

Grids and System Virtualization

BOF

Franco Travostino

travos@nortel.com

Director, Advanced Technology and Research, Nortel
Area Director for Infrastructure, OGF

Intellectual Property Policy

- I acknowledge that participation in OGFx is subject to the OGF Intellectual Property Policy.
- Intellectual Property Notices Note Well: All statements related to the activities of the OGF and addressed to the OGF are subject to all provisions of Section 17 of GFD-C.1 (.pdf), which grants to the OGF and its participants certain licenses and rights in such statements. Such statements include verbal statements in OGF meetings, as well as written and electronic communications made at any time or place, which are addressed to: the OGF plenary session,
 - any OGF working group or portion thereof,
 - the GFSG, or any member thereof on behalf of the GFSG,
 - the GFAC, or any member thereof on behalf of the GFAC,
 - any OGF mailing list, including any working group or research group list, or any other list functioning under OGF auspices,
 - the GFD Editor or the GWD process
- Statements made outside of a OGF meeting, mailing list or other function, that are clearly not intended to be input to an OGF activity, group or function, are not subject to these provisions.
- Excerpt from Section 17 of GFD-C.1 Where the GFSG knows of rights, or claimed rights, the OGF secretariat shall attempt to obtain from the claimant of such rights, a written assurance that upon approval by the GFSG of the relevant OGF document(s), any party will be able to obtain the right to implement, use and distribute the technology or works when implementing, using or distributing technology based upon the specific specification(s) under openly specified, reasonable, non-discriminatory terms. The working group or research group proposing the use of the technology with respect to which the proprietary rights are claimed may assist the OGF secretariat in this effort. The results of this procedure shall not affect advancement of document, except that the GFSG may defer approval where a delay may facilitate the obtaining of such assurances. The results will, however, be recorded by the OGF Secretariat, and made available. The GFSG may also direct that a summary of the results be included in any GFD published containing the specification.

OGF Intellectual Property Policies are adapted from the IETF Intellectual Property Policies that support the Internet Standards Process.

Agenda

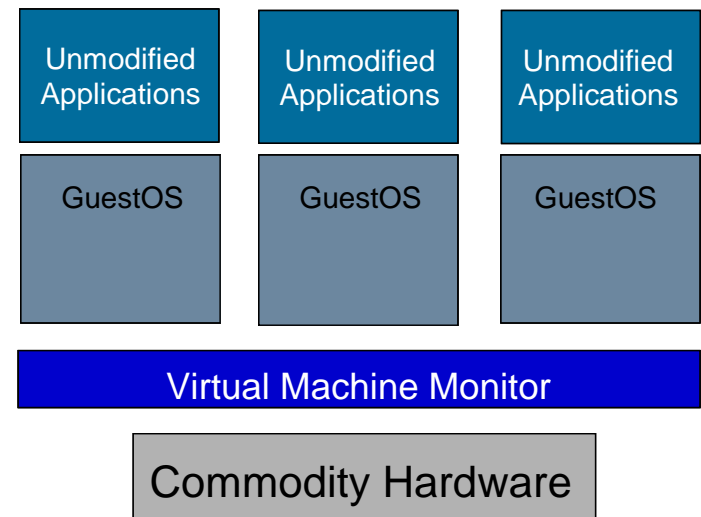
- i. Agenda & administrivia (Franco, 5')
- ii. Introduction and scope definition (Franco, 10')
- iii. A use case in focus: power-aware placement of Grid computation (Cees, 10')
- iv. Technical review of proposed work: functional description, approach discussion (Yaron, 30')
- v. Intercepts with DMTF's work and the SVPC group in particular - the viewpoint of the SVPC Chair Ron Goering (10'-15')
- vi. Proposed charter and milestone review
- vii. Discussion 'till the end

1 – Virtual Machines

the “new” workhorse of computation

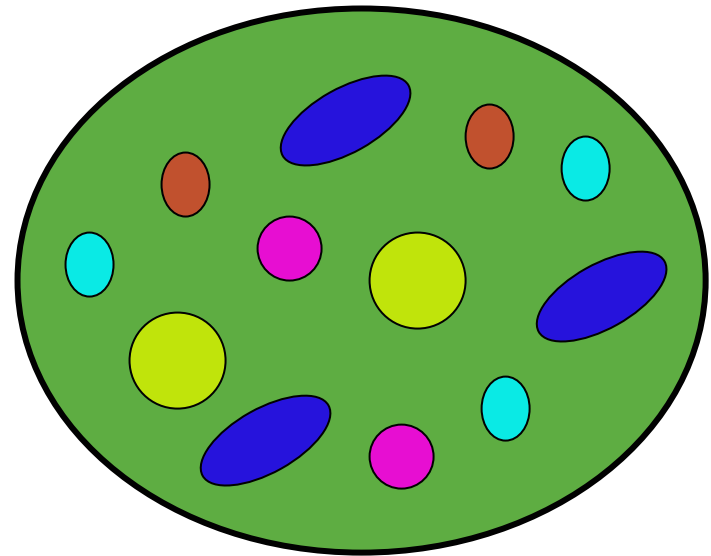


- With VMs, multiple guest Operating Systems coexist on the same physical hardware, unbeknownst to one another
- Use VMs for consolidation, isolation, migration
- A vision that came a long way
 - 1967: cp67 system for IBM 360
 - '80s: Microkernels (Minix, Mach)
 - '90s: the Java Virtual Machine
 - 1998: VMware
 - 2003: Xen’s “paravirtualization”
 - 2006: VM-ready chipsets



2 – Virtualized I/O

- **A capability to TX or RX bits w/ given SLA**
 - statistic or deterministic fraction of system's I/O estate
- **Can be expressed by way of:**
 - IP endpoints
 - TOE handles
 - IPsec SAs
 - Infiniband QPs & VLs
 - Etc.



1 + 2 create new opportunities for Grids



- Although the lifecycle of 1,2 may be sysadm as usual
- ... there are cases wherein a Grid infrastructure benefits from actively driving aspects of 1+2. E.g.:
 - I want my workload carried exclusively by just-in-time instantiated virtualized systems (e.g., for greater isolation, malware avoidance, etc)
 - I want my workload to be executed by virtualized systems because I anticipate the need to perform suspend/resume operations on the workload prior to its completion;
 - I want to migrate my running workload, inside or outside a data center, because I need to optimize data/instrument locality, harvest "track the moon" efficiencies, increase geographical spread, etc.

Chief points before us...

- a) Are Grid specifications neutral to 1+2?
 - i.e. do they apply regardless of compute and I/O resources being physical vs. virtual

- b) How should we enhance the Grid specifications to drive the new features enabled by 1+2?

In the larger schema of things

Besides #1 (Virtual Machines) and #2 (Virtual I/O)

- There is a #3, Virtualized Network Resources
 - OGF's GHPN-RG tackling Grid Network Services (e.g., end-to-end reserved bandwidth)
- There is a #4, Trusted Platform Module
 - OGF's TC-RG exploring Trusted Computing avenue

Outside OGF

- There is DMTF Sys. Virtualization Partitioning Clustering
 - Will be hearing today from its Chair, Ron Goering

Backup