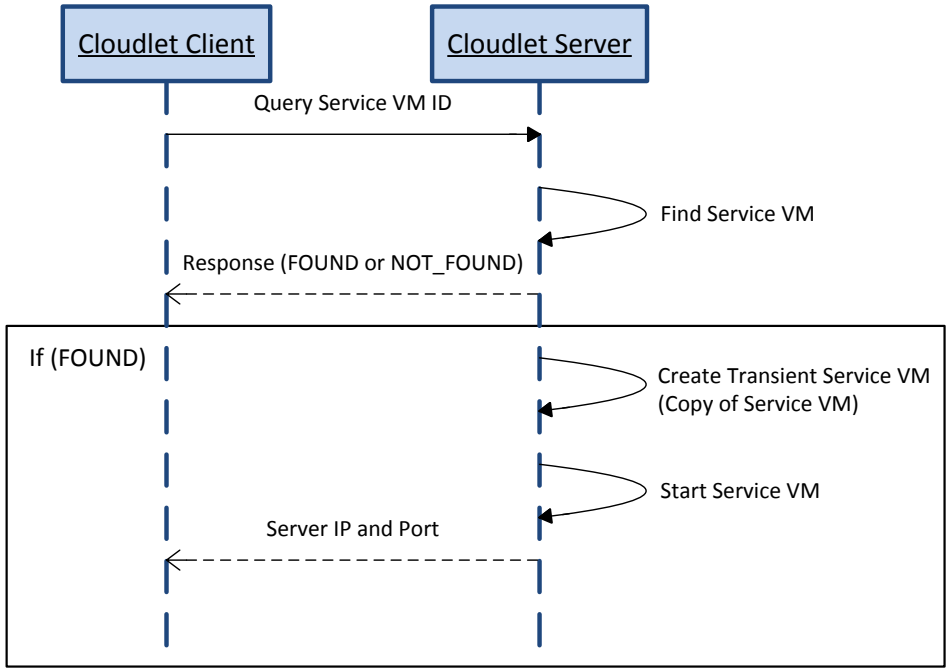

Applications	Payload Size (MB)*	Application -Ready Time (s)	Client Energy (J)
FACE (Windows)	0.00	8.2	10.3
OBJECT (Linux)	0.00	11.6	13.5
SPEECH (Windows)	0.00	12.2	14.7
SPEECH (Linux)	0.00	12.2	14.9

* Size of payload is less than 1KB

Cloudlet Content	Service (VM) repository
Mobile Device Content	Client App + Metadata
Payload	Service ID
Advantages	Supports server code updates as long as service interface remains the same
Constraints	Cloudlet is provisioned with service VMs required by client apps (or has access to them)

Cloudlet Push

Cloudlets are pre-provisioned and corresponding client apps are pushed to the mobile device at runtime

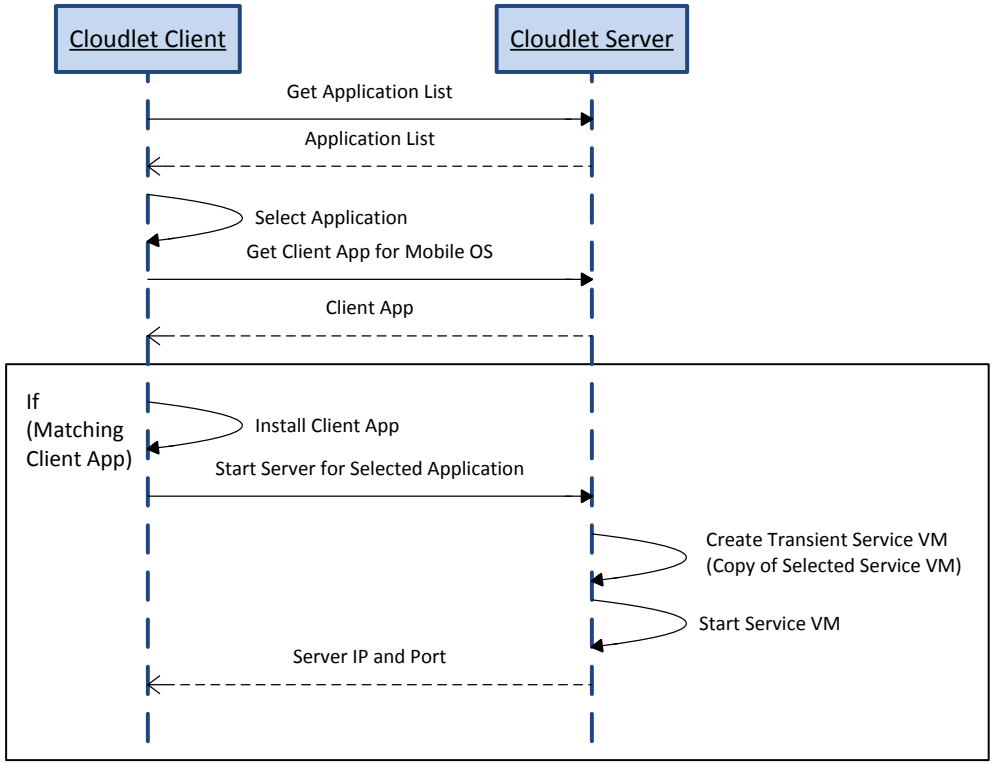

Applications	Payload Size (MB)*	Application -Ready Time (s)	Client Energy (J)
FACE (Windows)	0.0	7.9	13.8
OBJECT (Linux)	0.0	11.7	16.9
SPEECH (Windows)	0.0	12.8	18.2
SPEECH (Linux)	0.0	12.8	18.2

Cloudlet Content	Repository of Paired VMs (Server code) and Client Apps
Mobile Device Content	None
Payload	Client App and Metadata
Advantages	Supports most client mobile devices with distribution at runtime
Constraints	Cloudlet has a client app version that matches mobile client OS version

On-Demand VM Provisioning

Cloudlets are provisioned by assembling a server VM at runtime using a cloud provisioning tool, according to a provisioning script sent from the mobile device

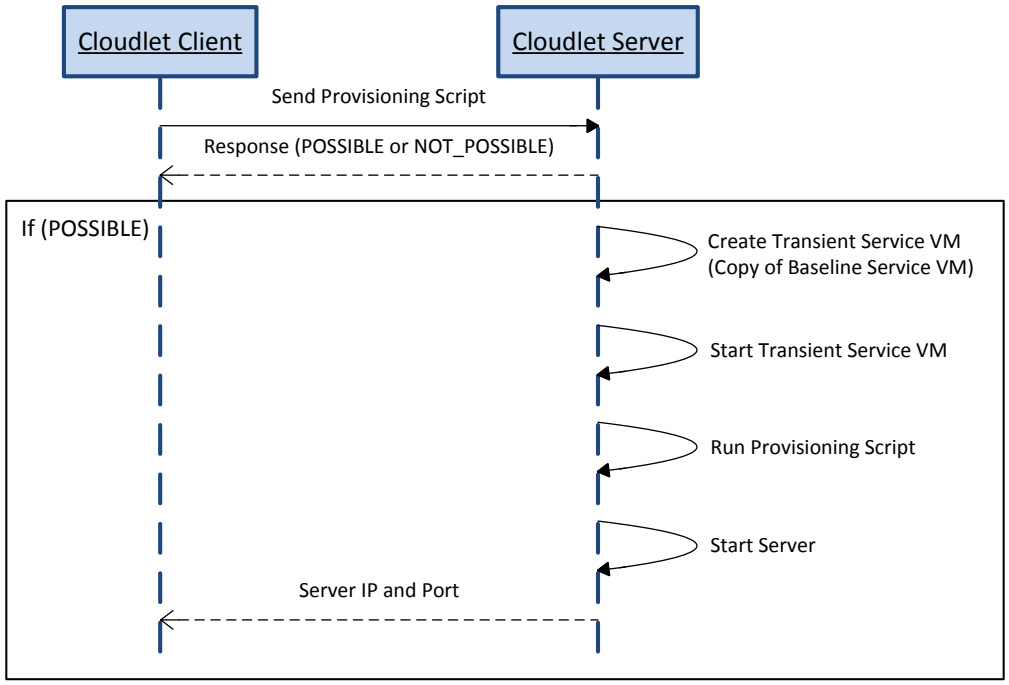

Applications	Payload Size (MB)*	Application-Ready Time (s)	Client Energy (J)
FACE (Windows)	0.0	112.7	129.1
OBJECT (Linux)	0.0	211.0	244.0
SPEECH (Windows)	0.0	237.6	269.2
SPEECH (Linux)	0.0	94.1	109.3

Cloudlet Content	• VM provisioning software • Server code components
Mobile Device Content	• VM provisioning script • Client App + Metadata
Payload	VM Provisioning Script
Advantages	Service VM with server code can be assembled at runtime
Constraints	Cloudlet has all required server code components (or access to them)

Application Execution

After receiving the Service VM IP address and port, the Cloudlet Client returns this information to the Cloudlet-Ready Client App

The Client-Ready App opens a socket to the Service VM IP address and port and executes in client/server mode until the app is closed

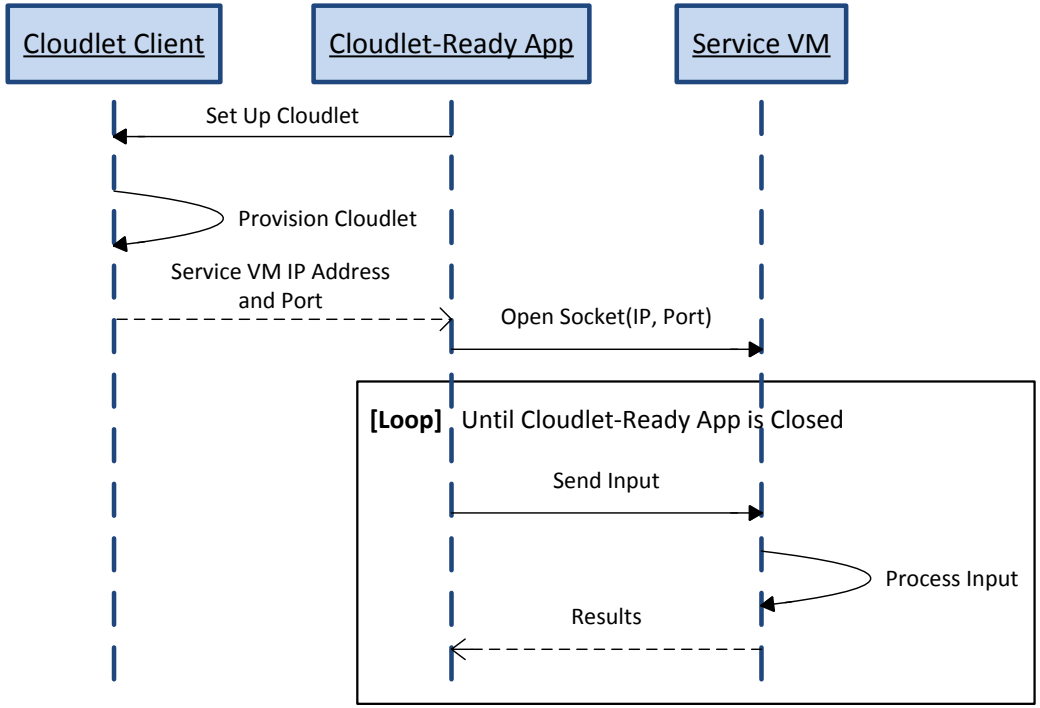

Selected Tactical Cloudlet Implementation

Combination of Cached VM with Cloudlet Push

- Lower energy consumption and less requirements placed on mobile device
- Simple provisioning — if the mobile device already has the client app it can simply invoke the matching Service VM; if not it can obtain the client app from the cloudlet — similar to accessing an app store — and then invoke the matching Service VM
- Promotes resilience and survivability by supporting rapid live VM migration in case of cloudlet mobility, discovery of more powerful or less-loaded cloudlets, or unavailability due to disconnection or disruption
- Supports scalability and elasticity by starting and stopping VMs as needed based on the number of active users
- Request-response nature of many operations lends itself to an asynchronous form of interaction in which the cloudlet can continue processing and send results back to a mobile device as network conditions change

Tradeoffs

- Relies on cloudlets that are pre-provisioned with server capabilities that might be needed for a particular mission, or that the cloudlet is connected to the enterprise, even if just at deployment time, to obtain the capabilities

Tactical Cloudlet Architecture

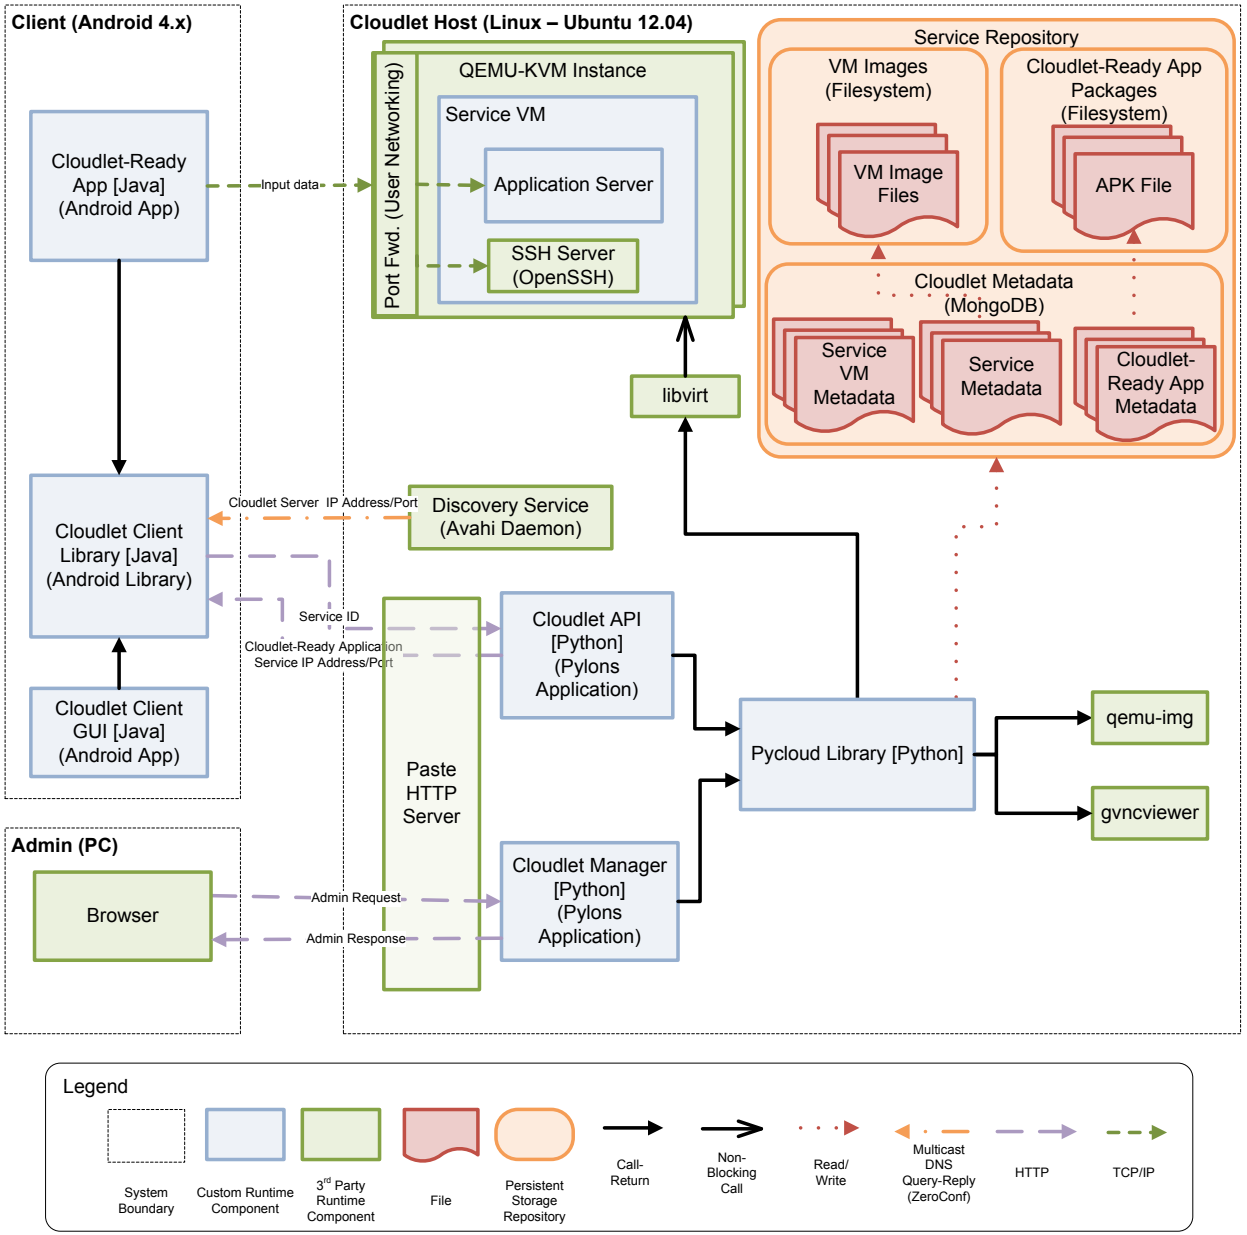

Execution from Cloudlet Client GUI

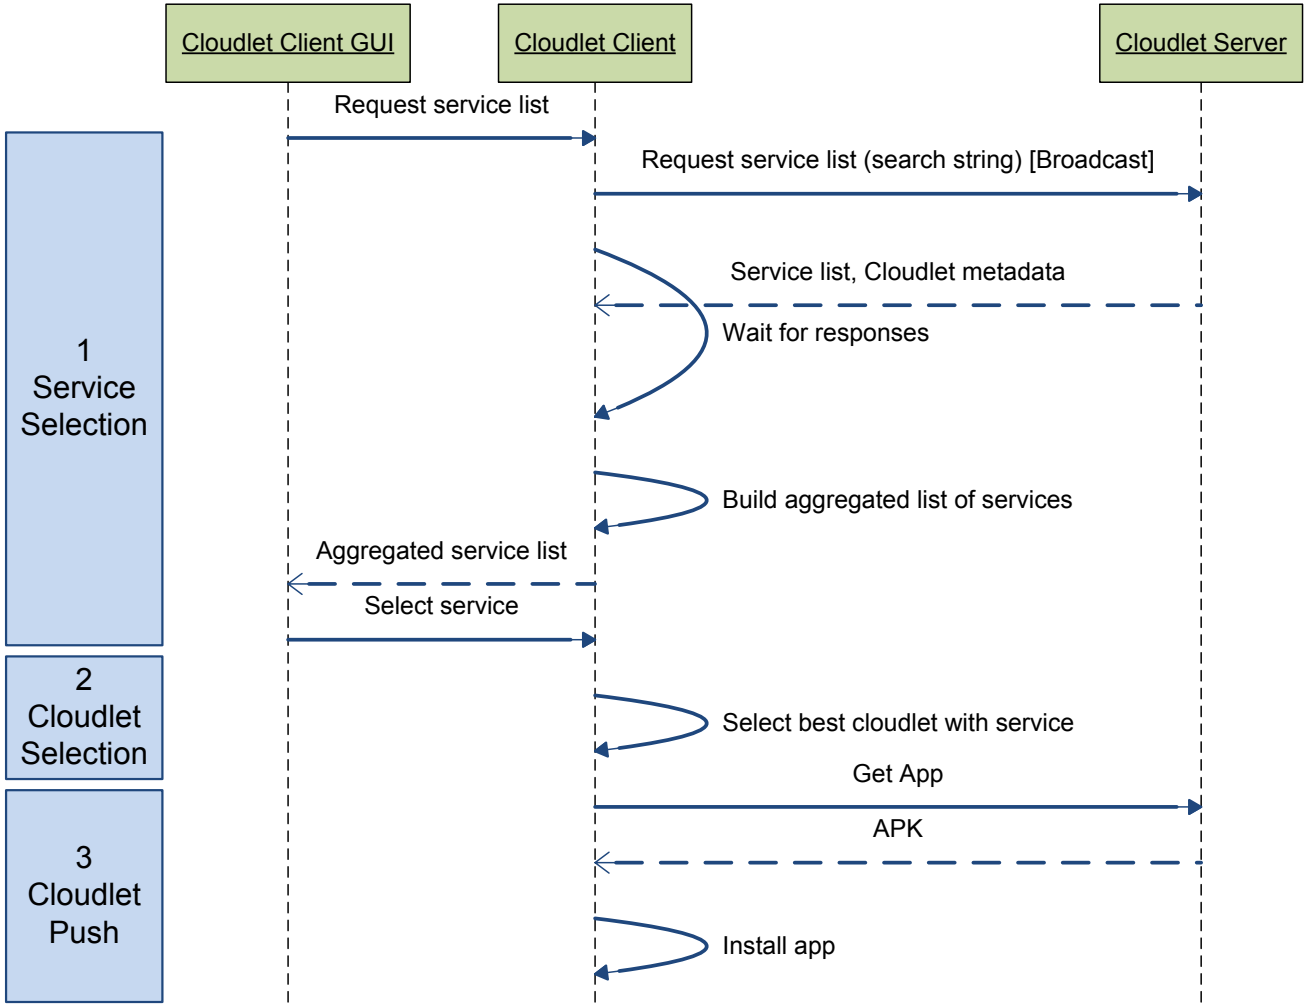

Execution from Cloudlet-Ready App

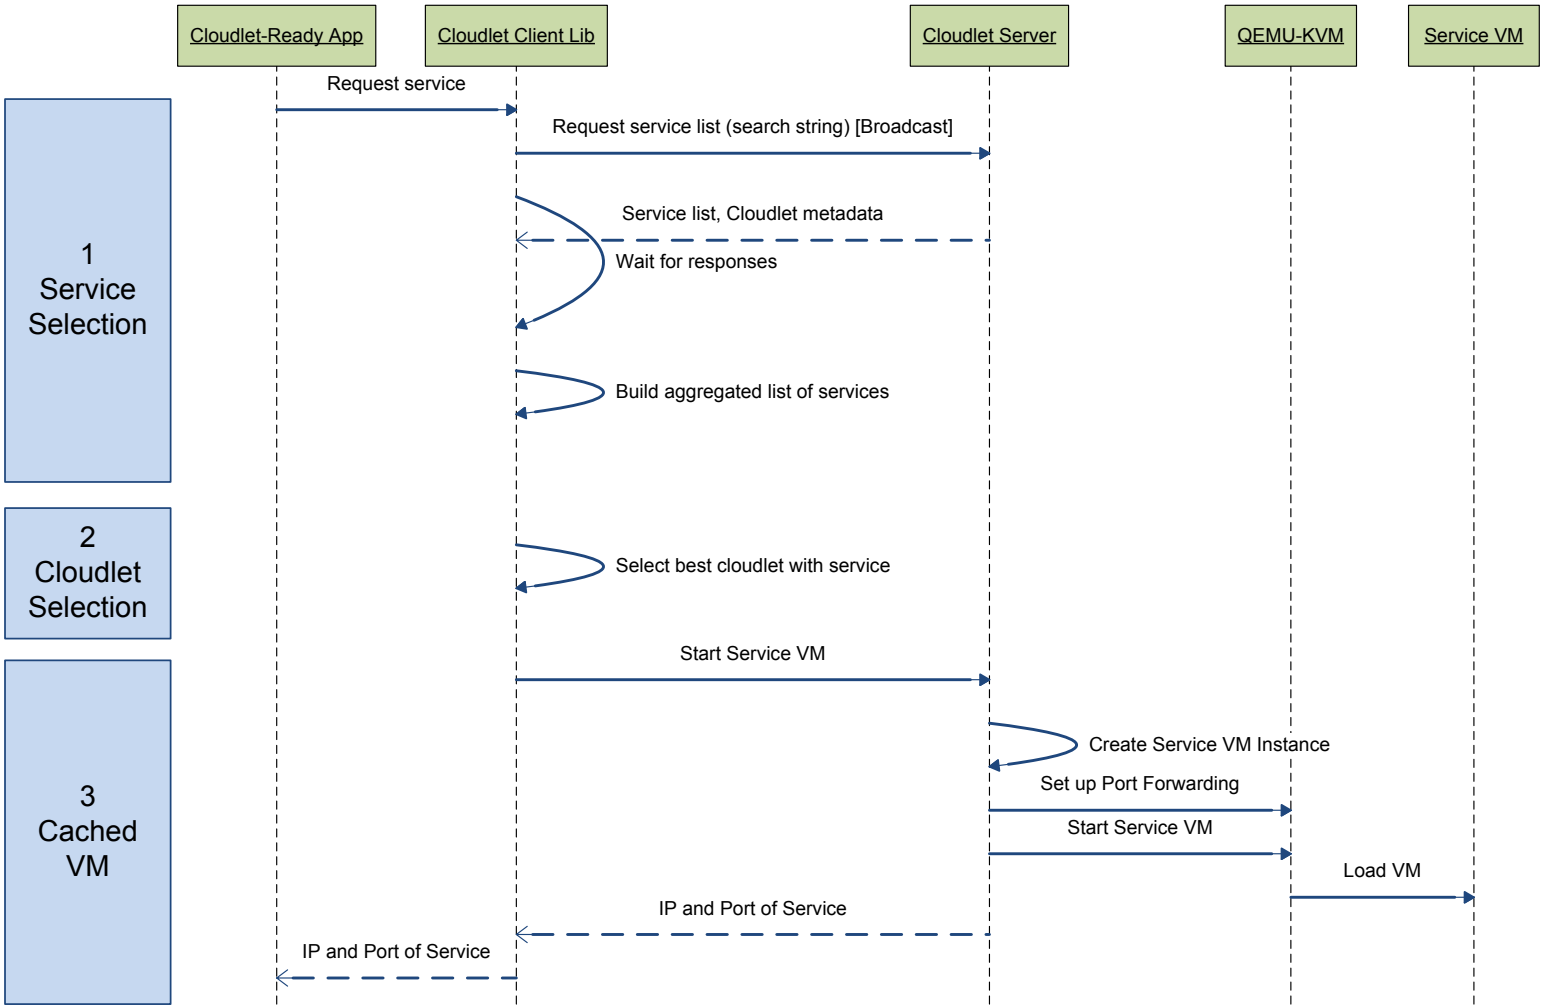

Current and Future Work

Standard packaging of Service VMs

- Installed from the cloudlet manager, enterprise Service VM repository, thumb drive, or the mobile device connected via USB to the cloudlet
- Capabilities to improve mobile systems survivability
- Optimal cloudlet selection
- Cloudlet handoff (live migration) — manual and automatic based on load and other attributes
- Support for data-reliant systems running on cloudlets disconnected from the enterprise

Focus for FY15 will be Trusted Identities in Disconnected Environments

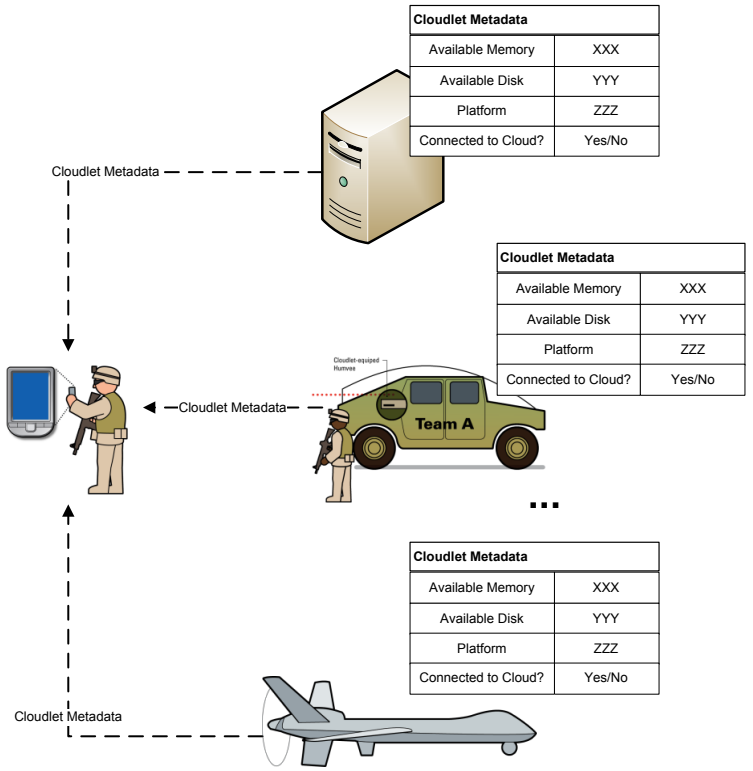

Cloudlets: Beyond Tactical Environments

Goals

- Bring the cloud closer to the user
- Support rich sensing and interaction capabilities of mobile devices seamlessly fused with compute-intensive and data-intensive processing

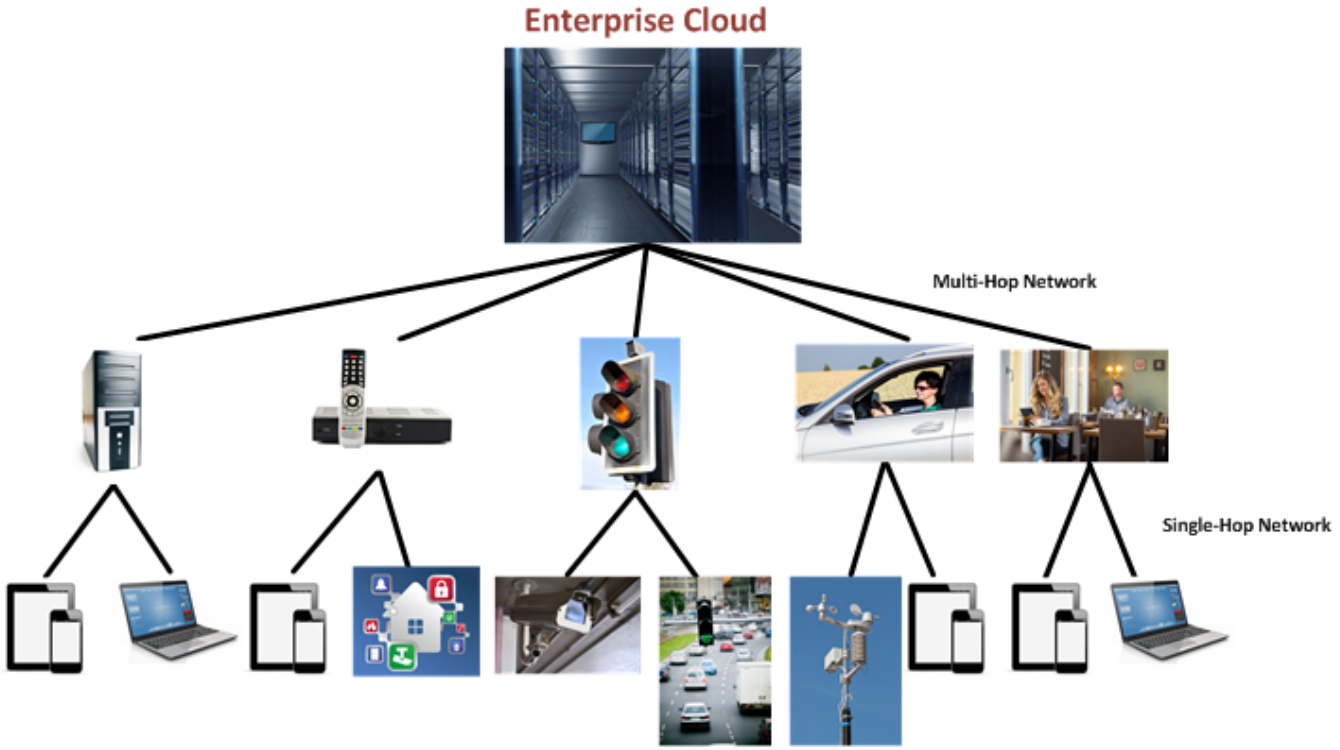

Mobile Device Trends

Smartphones and tablets have become for many the preferred way of interacting with the Internet, social media and the enterprise

- Number of smartphones has passed the number of laptops and desktops[1]
- Growth rate of e-readers and tablets is higher than smartphones, and if continued will reach high numbers soon[2][6]
- Smartphones and tablets are becoming the main computing device for many users[3][6]
- Not uncommon for there to be multiple mobile devices per user and household[4][5]

Organizations are pushing out more and more content and functionality to mobile users

Therefore …

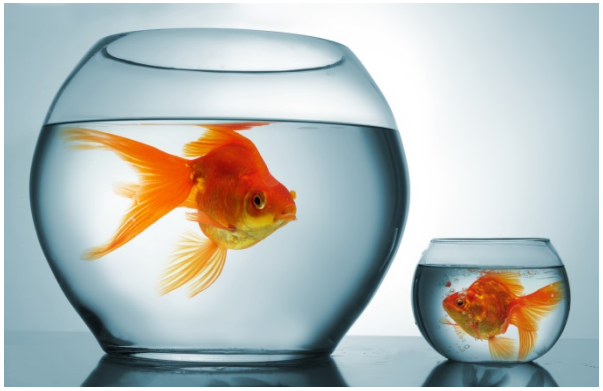

Not unreasonable for users to expect the performance and capabilities of mobile devices to be equal to laptops and desktops

However …

- Mobile devices will always lag behind their PC counterparts due to size and battery limitations
- Large and variable end-to-end latency between mobile device and cloud, and the possibility of disruptions, have a negative effect on user experience
- Will only get worse with the amount of network traffic generated by IoT

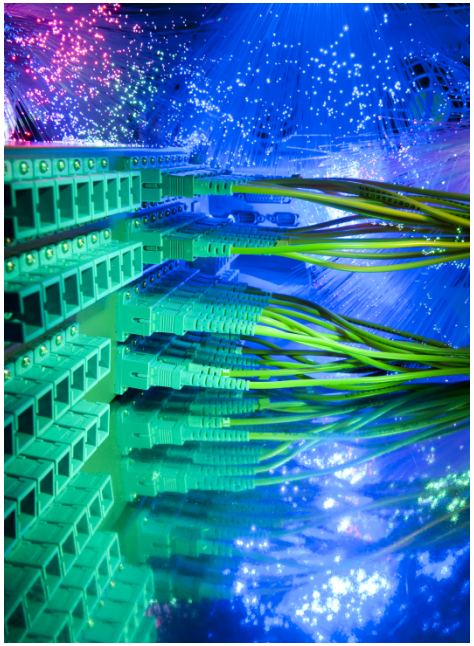

Food for Thought

With increasing number of mobile devices and users, increased network traffic cause by IoT, and increasing complexity of user experience, cyber-foraging will become a standard feature of mobile applications

Requires mobile applications and infrastructures to be architected and designed to adapt to a changing environment in which resources with greater computing power are discovered and used opportunistically

While the benefits in terms of mobile user experience and new business opportunities are huge, it requires a different paradigm in mobile software engineering

Contact Information

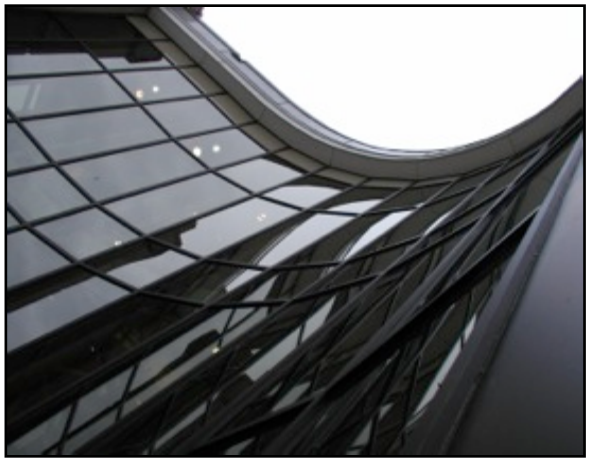

Grace A. Lewis

Advanced Mobile Systems (AMS) Initiative

Software Engineering Institute
4500 Fifth Avenue
Pittsburgh, PA 15213-2612
USA

Phone: +1 412-268-5851
Email: glewis@sei.cmu.edu
WWW: http://www.sei.cmu.edu/staff/glewis/

References

[1] mobiThinking. Global Mobile Statistics 2014. http://mobithinking.com/mobile-marketing-tools/latest-mobile-stats (2014)

[2] Gartner. Gartner Says Worldwide PC, Tablet and Mobile Phone Shipments to Grow 4.5 Percent in 2013 as Lower-Priced Devices Drive Growth. http://www.gartner.com/newsroom/id/2610015 (2014)

[3] K. McGrane. The Rise of the Mobile-Only User. Harvard Business Review. HBR Blog Network. March 2013. http://blogs.hbr.org/2013/05/the-rise-of-the-mobile-only-us/ (2014)

[4] CTIA. Wireless Quick Facts. http://www.ctia.org/your-wireless-life/how-wireless-works/wireless-quick-facts (2014)

[5] Mobile Operators Association. Stats and Facts. http://www.mobilemastinfo.com/stats-and-facts/ (2014)

[6] KPCB. Internet Trends 2014. http://www.kpcb.com/internet-trends (2014)

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM-0001955