



# The SC|06 VMT Demonstrator

Franco Travostino ([travos@nortel.com](mailto:travos@nortel.com))

Phil Wang ([pywang@nortel.com](mailto:pywang@nortel.com))

Inder Monga ([imonga@nortel.com](mailto:imonga@nortel.com))

Paul Daspit ([pdaspit@nortel.com](mailto:pdaspit@nortel.com))

CTO Advanced Technology

November, 2006





**General concept:** The GLORIAD VMT demo at SC|06 was the latest in a series of demonstrations aimed at showing how virtualization of resources can be made adaptable and dynamic to provide a richer experience for the user. Employing Web Services to simplify the various “mashups,” computation can be delivered at the right place and time. The SC|06 demonstration added two new dimensions: (1) Sensor-triggered initiation of the VM migration and (2) VM clones – send one, receive many.

**Demonstration network topology:** VMT host nodes were located at Amsterdam, Daejeon (Korea), Chicago, and Tampa. GLORIAD supported 2xGbE connections between Amsterdam, Daejeon, Chicago, and New York, respective for data and control. Atlantic Wave and Florida LambdaRail provided 2xGbE connectivity between New York and Tampa. Nortel’s Dynamic Resource Allocation Controller (DRAC) orchestrated the “just-in-time” lightpath cross-connects at switching nodes in Seattle and Chicago using UCLP and TL-1 proxy interfaces. Nortel’s sensor virtualization and event processing prototype was leveraged to monitor per-rack temperature and UPS events within the Tampa data center to generate the sensor triggers.

**Demo description:** More than 50 migration experiments were conducted during scheduled demonstration timeslots over the five day period. The VMT demo performed three types of experiments, including live VM migration (manually initiated 1:1 migrations), Sensor-triggered VM bailout (1:1 migrations) and VM clone (manually initiated or sensor-triggered 1:2 and 1:3 migrations). An image pattern search application was used as an example for demo purposes. Large data sets of images were placed at each VMT location. Image searches were performed as the VM delivered the application to each VMT location. Successful matches resulted in the image file being “tagged” for later identification. The demo network, VM application and sensors at each VMT locations were globally monitored and displayed in real-time.

**Demo results:** The results from the experiments successfully verified VMT technology effectiveness. First, a live migration from Tampa to Chicago took 3.06s to transfer the complete VM data and caused an application interruption time less than 50ms (CPU interrupt). Second, the Sensor-triggered VM bailout took 986ms to detect a sensor alarm, 10ms to initiate a VM bailout action, 1.35s to verify the lightpath from Tampa to Chicago, and 3.62s to restore the bailout VM at Chicago. Last, the VM clone sourced from Daejeon, was received at the other VMT locations in under 10s:

- 0.00s, VM clone action issued
- 3.35s, VM checkpointed at Daejeon
- 5.34s, VM cloned at Daejeon
- 6.56s, VM received at Chicago
- 6.98s, VM received at Amsterdam
- 8.82s, VM cloned at Amsterdam
- 9.41s, VM cloned at Chicago

**Summary:** The VMT experiments show “CPU + Data + Network + Sensors” multi-resource coordination in action – demonstrating that a technology once restricted to the data center is now useful in the MAN and WAN. We can steer computation resources towards data, in planet scale, unbeknownst to applications and users. Deterministic and dynamic lightpaths are the key enablers rather than sheer throughput.

# The SC|06 VMT Demonstrator



Orchestration of:

- Computation
- Storage
- Lightpaths
- Sensor events

Points to new capabilities in datacenter automation and to virtual datacenters with a span of many thousand miles by orchestrating diverse resources:

- virtual machines
- optical paths
- sensor data reflecting physical events
- all within SOA/Web Services framework.

# The SC|06 VMT Demonstrator



A representative use case:

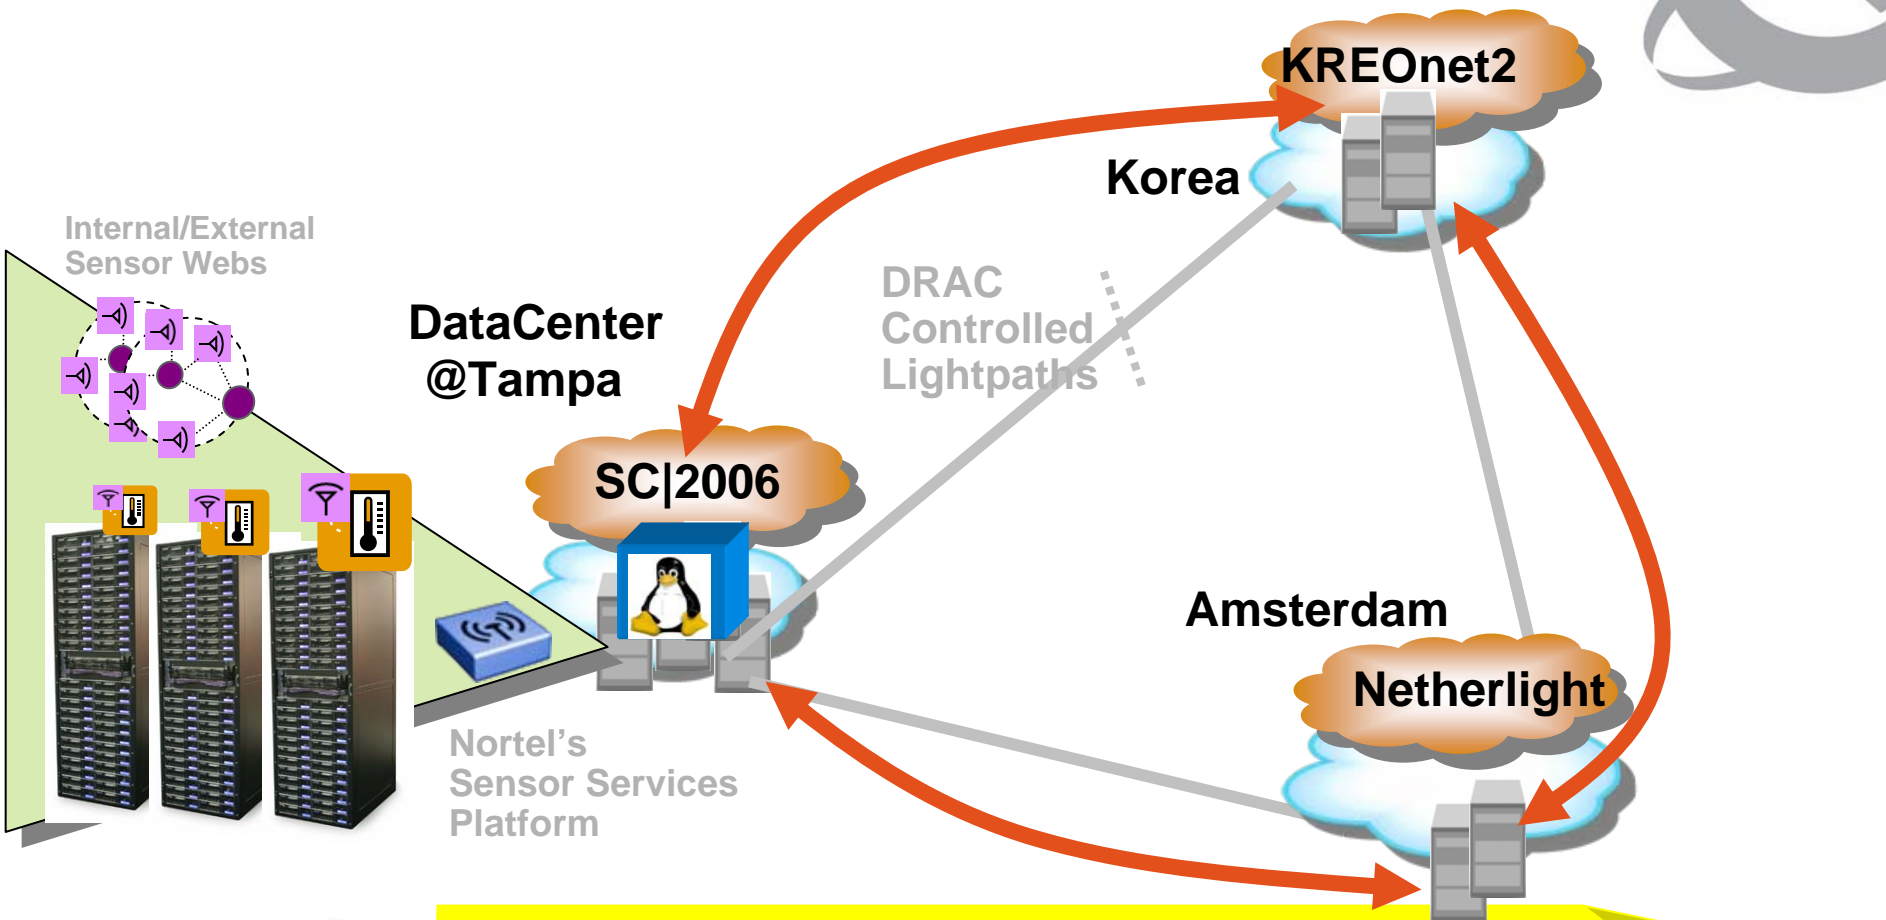
- temperature sensor warns of cooling capacity limits
- triggering the execution of a workflow migration to a data center with more favorable readings for electrical power, cooling, and compute load
- the "live" migration of computation causes the set-up of lightpaths to the distant data center for the duration of the migration workflow

Live from the Tampa SC06 floor:

- the computation is migrated over planet-scale distances
- from Tampa to Chicago, S.Korea and the Netherlands
- transparently to any running application and external client, with a mere 1 second glitch in wall-clock execution time.



# The SC06 VMT Demonstrator

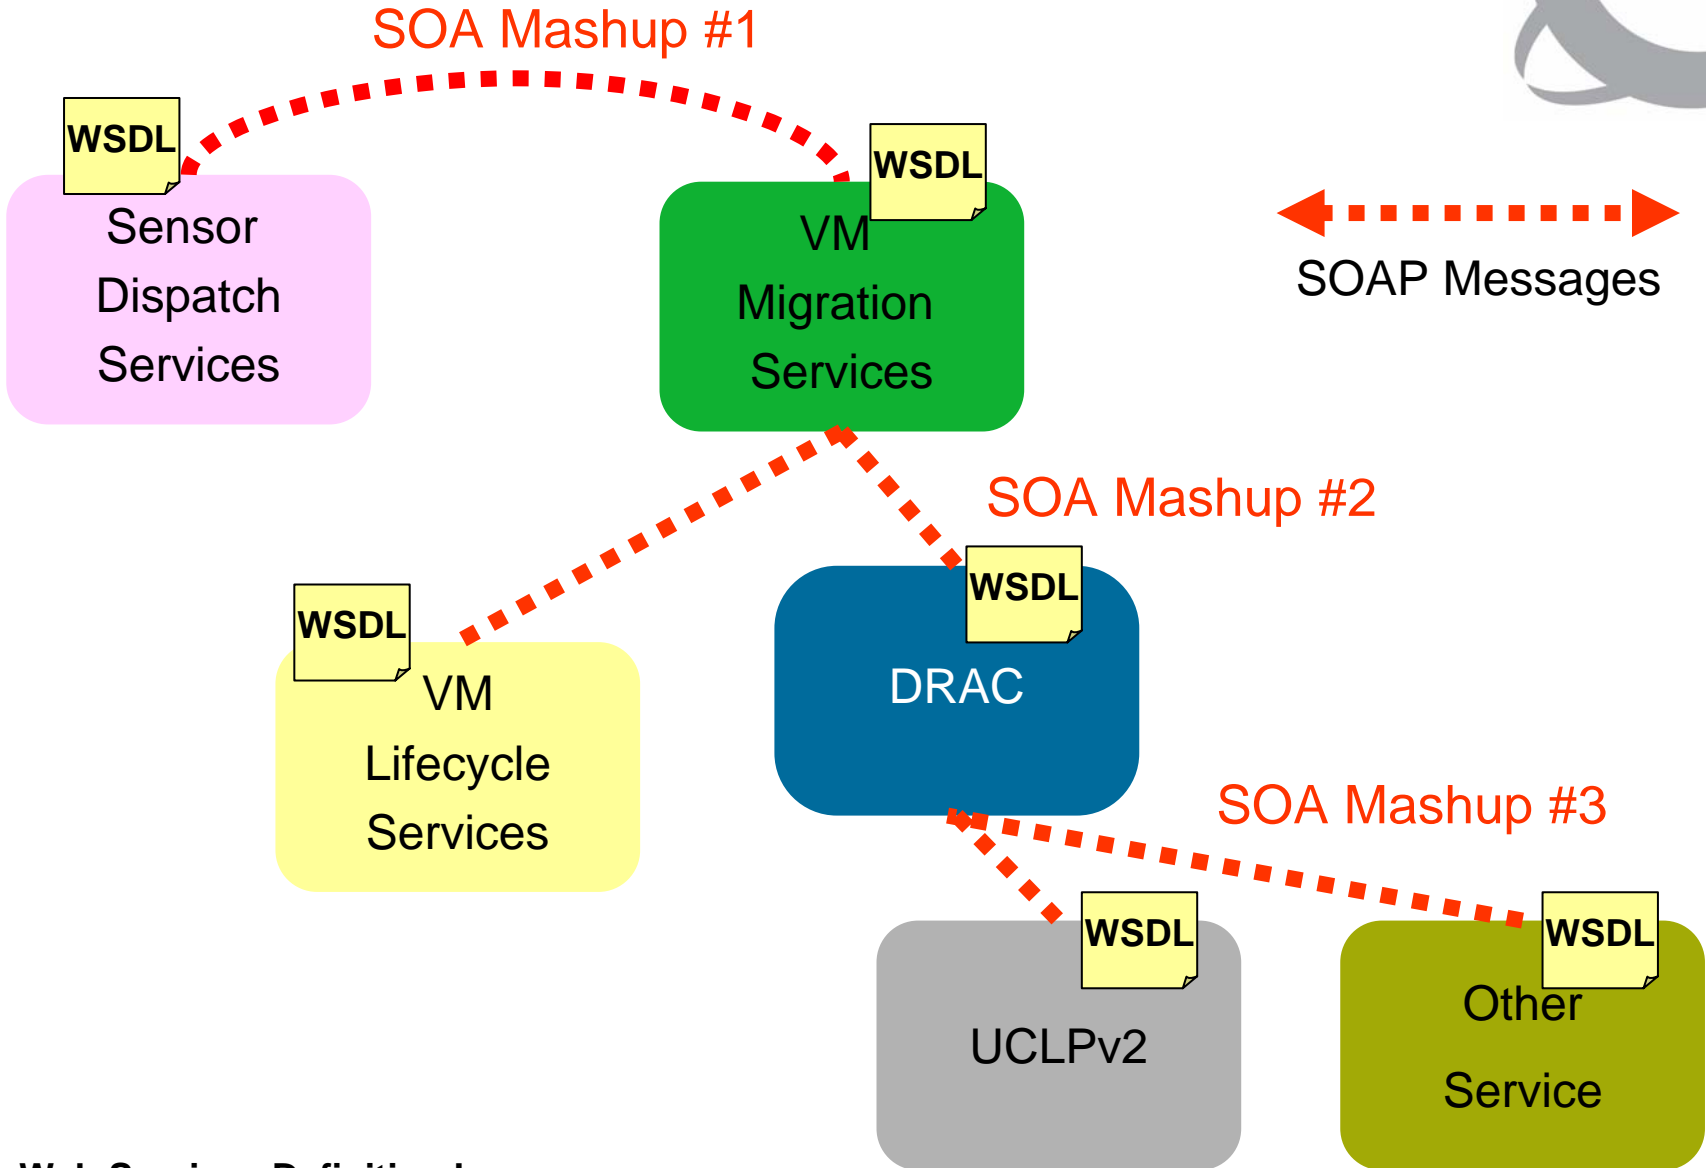


Computation at the Right Place & Time!  
We migrate live Virtual Machines, unbeknownst to applications and clients, for data affinity, BC/DR, load balancing, or power management





# Control Flow



WSDL=Web Services Definition Language

USA-RUSSIA-CHINA-KOREA-NETHERLANDS-CANADA-DENMARK-FINLAND-ICELAND-NORWAY-SWEDEN

# GLORIAD 2006



Global Ring Network for Advanced Applications Development



Produced by Natalia Buleshova, GLORIAD, 2006, Earth Imagery visibleearth.nasa.gov  
Ver. 2006-1, issue 2006-11-03

**VMT demo supported by GLORIAD lightpaths**