Dynamic execution context management in heterogeneous computing environments, towards persistent computing environments

**NOKIA** Connecting People

Serguei Boldyrev Senior Manager PhD, MBA

Location & Commerce BU Common Technologies & Architecture

NVRAMOS 2011, November 9th

©2011 Nokia

# The nature of the consumer's relationships with companies is changing



From a monologue...



to a conversation...



into continuous relationships



From a unified...

Rational Restored Former Restored Form



to a segmented...

Into a dynamic personalized offering together with ecosystem

#### Content and Consumer Lifetime Value are MAIN DRIVERS





# **Enhancing user experience**



#### **Towards inclusive and sustainable ecosystem**

- Nokia targets 300M active Nokia Services users by the end of 2011
- 800M Cloud-compatible devices exist in <u>2012</u>
- Persistent computing, energy and cost effective HW is on the vendor's roadmaps, 2013 and onwards
  - ARM big.LITTLE, GPUs, SSDs NAND, PCM, MRAM
- Cloud Computing subscribers would total nearly one billion by 2014



#### **Outline**

- From Web of Information
- To Web of Computing
- Explore & Share
- Persistent Computing
- Opportunities & Challenges
- Data
- Computation
- Managed Performance
- Privacy
- Conclusions



# From a Web of Information

- Cloud computing today serve email, apps, downloads and storage
  - Current "cloud apps" relying more on server-side processing
  - SaaS, PaaS, NaaS and other service models to emerge
  - Thin client, network as a computer, etc
  - Seamless information management between Client and Cloud



# **Towards a Web of Computing**





#### Why is this happening? (we agree with "Others" (19)



Because digital transformation of products and services are demanded by consumers, with increasing **complexity driven by** new customer **expectations** 



NDKIA

### **Beginning: Explore & Share**

Deliver services and content



Ovi at Retail concept (see in YouTube)

# Share content between devices



#### **Device to Device**

#### Explore and Share new content in tags



#### **RF** memory tags



#### **RF Memory tag technology**



Applications: wireless memory in batteryless objects, ambient intelligence => **Universal Local Storage** 





#### **Positioning of wireless technologies**



### **Explore & Share, architecture**





MEMORY

SUBSYSTEM

### **Explore & Share, functional blocks in**



©2011 Nokia

### **Explore & Share, Computation migration, sort of**













#### **Ideal Scenario for the Future Solution**

- Mobile device is a true contributing participant , not just a depending client
- Mobile device connects to multiple data sources (possibly independent devices or clouds), forming one logical Computation/Data and Service space
- From any connected device, the user deals with the same data regardless of its physical location

- Solution provides both data and distributed computing services
- Solution access is managed by fine-grained security policies applied according to the user's context
- From the design perspective, the Solution manages and hides complexity; does not increase it
- Solution data and services decrease development and maintenance costs
- Solution data and service pricing model creates new money flows between players



#### **New Digital Experience**

"My digital experience follows me regardless the computing environments around me"

- take advantage of finely granular and accountable processing services,
- integrate heterogeneous information sources, and
- ultimately free users of mundane challenges and details of technology usage

#### <u>Greater reuse of information and computational tasks is just few steps</u> <u>away</u>

# **Opportunity of New Digital Experience**

World of fully deployed and available Computing – Persistent Platforms

The emergence of Platforms that can <u>seamlessly span the</u> <u>"information spaces"</u> of multiple

- hardware,
- software,
- information and
- infrastructure providers





#### **Technological enablers for Persistent Computing**

#### Utilizing the best parts of the Web development paradigm

- ...Web-based applications, not Web sites...
- and leveraging the special contextual capabilities of mobile devices (ex. WURFL databases on sensors)

#### Moving from a Device, Cloud and Web-centric models to a Persistent model (Persistent computing)

- Reuse (partition) information and persistent computation in the Cloud, Infrastructure and Devices
- Elimination of special-purpose software to download, install applications, Service oriented infrastructure constructed as a functional flows requested on demand
- Enabling balancing of computation and relevant data between heterogeneous Cloud Back-Ends, Infrastructures (ex. Spanner, HDFS, Web Intents etc) and devices

#### Fostering faster, easier, richer application innovation and deployment through SW and HW Persistent computation recycling

- AMP processors, non-volatile execution memory
- Diverse Cloud Back-Ends can all leverage the infrastructure capabilities
- Atomic computations are now deployed once to the Back-Ends and composed down to Infrastructure and clients, where a set of functional flows or description of those form the actual service
- Allowing Network Infrastructure to leverage data and computation workload, by taking care of services or provider of services capabilities, distributing computation between Back-Ends and Infrastructure
- Allowing more efficient contextual composition of services than purely device or Web-centric models



NDKI

#### **Persistent Computations**

- A next generation distributed systems are dynamically spanned around virtual and physical infrastructures
- Device becomes a known and contributing node of the cloud infrastructure, service provisioning element

Examples: Services based on crowdsourced information on traffic, pollution, weather, body temperature, context intelligence, environments adapting to user context, efficient computing





### **Persistent computing, Requirements**

The new emerging paradigm of Distributed Systems requires an explicit orchestration of computation across

- Mobile Devices (D2D connectivity)
- the Edge, (thin elements of the Cloud "skin")
- Cloud and
- the Core

Orchestration depends on

- components' states (persistency!) and
- resource demand,

through the proper exploitation of granular and therefore accountable mechanisms





#### Design framework should be re-defined!

#### Persistence is the key attribute of ALL future Data & Computation Management systems





# **Opportunities and Challenges**

Main aspects of

- Data and Computation semantics,
- Performance and Scalability of computation and networking,
- Role of Security and Privacy



### **Data - Common Data Model (CDM) enables Shared data across application domains**



Client Applications should not "own" data, instead, they should reflect the users' desire to *do something*, to *accomplish certain goals* 

- to call someone, you need **Contacts**,
- to know what to do next, you need Calendar, etc.

Data is still the same...

- for example, it would be helpful for Contacts to know who is the organizer of your next meeting
- EXAMPLE: Apple Siri, Google Chrome



## Shared data, layers view



#### Shared data demands

- steep learning curve (new ways to think about apps, platform, UX, ...)
- represents a shift from "device-centric" and "Cloud" thinking to "Future" thinking
- risks, uncertainties, … (also big potential reward)



### **Computation and Data perspective**

#### Prerequisites

<u>Core Data assets are encapsulated into results of Computation tasks</u>

**Objectives** (but not limited to listed below)

- Design and implementation of scenarios where computation can be described and represented in a system-wide understandable way
  - enabling computation to be "migrated" (transmitted) from one computational environment to another
  - such transmission would involve reconstruction of the computational task at the receiving end to match the defined semantics of task and data involved

• Construction of a system where larger computational tasks can be decomposed into smaller tasks or units of computation

 independent of what the eventual execution environment(s) of these tasks (that is, independent of hardware platforms or operating systems)



# **Computation, the components**

#### Computation closure –

- current state of execution and context
- automatically generated during the development phase
- composed out of anonymous function object

**Function object** – consists of:

- class template result\_of,
- function template mem\_fn,
- function object binders, and
- polymorphic function wrapper function



# **Computations, zoom in**



Computation Closures set are backed by a Recycling or Marshaling Service, with a load-balancing facility that serializes Closures to be executed on remote computing environments



Treating Closures as Data entities, letting us compose Closures, chain them, and project computations into them





# **Computations, Reflective Run-time environment**



Distributed information management framework enables reflective process/context migration

- Marshalling/UnMarshalling
- Alignment and enhancement with distribution

Distributed Computation and Data model allows dynamically balanced load between concurrent execution environments taking into account the user's current context



#### **Persistent Computations**





# **Construction of the computational ontology**



### **Developer Experience – integration with SDK**

Closures allow to compose function expression that can be wrapped and exposed to Developer as Data Closures can be passed to API, stored, shared or spread the computation

- The best estimated result based on the first returning closure can be seen as a part of functional flow
- When closures "return", results can <sup>cc FL</sup><sub>cc</sub> be refined (parallelism)
- Computation recycling is possible
- Builds up a cache of results/partial results ("fast computation")





#### Each branch represents Computation associated with Data, where Computational model is encapsulated into Data model



#### **Performance/Scalability, Requirements**

- Granular latency control for diverse Data and Computational load in heterogeneous environments
- Resolution of short term capacity decisions, including
  - the determination of optimal configuration,
  - the number of (live) Persistent nodes or
  - the migration of computation or data
- Resolution of long term capacity decisions
  - decisions concerning the development and
  - extension of data- and service architecture, choice of technology, etc
- Control of SLAs at different levels,
  - e.g. consumer services, platform components, etc
- Quality assessment of distributed software architecture

#### Test setup, SSD based

- We are running tests on MLC SSDs for the production Common Data Platform
- Initial results from performance testing
  - very promising ☺,
  - next wear testing the drives
- Performance test from 3 node cluster with 5-6 MLC SSDs on each
- 3 node cluster to over 8000TPS with constant sub 20ms latency
- safe load to achieve ~5ms
- 90% latency is around 6000TPS
- Max TPS is way higher but latency grows to 10-20ms with it (can be adjusted by playing with concurrency)



#### **Next steps PCM/MRAM evaluation**

Memory blocks modules could be packed in

- 1. xGB PCIx board with Memory blocks in single server (small size frontside cache)?
- 2. Mountable as a block device (due to SW reasons), SATA or PCI etc.?

So far, SSDs sounded like a perfect fit for reducing our latencies and cost for the Data Endpoints, but we could improve it even further and enable certain number of new business case





#### **Motivation**

#### Why we are doing ALL this?



#### **User, Trust, Privacy the placement and scope**

#### **Content and Consumer Lifetime Value are MAIN DRIVERS**



#### (note that in this diagram, by "ontology" the semantics is meant)



# **Conclusion - Driving agenda**

- Creation of **approaches** and effective **tools** for the next developer and consumer applications, tackling
  - **Fuse & Lift of Big Data** (Bring algorithm to the data, Knowledge creation)
  - Social graph oriented patterns (social media, chat, calendar, etc)
  - Tele immersive environment (AR, 3D maps, etc)
- Persistent Computing **platform** for the Continues Operations
  - the **fine-grained** (and thus easily accountable) software **frameworks** developments
  - o anything as a persistent service, e.g. sometimes giving hardware for "free"
- Mobile devices supported by Backend infrastructure technologies to enable qualitative leap in ICT industry offerings, from the perspectives of
  - o scalable technology and
  - o **multisided** business models with high elasticity

#### <u>Aiming and supporting future strategic growth areas and maintenance</u> <u>requirements</u>



# Q&A [](){} sergey.boldyrev@nokia.com



# Thank you

