

Mobile cloud computing

Vivek Shrivastava
Computer Sciences

Key Papers

- The case for cyber foraging
 - Balan et al, CMU
 - Sigops 2002
- Slingshot: Deploying stateful services in Wireless Hotspots
 - Su et al, University of Michigan
 - Mobisys 2005
- CloneCloud: Augmented smart phone applications
 - Chun et al, Intel Labs
 - HotOS 2009

What is cloud computing?



- *I don't understand what we would do differently in the light of Cloud Computing other than change the wordings of some of our ads*

Larry Ellison, Oracle's CEO

- *I have not heard two people say the same thing about it [cloud]. There are multiple definitions out there of "the cloud"*

Andy Isherwood, HP's Vice President of European Software Sales

- *It's stupidity. It's worse than stupidity: it's a marketing hype campaign.*

Richard Stallman, Free Software Foundation founder

What is cloud computing ?

- Application is used as an **on demand service**.
Often provided via the Internet
 - Think on-demand TV programs
- Example: Google App (online office)
- Benefits to users
 - Reduce expenses: multiple computers, multiple users
 - Ease of usage: easy installation, access everywhere
- Benefits to providers
 - Easier to maintain
 - Control usage (no illegal copies)



What is cloud computing ?

- Computing resources are provided as on demand service
 - CPU hour
 - Memory
 - Network
 - Platform to run software
- Examples of cloud service providers: Amazon EC2, Google AppEngine ...

Now to mobile computing ...

Mobile usage is growing rapidly !

More People



1.7 Billion Internet Users by End 2009

More Places



Many of these are increasingly Mobile users

**Spending
More Time**



Average User Spends 32 Hours per Month Online

**Doing More
Things**



**Over 85000 iPhone Applications, 2 Billion
downloads
More collaborative and interactive than ever before**

**Consuming
Tons of data**



**Predicted to touch 1 ZettaByte by 2015
HD audio/video; virtual worlds; online file servers,
Telepresence**


Anytime-anywhere access becoming expected

2 Broad class of applications

- Personal Productivity Applications
 - Email
 - Calendar
- Computationally Intensive Interactive Applications
 - Speech Recognition
 - Language Translation
 - Augmented Reality

Motivation: Handhelds are weak!

- Resource intensive App
- Huge Data Sets



2 GHz, 1 GB,
3-D graphics
2 GB of data

200 MHz, 32 MB,
no 3-D, no FPU
32 MB Flash

Resource-poor
wearable

Poor performance!

Solution: Cyber Foraging



- “To live off the land”
- Use resources in environment to augment device capabilities by using surrogates
- 2 methods
 - Data Staging
 - Remote Execution

The Big Picture

■ Data Staging

- Caching of large amounts of data
- Handhelds with limited storage can access this data fast
- Security and authentication

■ Remote Execution

- Uses remote servers to augment computational capabilities of handhelds
- Enables computationally intensive applications

Roadmap

- Data Staging
- Remote Execution
 - Slingshot
 - Clonecloud

Data Staging: Motivation


- End-to-end latency across the Internet isn't getting better
 - Physical limits
 - Routers, firewalls
 - Shows up in interactive file access delays
 - Crucial for small to medium files
- Can overcome this by caching & prefetching, but ...
 - Handheld clients don't have enough resources
 - Cache consistency
- *Can untrusted and unmanaged computers help?*

Yes!!



Data staging: Mechanism

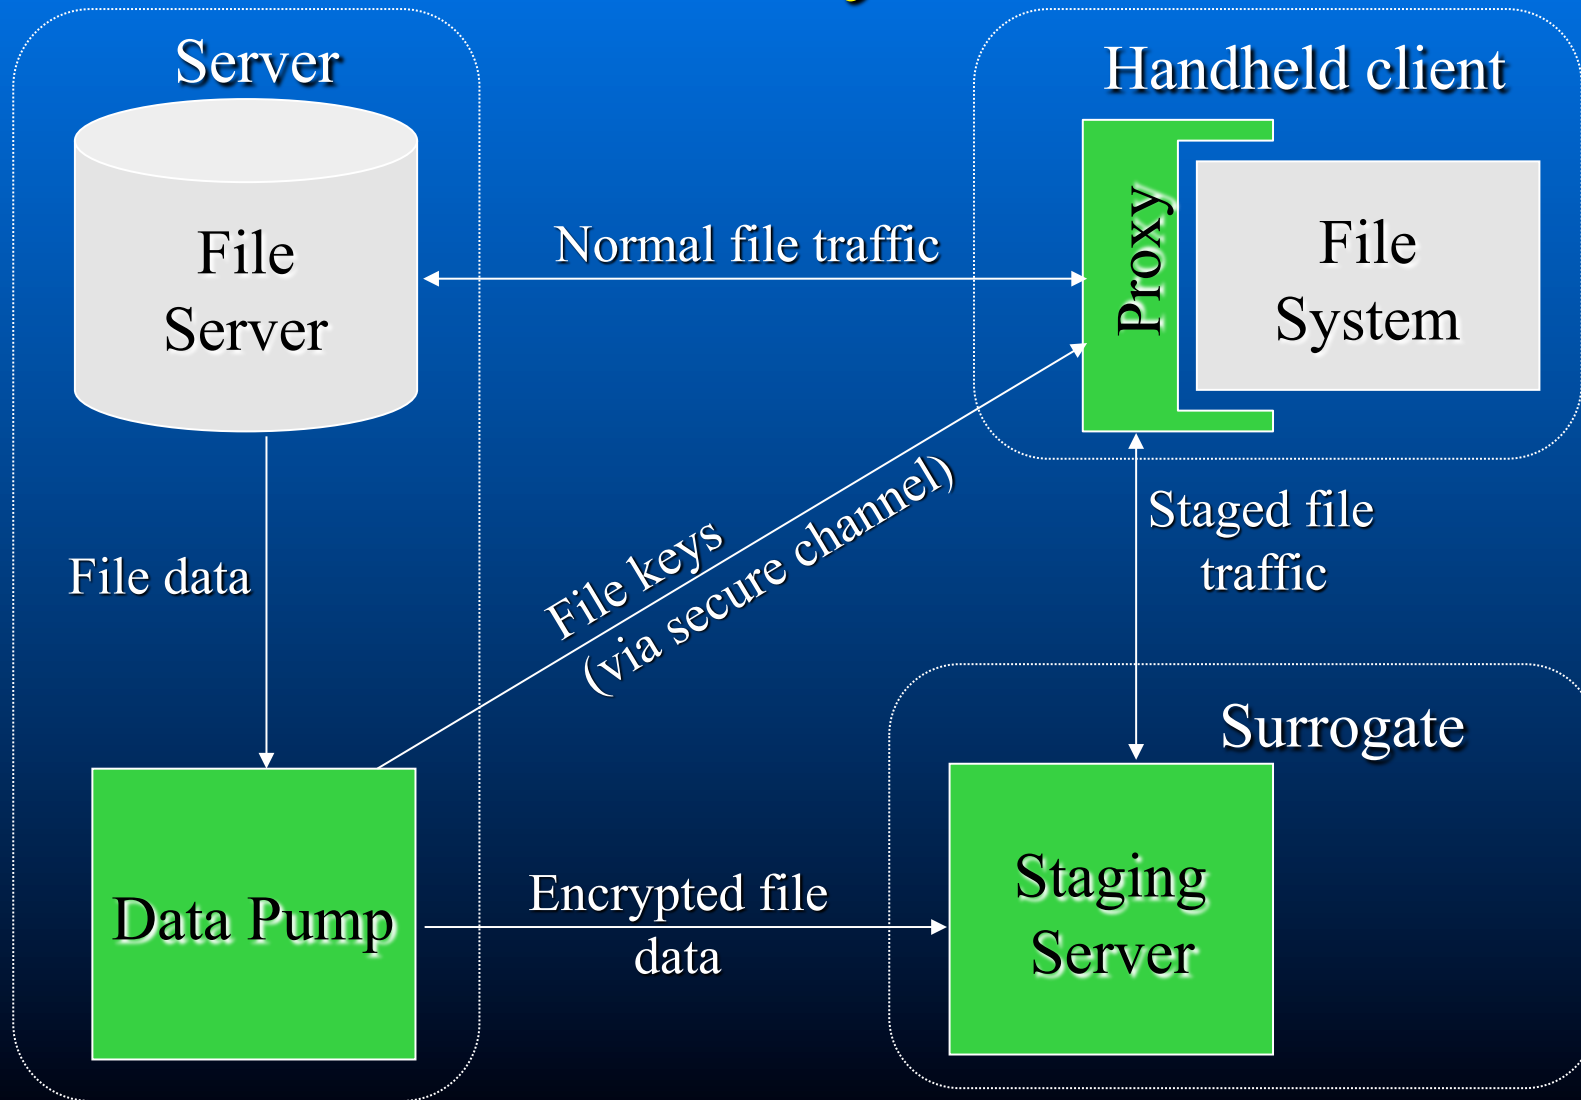
- Coda clients speculatively prefetch data :
 - Nearby **surrogate** runs **staging server**
 - Used like a second level cache
 - Cache misses serviced by staging server
- Surrogates deployed in high-usage areas



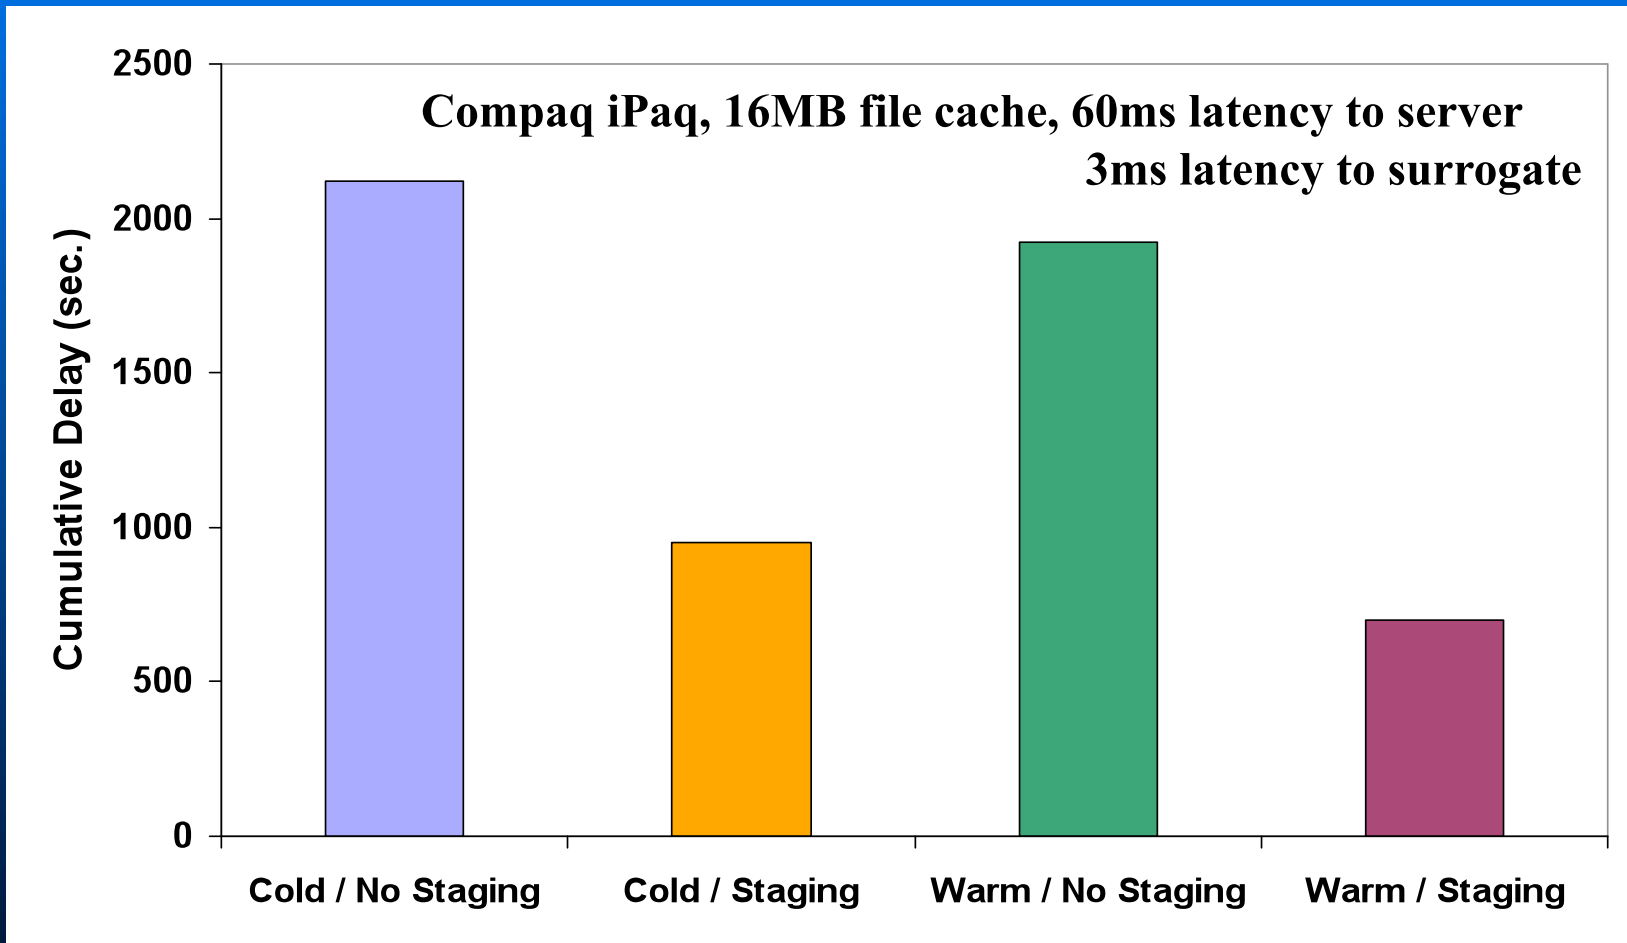
Security

- Must provide level of security users expect
- *But surrogate is untrusted*
- Use end-to-end encryption
 - Only store encrypted data on surrogate
 - Client caches keys and checksums
 - Only need access control for keys

The “Gory” Details



Benefit for image viewing




Data staging reduces cumulative delay up to 64%

Roadmap

- Data Staging
- Remote Execution
 - Clonecloud
 - Slingshot


Motivation: mobile interactive applications

- speech recognition, language translation, augmented reality, ...
 - Resource-heavy, but need bounded response time



Solution: Remote Execution

- Augment capabilities of handhelds by using nearby servers



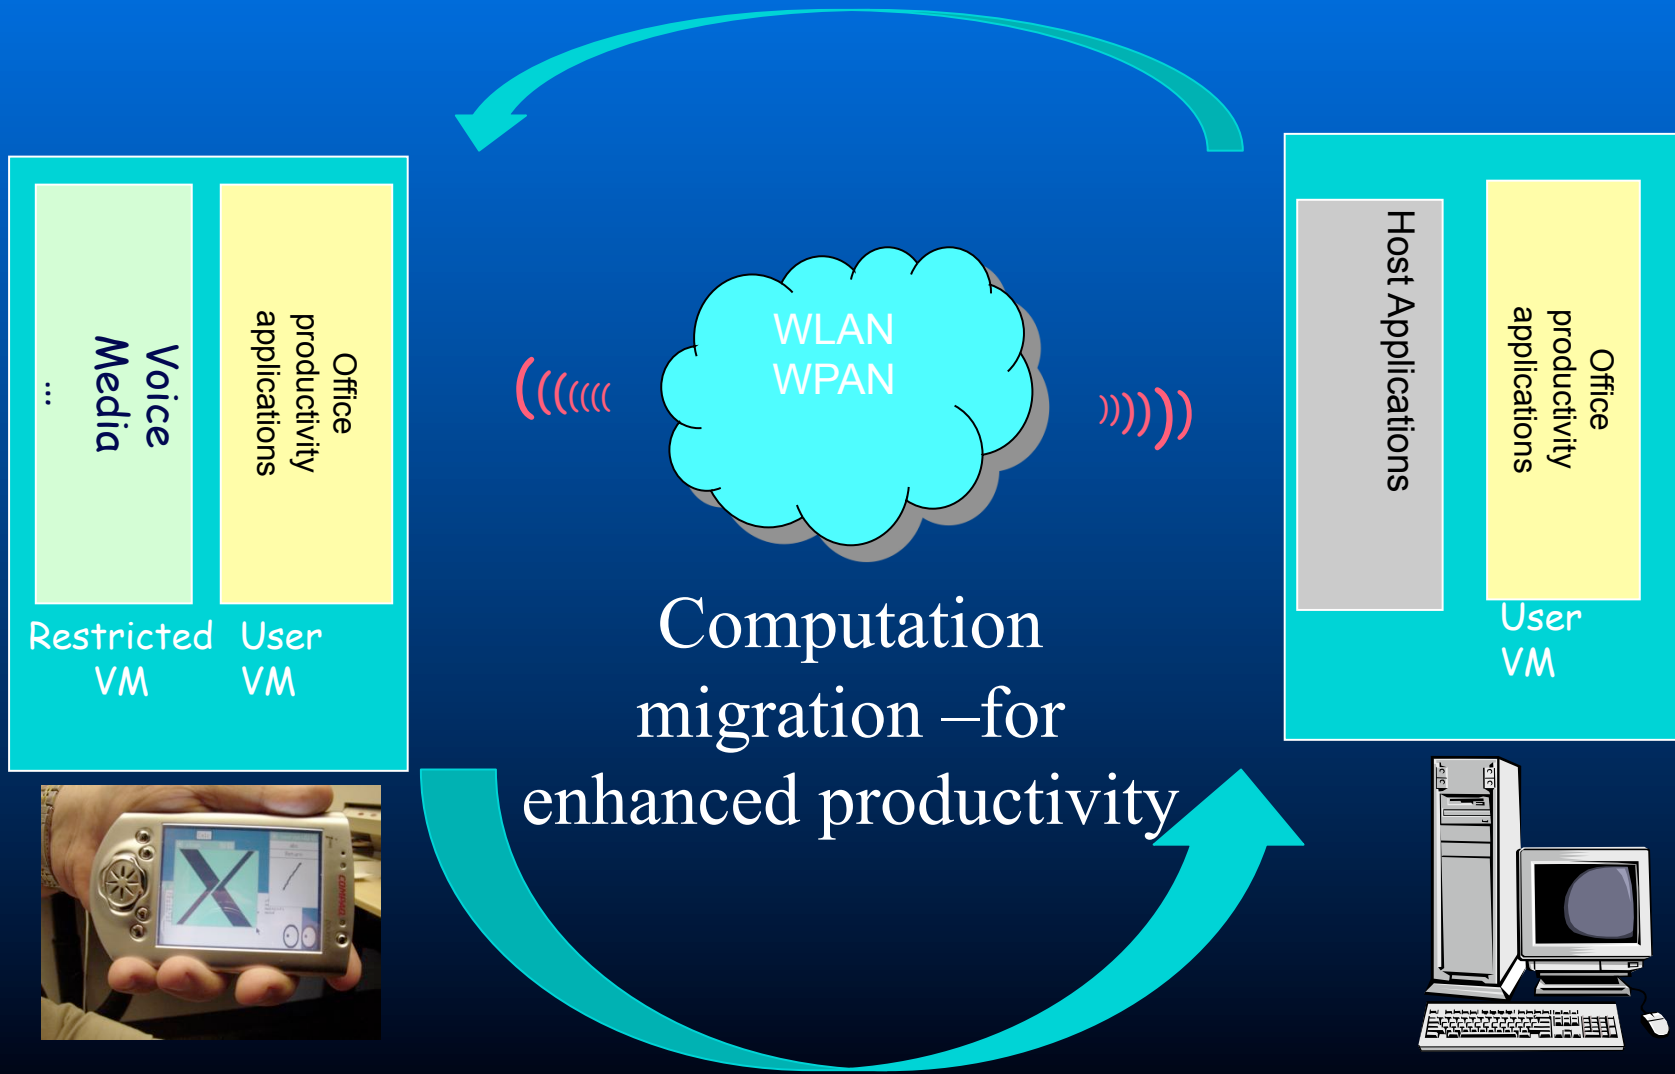
- But how do you make legacy applications use remote execution?
- And get good performance as well?



Strawman Mechanism

- Heavily modify each application to use remote execution
 - Tweak every last drop of performance
- Requires ~3- 4 grad student months per reasonably sized application
 - Grad students have nothing else to do anyway right?? ☺
- Method does not scale and is not agile


Mechanism: CloudClone




CloneCloud (Chun et al.)

- Developers write applications once for their Smartphone platforms
 - CloneCloud morphs applications automatically
 - Clones software of the Smartphone
 - Synchronizes image incrementally and offloads execution in clone
 - Merges results back retroactively or not

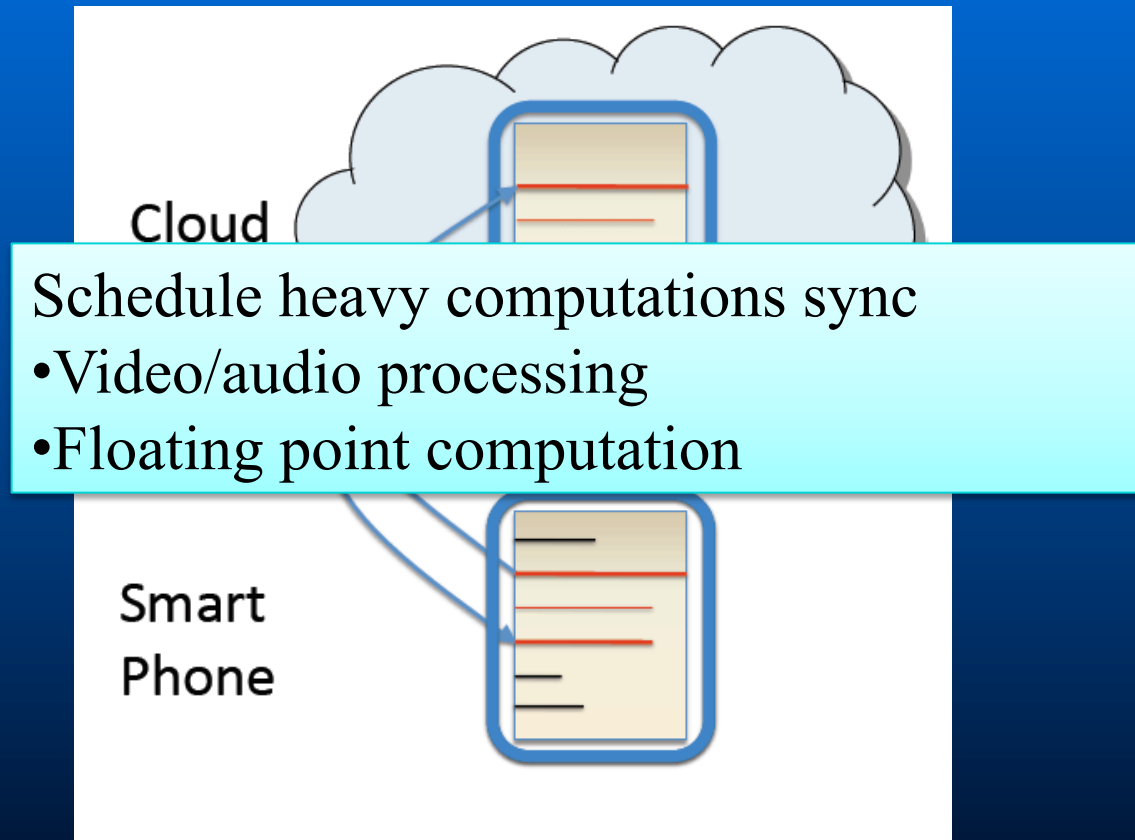
Augmentation types




Primary functionality outsourcing



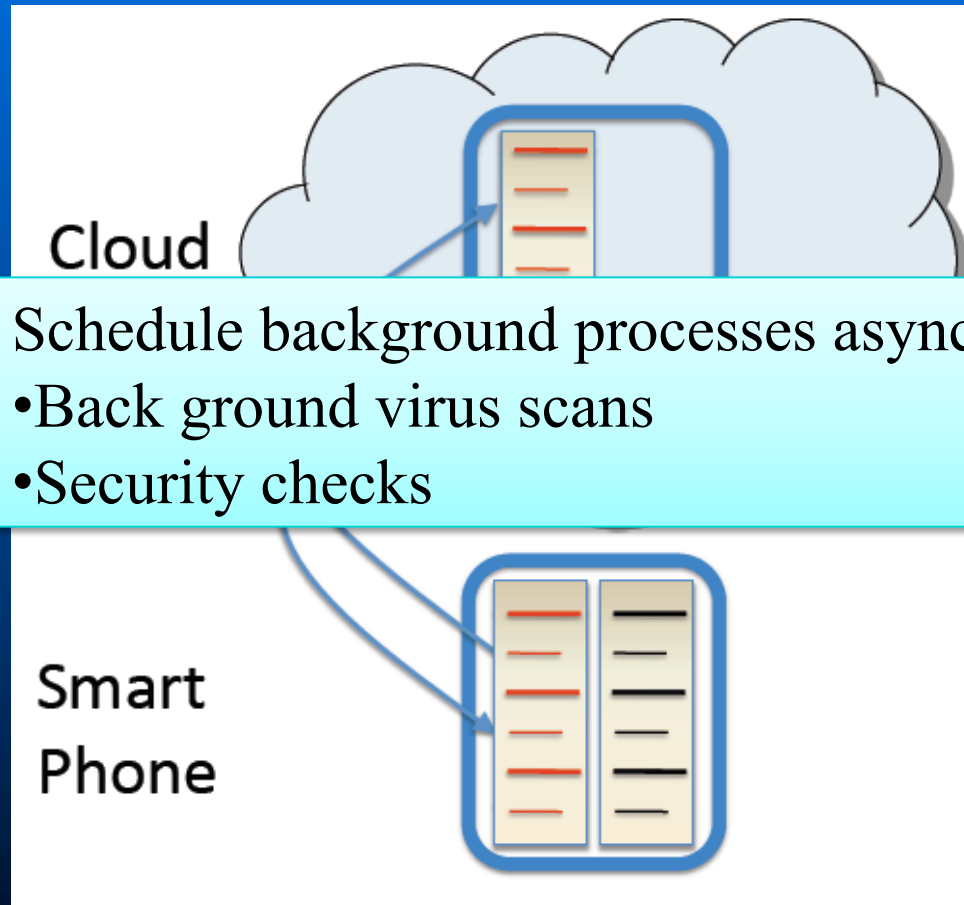
Primary functionality outsourcing




Background augmentation



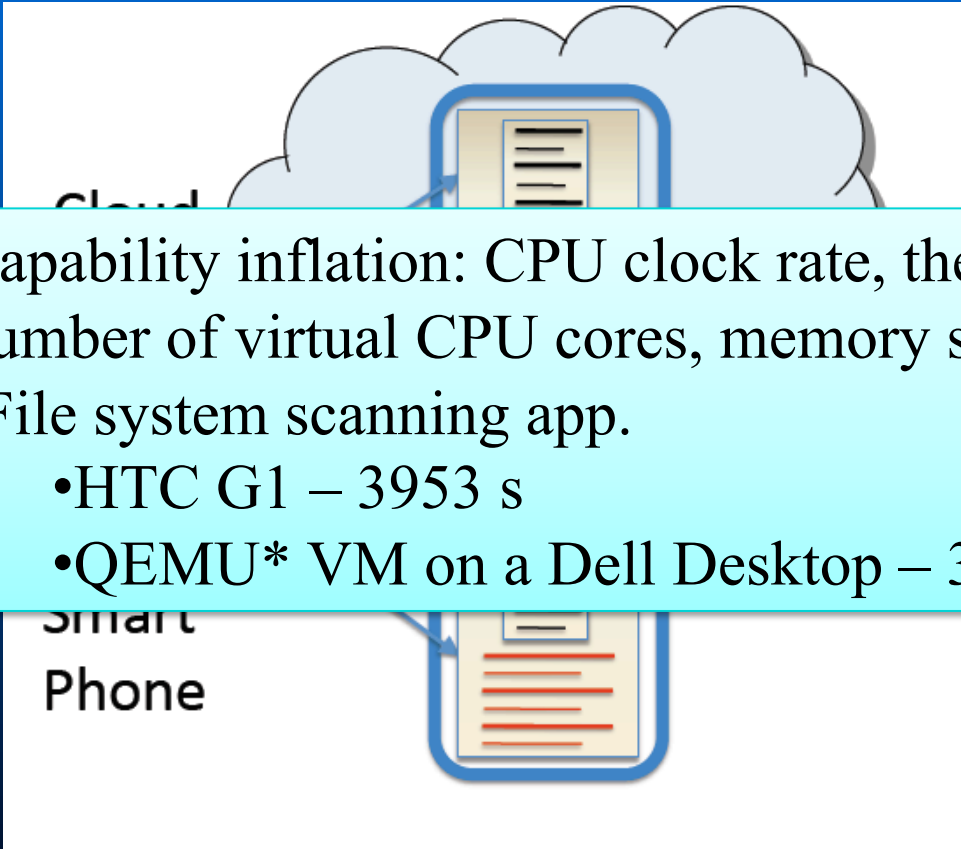
Background augmentation



Hardware augmentation



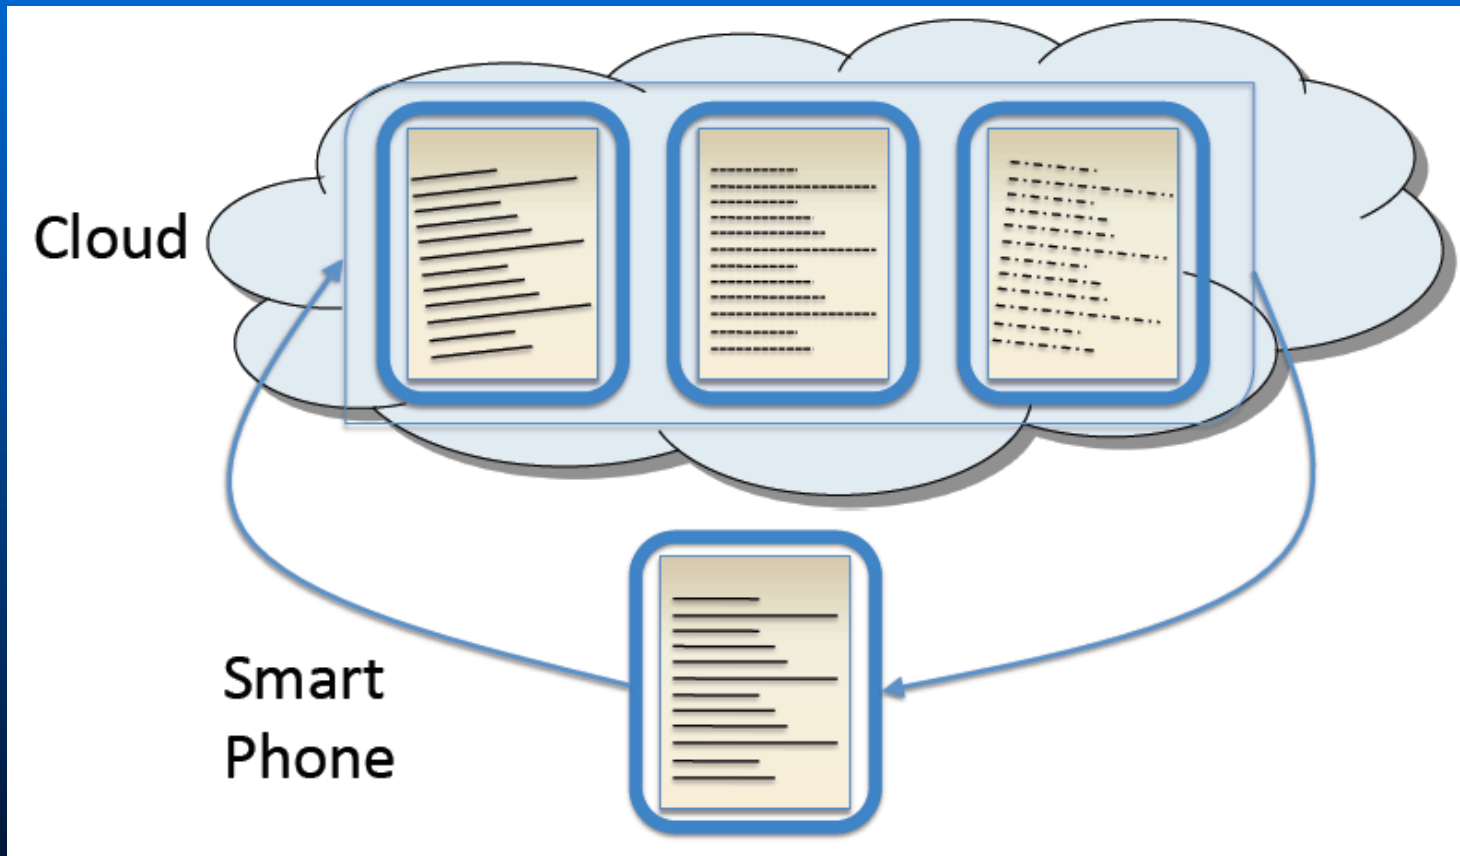
Hardware augmentation



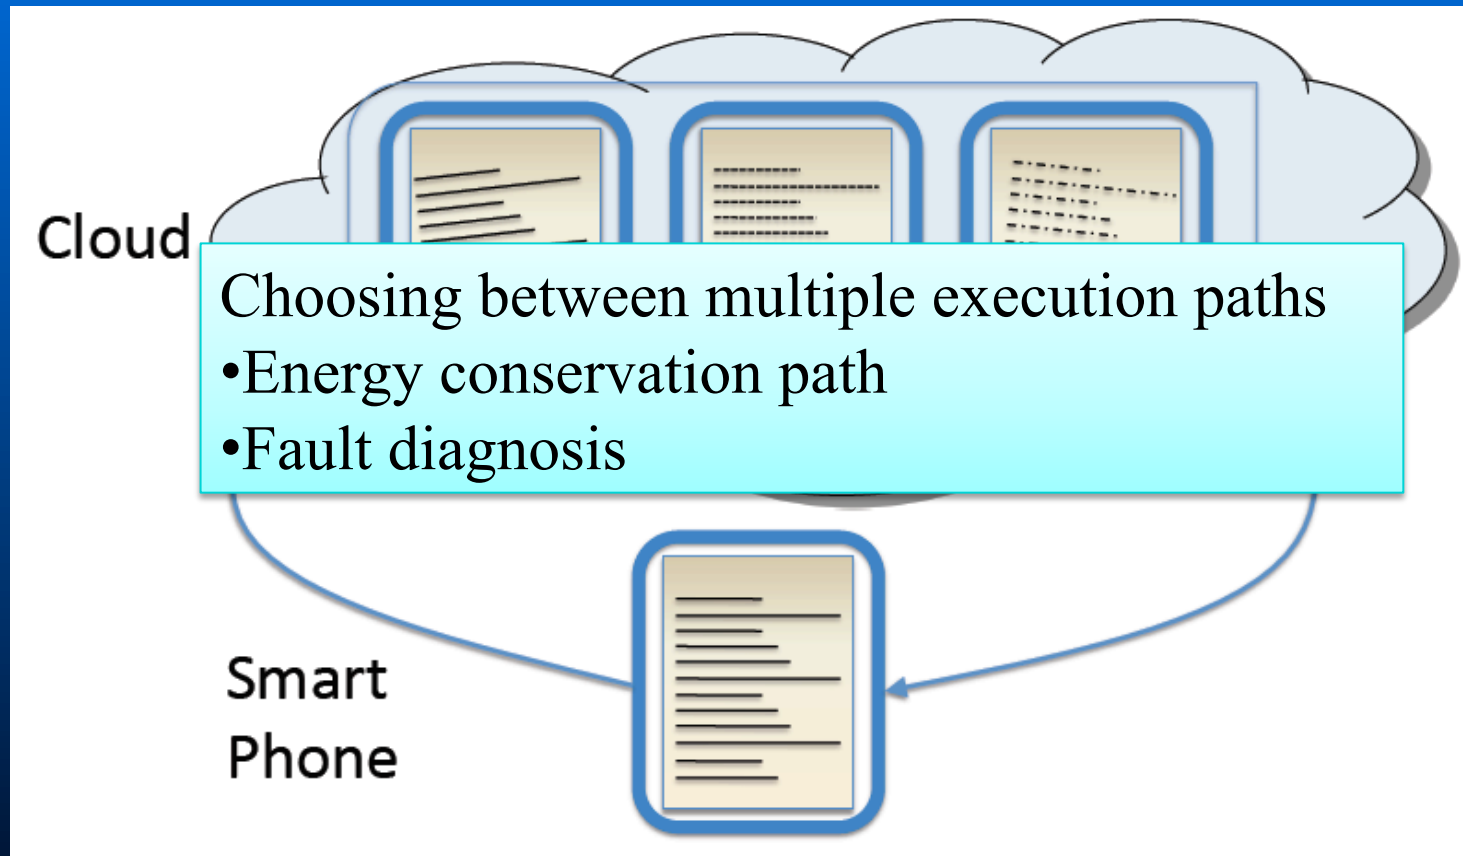
Capability inflation: CPU clock rate, the number of virtual CPU cores, memory size

- File system scanning app.
 - HTC G1 – 3953 s
 - QEMU* VM on a Dell Desktop – 336 s

Augmentation through multiplicity



Augmentation through multiplicity



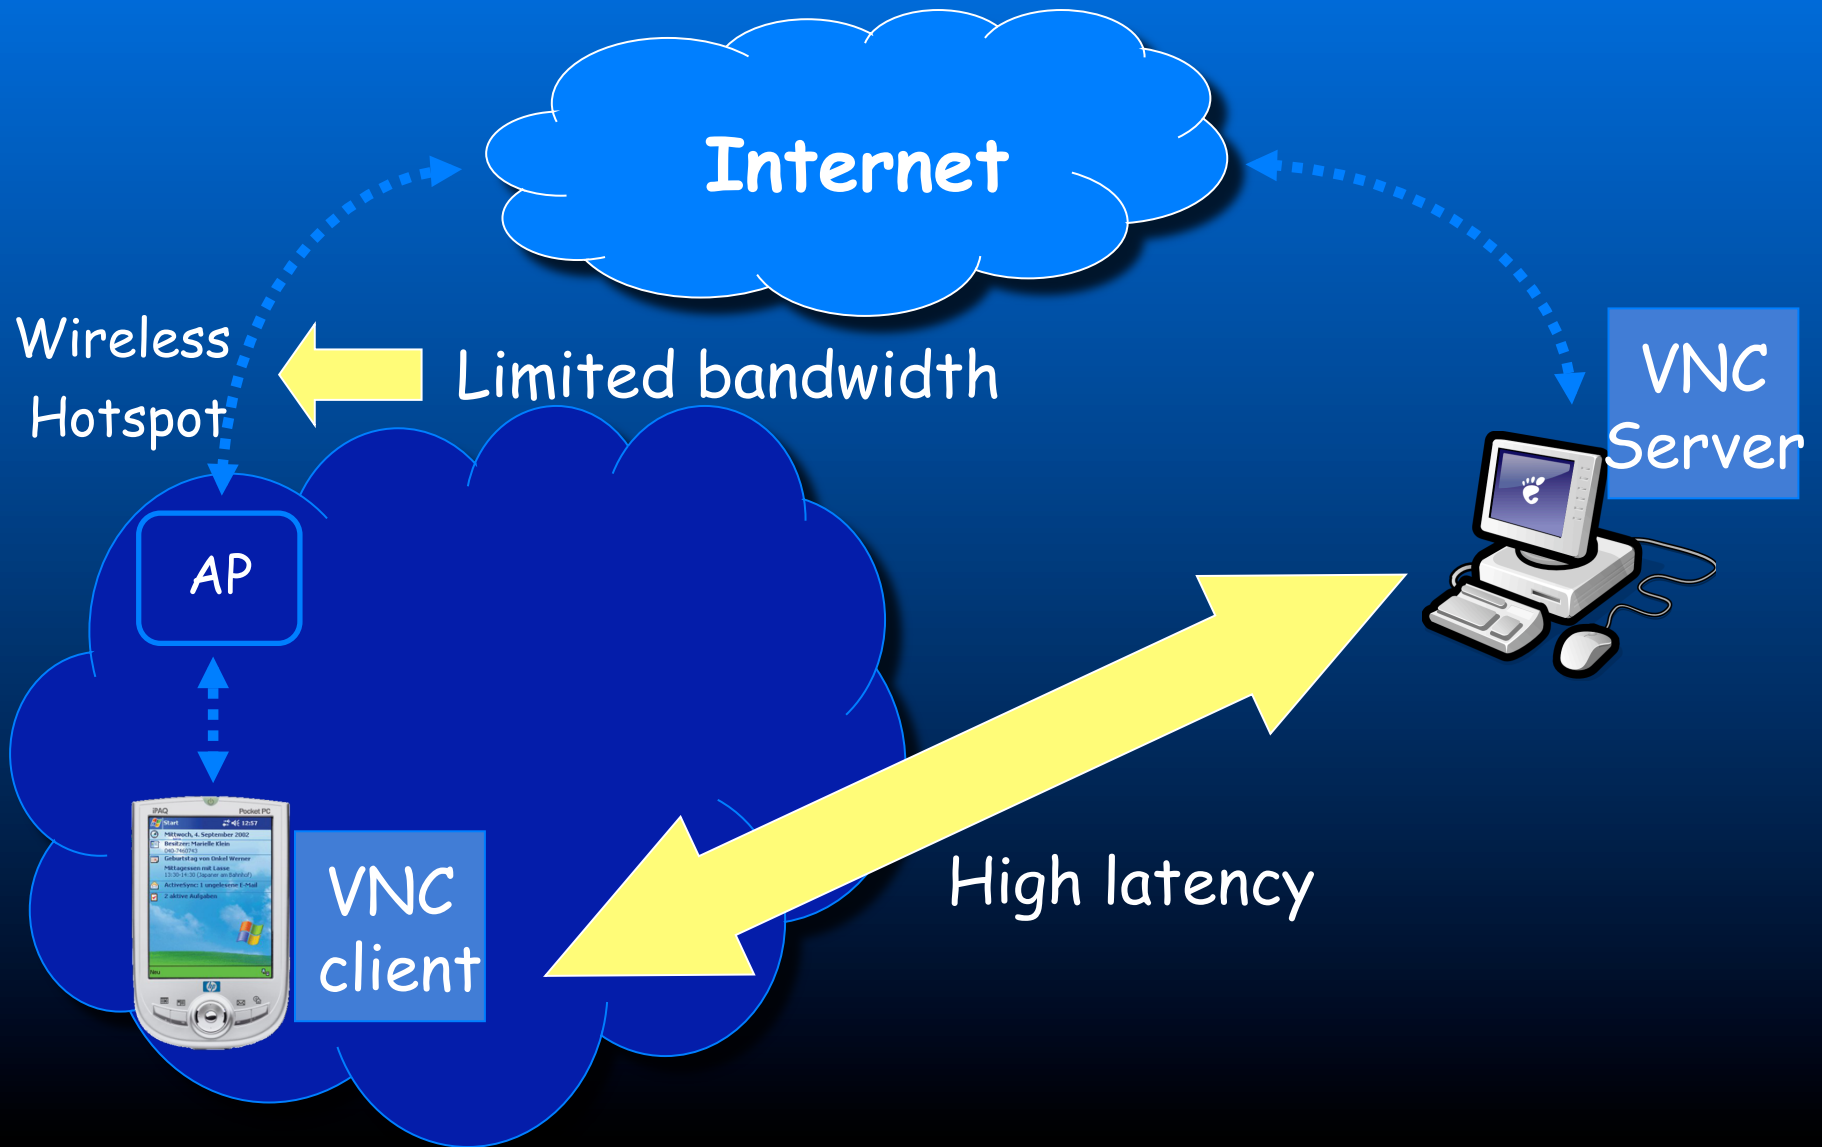
Cloudclone: Key challenges

- When to augment?
 - Applicability is application dependent
- Which part to augment?
 - Consider local computation power
 - Resource usage such as power
- Where to augment?
 - Choose remote location intelligently
 - Network latency

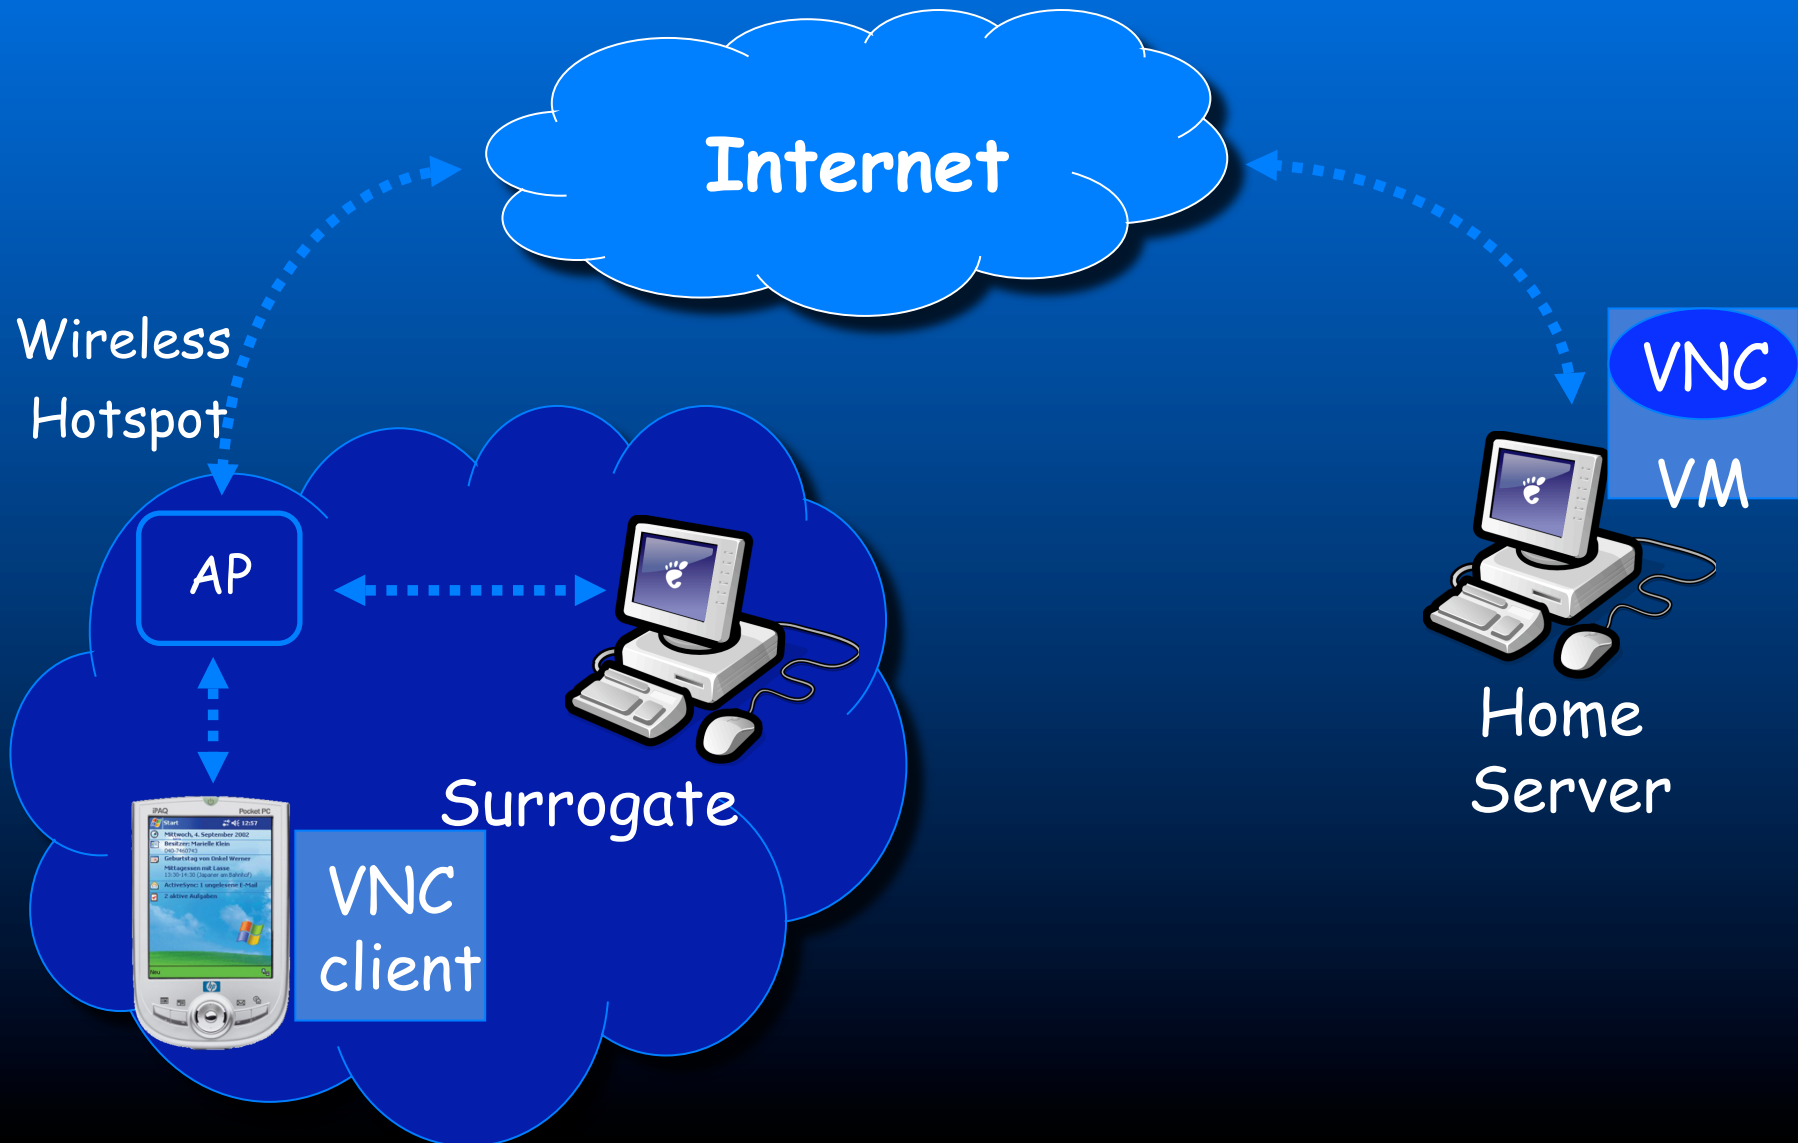
Cloudclone: Key challenges

- When to augment?
 - Applicability is application dependent
- Which part to augment?
 - Consider local computation power
 - Resource usage such as power
- Where to augment?
 - Choose remote location intelligently
 - Network latency

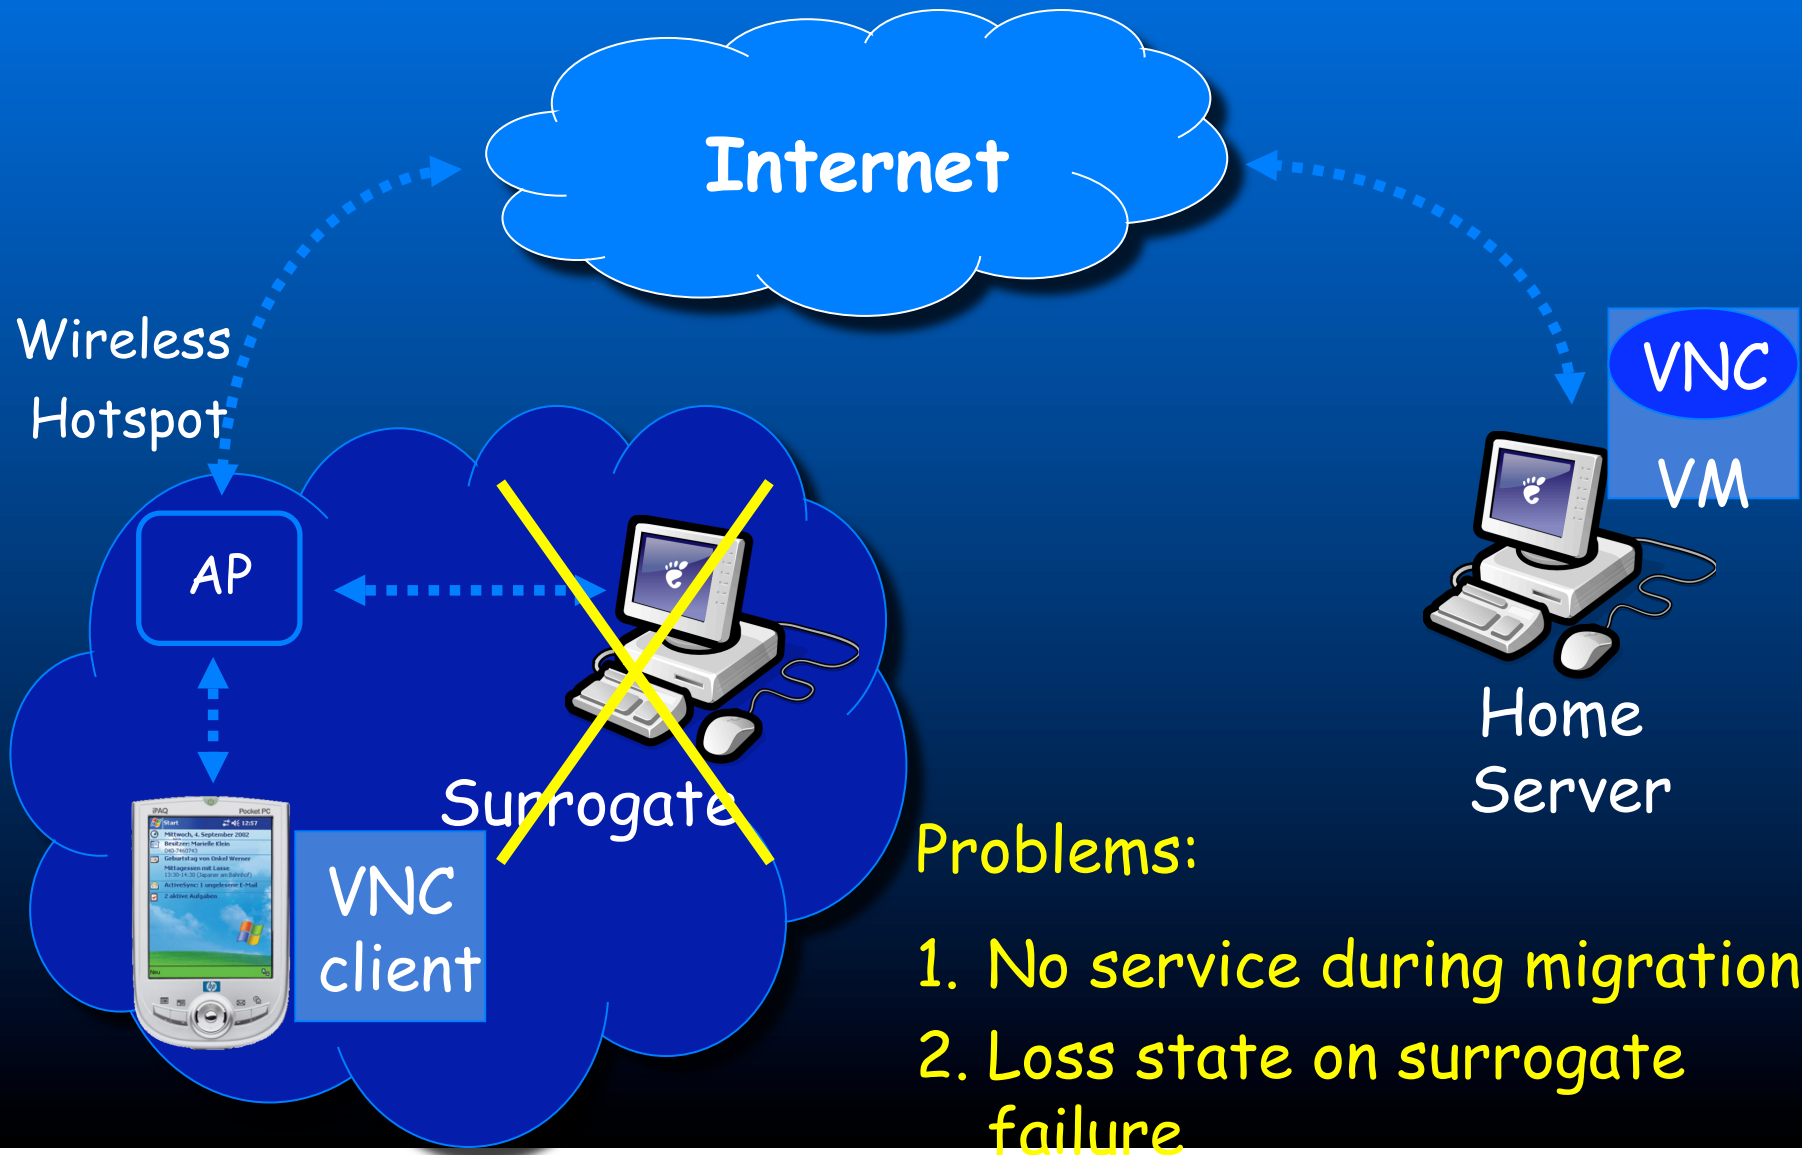
Remote execution



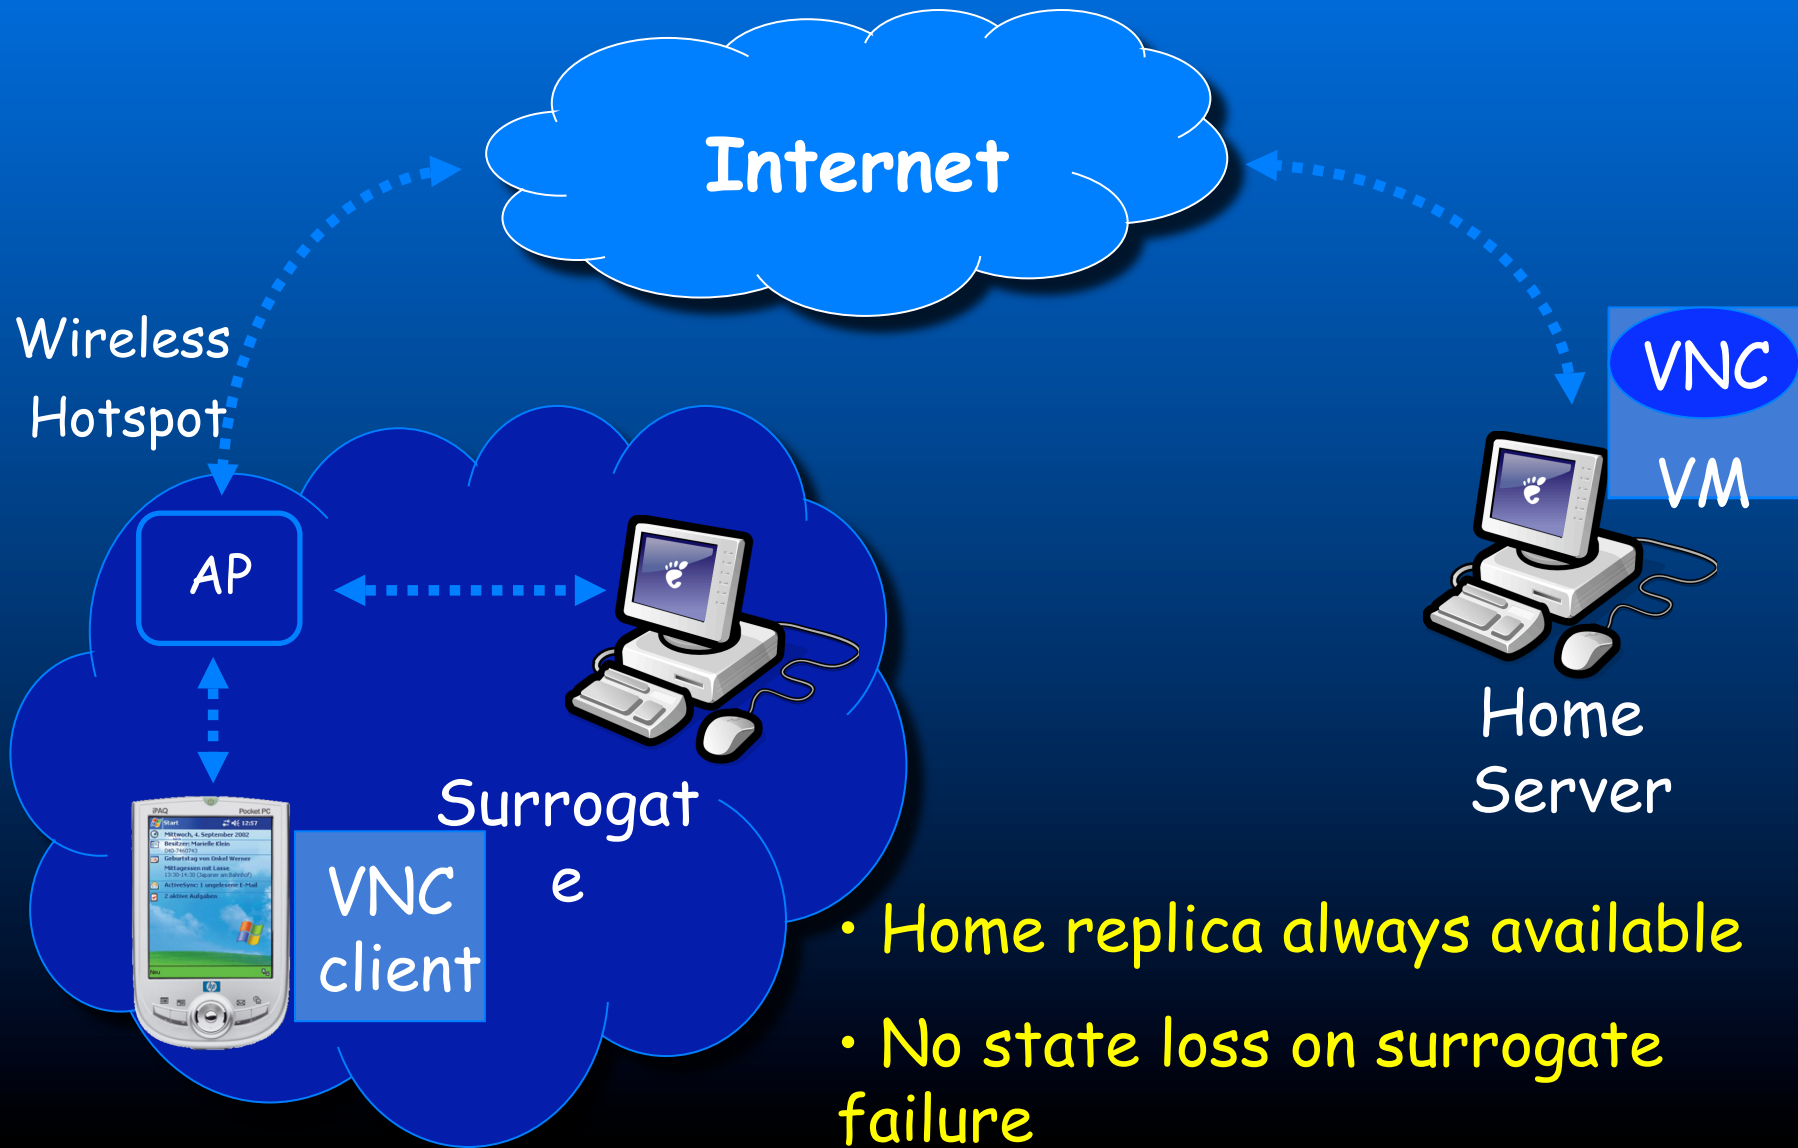
Slingshot: Smart Remote execution



Slingshot: Smart Remote Execution



Slingshot: Smart Remote Execution



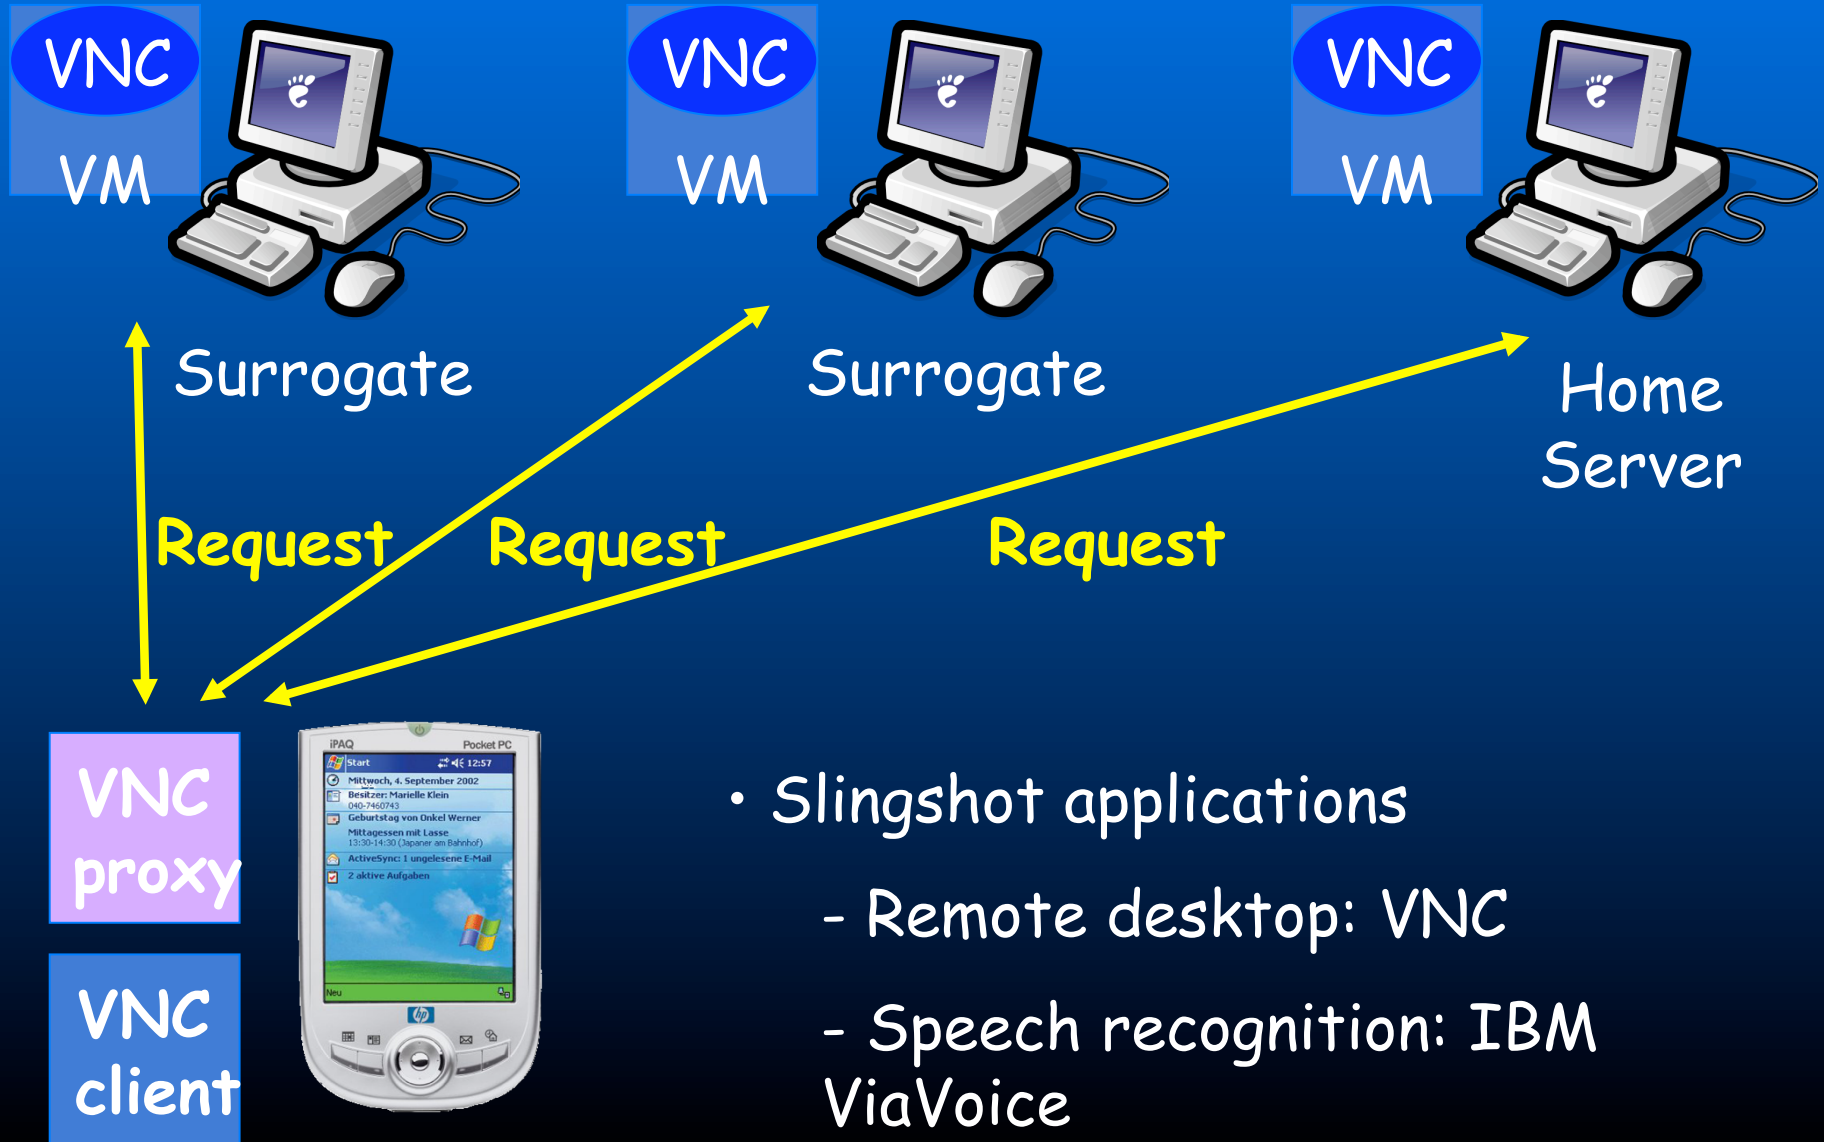
- Home replica always available
- No state loss on surrogate failure

Ease of Management

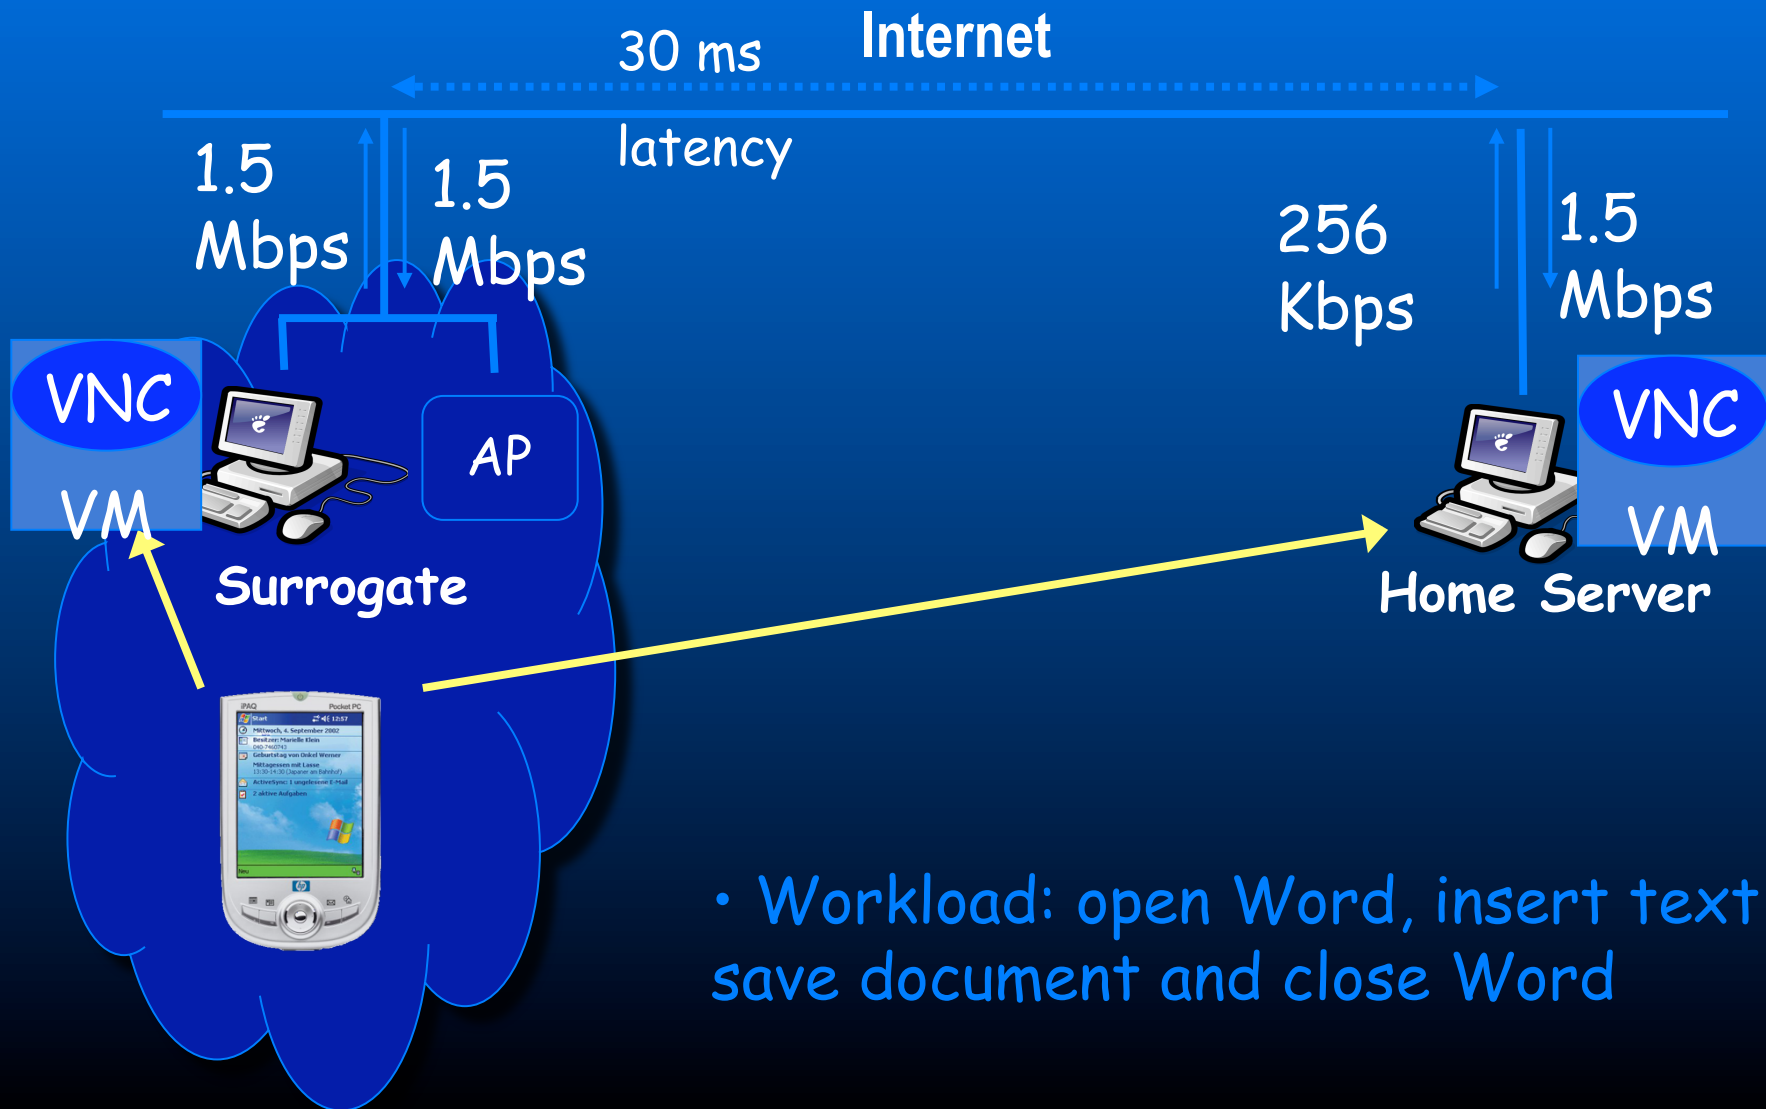
■ Slingshot

- Minimizes the surrogate computing base
- Uses a heavyweight virtual machine
- Places no hard state on surrogates

Slingshot Overview




Network Topology




- Workload: open Word, insert text, save document and close Word

Benefit of Slingshot

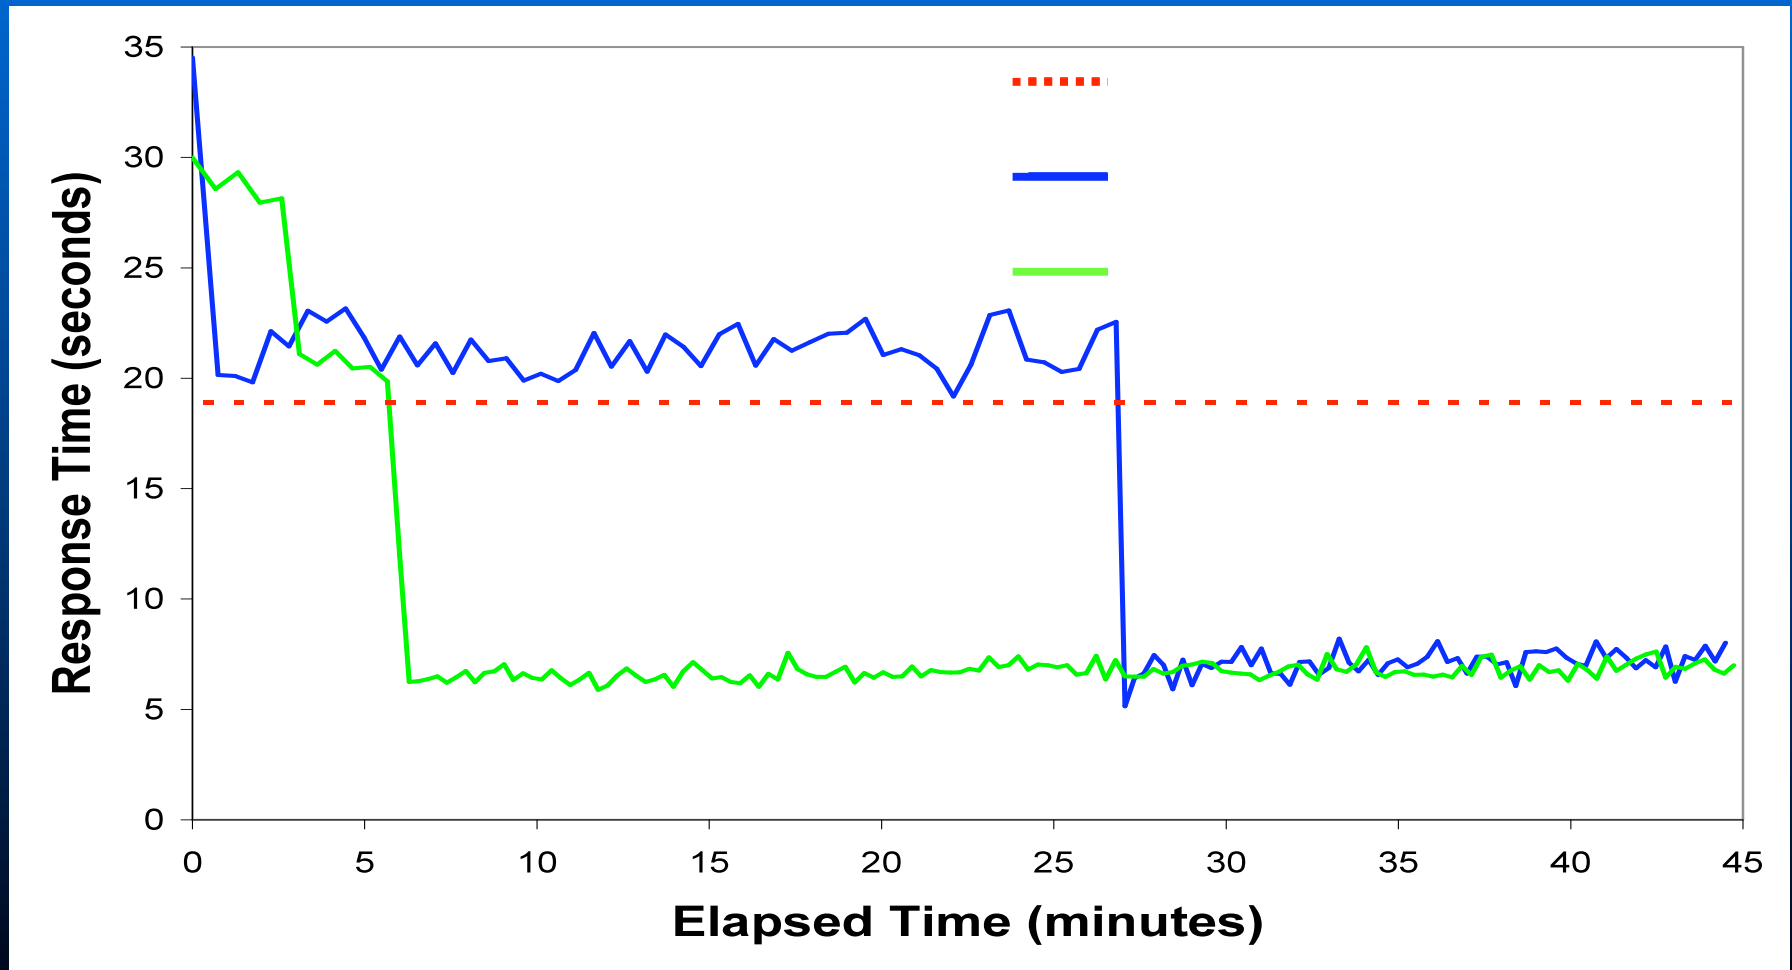


- Slingshot: 2.6 times faster than remote execution

Network Topology




Instantiating the First Replica



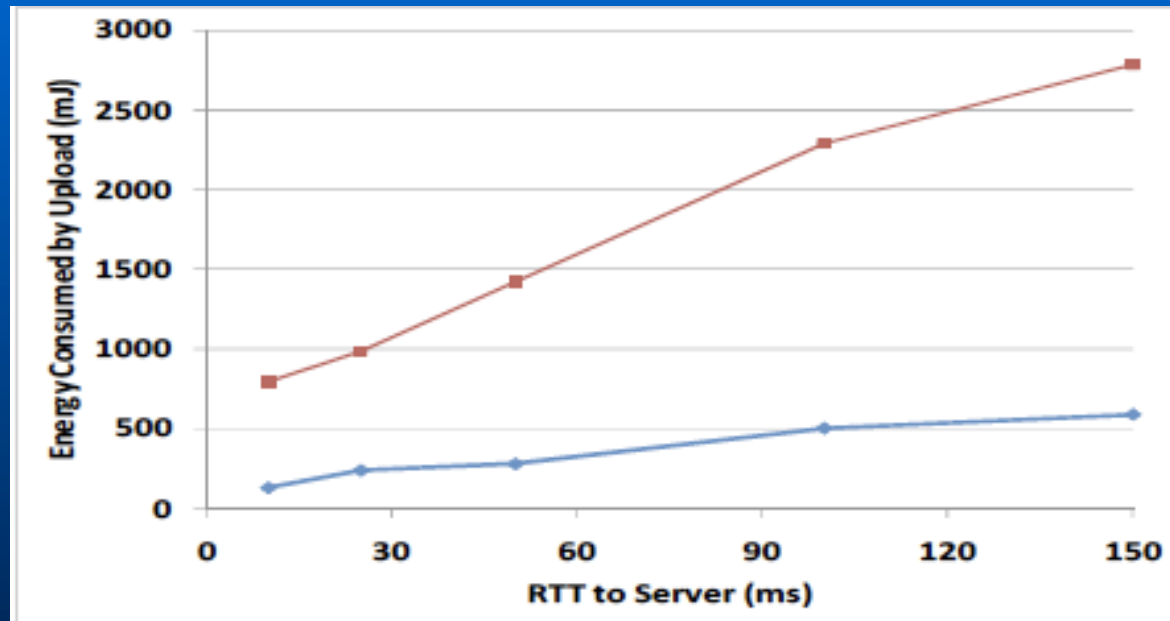
- Slingshot executes 2.6 times faster than remote execution

How about energy savings ?



The Energy Consumption of Wi-Fi vs. 3G for
Connectivity to the Cloud

Energy consumption is function of RTT

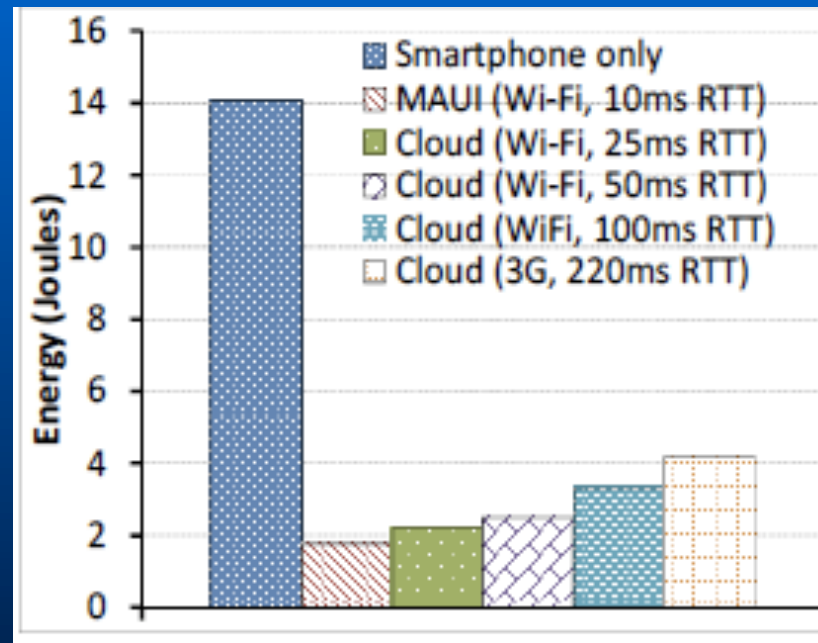


The Energy Consumed by Offloading 10KB and 100KB of Code to the Cloud as the RTT Increases

How do we conserve energy then ?


- Use close-by surrogates for offloading computation
- Choose between modes of transfer (3G vs WiFi)
 - Profile WiFi/3G/DSL links
 - Choose the one with the best RTT
 - See work like “Context for wireless”

Energy savings from smart offloading



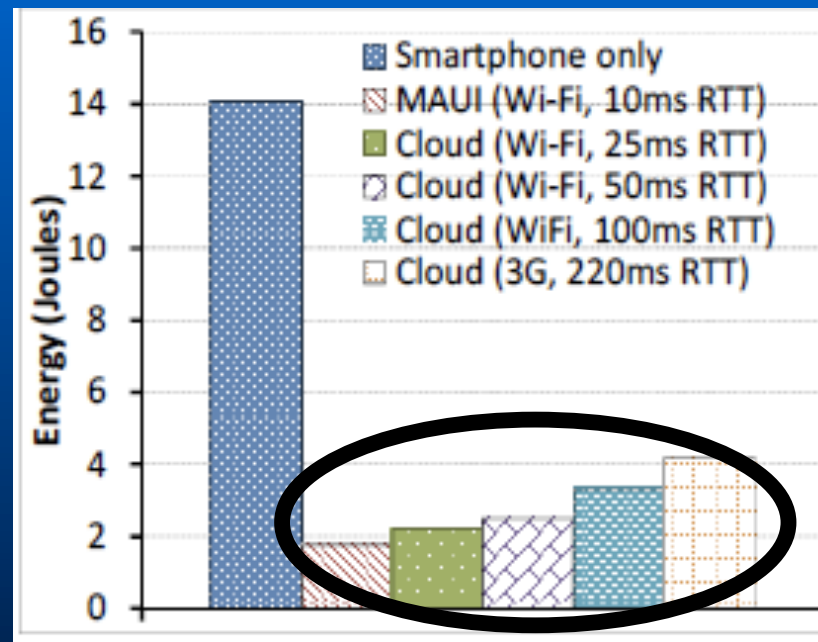
A comparison of energy consumption for image recognition software

Energy savings from smart offloading



A comparison of energy consumption for image recognition software

Energy savings from smart offloading



A comparison of energy consumption for image recognition software

Conclusions

- Mobile cloud computing extends limits of mobile devices
- Presented some key mechanisms to facilitate such mobile computing in the cloud
 - Cloudclone
 - Slingshot