Distributed Virtualization for Net-Centric Operations

Bob Marcus

Definitions

Distributed Virtualization is the transparent sharing of distributed resources by multiple clients.

Transparent sharing means clients are not directly aware of each other and the underlying physical resources

Virtualization Layers

Virtualization enables access at any layer while hiding the layers below

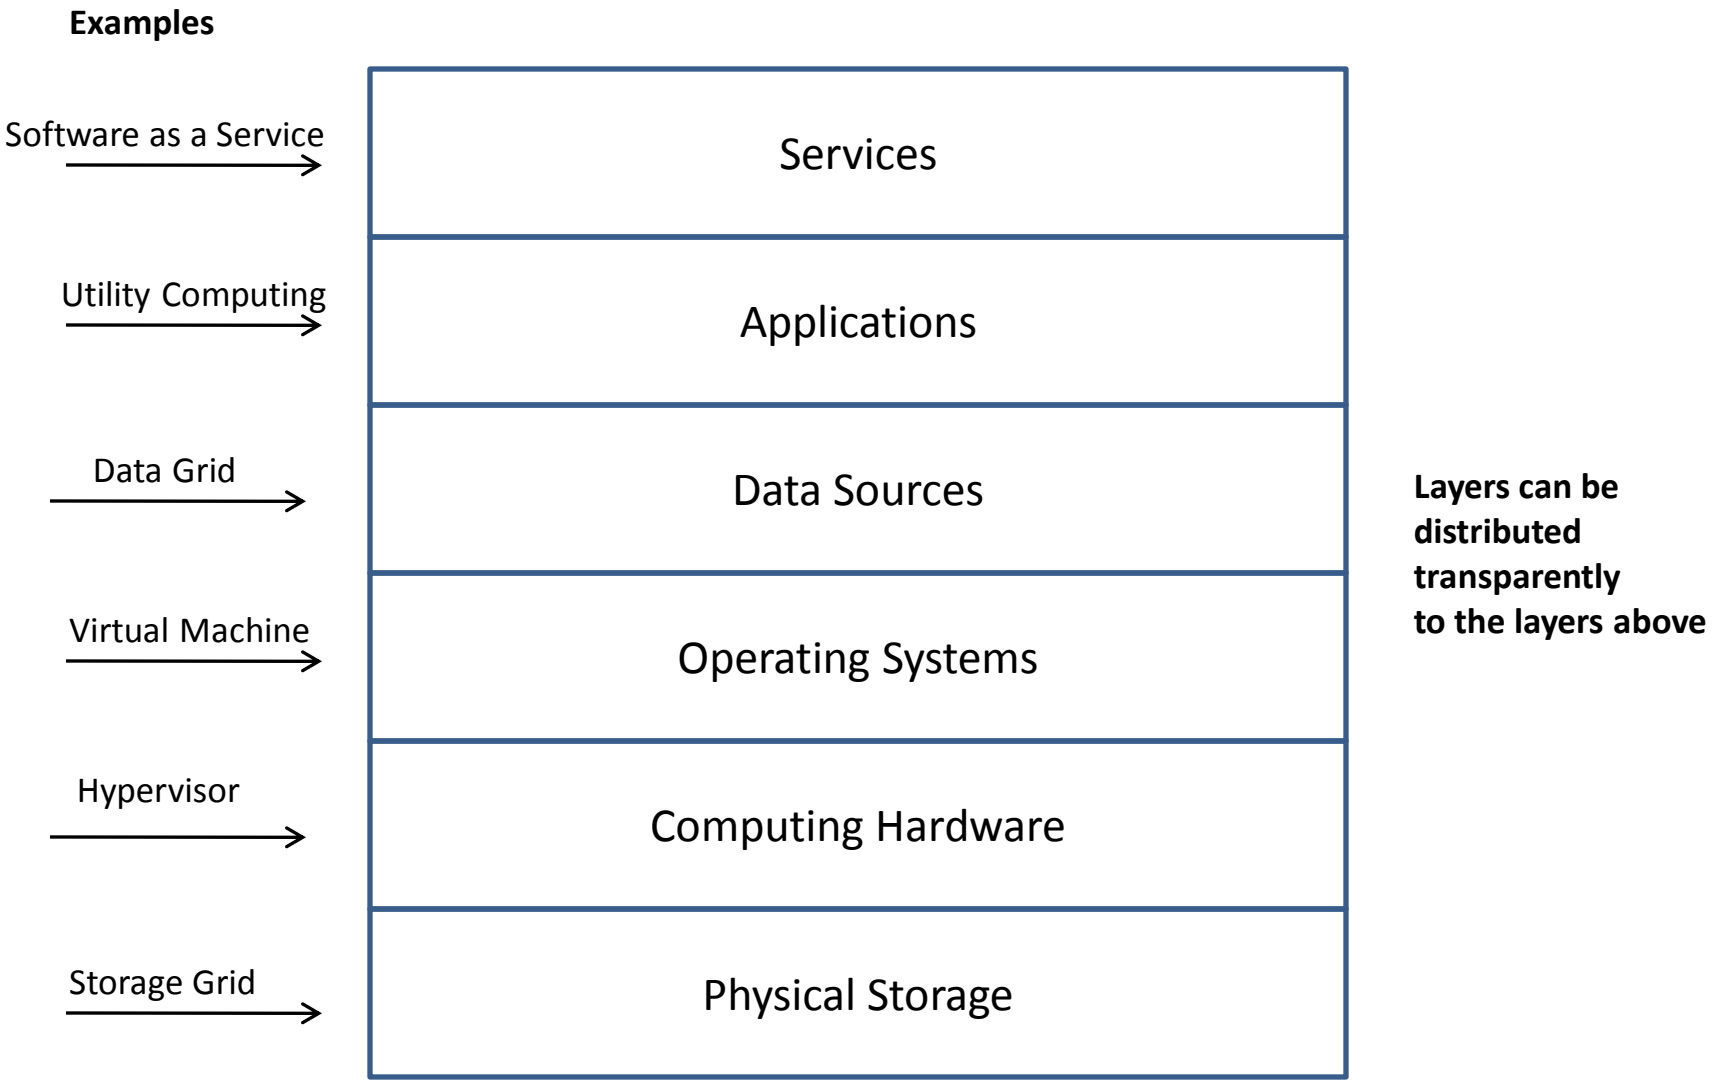

Virtualization Alternatives

- 1. Software as a Service
- 2. Utility Computing
- 3. Computational Grids
- 4. Transaction Grids
- 5. Data Grids
- 6. Storage Grid or Utility
- 7. Virtual Machine
- 8. Virtual Server
- 9. Virtual Machine Monitor (Hypervisor)
- 10. Virtual Appliance

Generic Virtualization

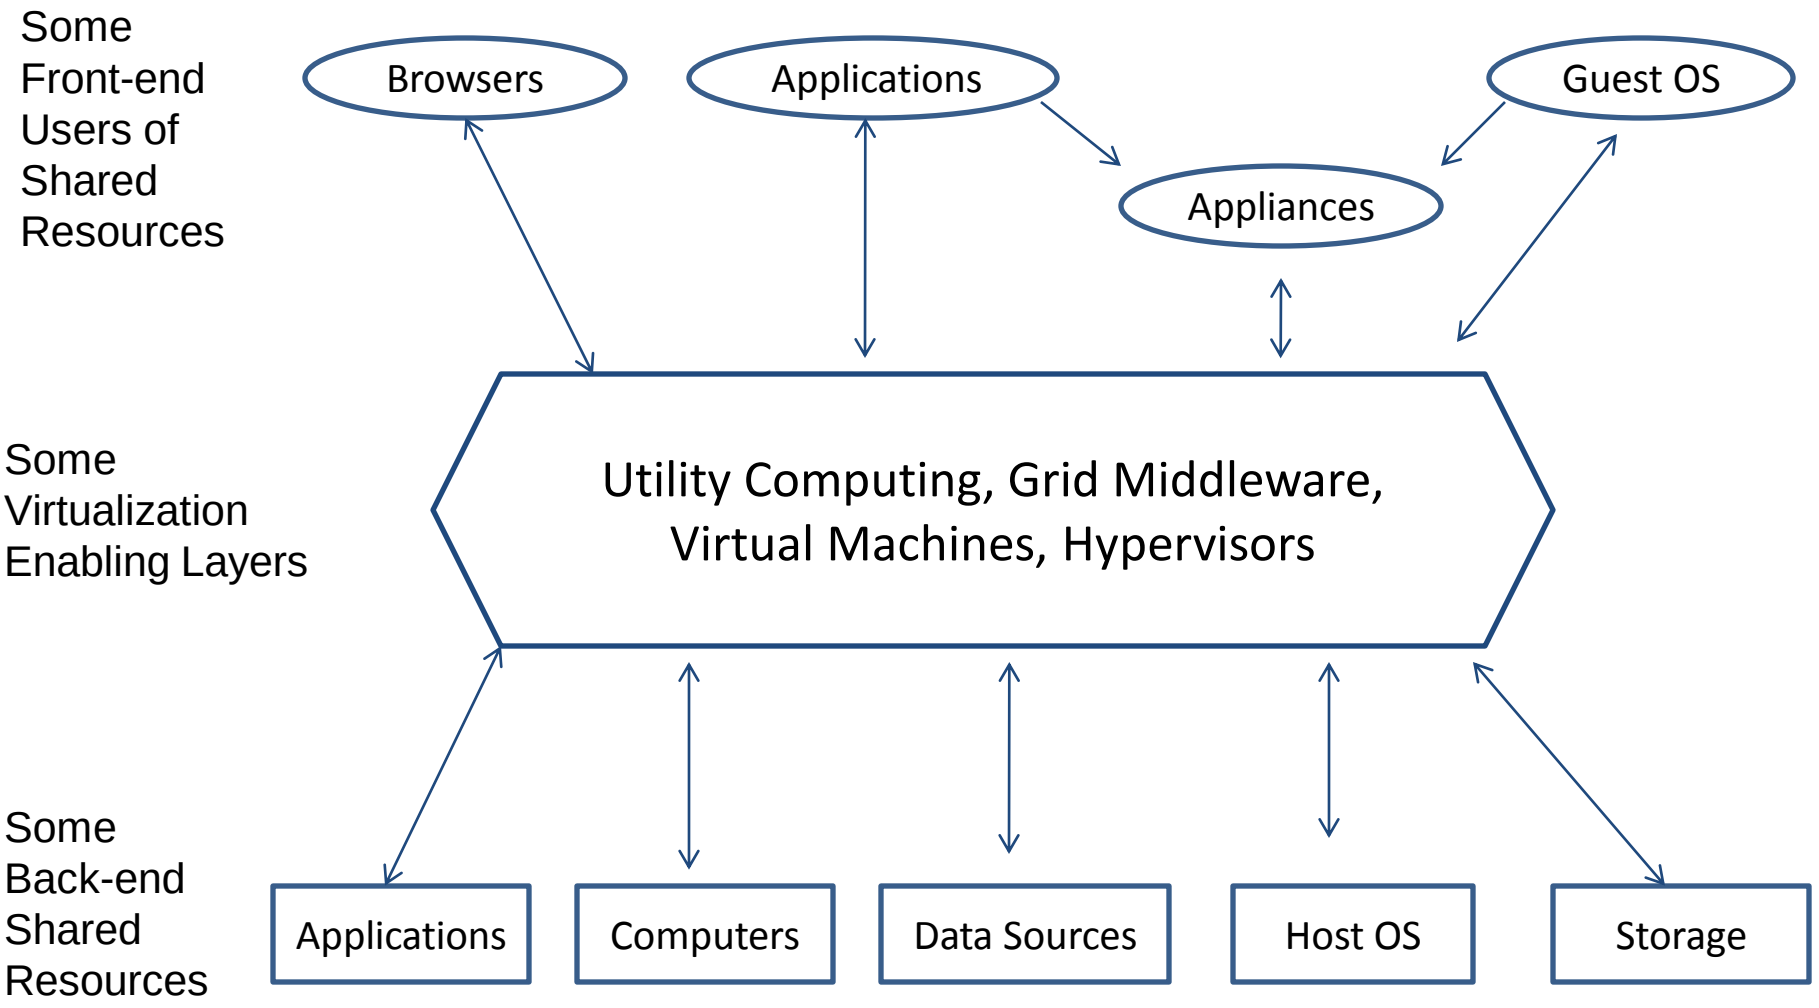

Virtualization Alternatives

Alternative	Shared Resource	Resource Users	Enabling Layer
Software as a Service	Application accessed as Web Services	Web clients paying per use	Multi-tenant architectures
Utility Computing (On Demand, Cloud)	Distributed data center software and hardware	Multiple clients renting resources	Distributed resource and workload managers
Computational Grids	Computers across locations or organizations	Multiple groups sharing computing resources	Grid middleware
Transaction Grids (Fabrics)	Hardware and software within an organization	Enterprise applications	Fabric middleware
Data Grids	Data sources across locations or organizations	Multiple groups creating a shared data capability	Data source resource broker
Storage Grids or Utilities	Storage hardware	Multiple applications and databases	Storage broker
Virtual Machine	Execution environment	Processes	Run-time support
Virtual Server	OS and CPU	Applications running in multiple partitions	Virtual server support
Virtual Machine Monitor	CPU	Multiple OS running on CPU	Hypervisor possibly with CPU support
Virtual Appliance	CPU	Bundled application, OS, database	Application virtualization Packaging and run-time

Horizontal and Vertical Virtualization

- Horizontal Virtualization is virtualization across distributed back-end resources
  - Software as a Service
  - Utility Computing
  - Grids
- Vertical Virtualization is virtualization across architectural layers
  - Virtual Machines
  - Hypervisors
  - Virtual Appliances
- There are a need for additional standards in both horizontal and vertical virtualization
- A key question is the integration of horizontal and vertical virtualization capabilities (e.g. virtual appliances and grids)

Horizontal Virtualization Alternatives

- Software as a service makes applications available in a remote data center through a service-based interfaces available to multiple external organizations
- Utility computing makes resources that are managed by a single organization available to multiple external organizations
- Grid computing combines distributed resources from multiple organizations into a shared resource
  - Computational
  - Transaction
  - Data
  - Storage

1. Software as a Service (SaaS)

- Description – Shared access to applications as remote services by different organizations
- Suppliers
  - Salesforce.com
  - Netsuite
  - Webex
  - http://saas-showplace.com/
- Benefits – Reduced cost for software and infrastructure
- Issues – Security across multiple uses

2. Utility Computing (Cloud)

- Description - Distributed data center resources made available as necessary.
- Suppliers
  - Amazon (Elastic Computing Cloud, Simple Storage Service)
  - IBM Distributed Computing Capacity On Demand
  - Sun Grid Compute Utility
  - HP Managed Capacity
  - Cisco Vframe
- Benefits – Reduced infrastructure cost
- Issues – Accounting, Resource management

3. Computational Grids

- Definition – Transparent sharing of computational server resources among multiple groups across or within an enterprise
- Suppliers
  - Univa UD
  - Globus Alliance
  - Platform
  - Data Synapse
- Benefits – Reduced cost of infrastructure
- Issues – Cross-organization management

4. Transactional Grids

- Definition – Sharing distributed hardware and software platform resources within an organization to support high performance transactional applications
- Suppliers
  - Appistry
  - Gigaspaces
  - Paremus
- Benefits – Reduced cost for transactional capabilities
- Issues – Lack of standards

5. Data Grid

- Definition – Transparent sharing of data servers among multiple groups
- Suppliers
  - UCSD – Storage resource broker
- Benefits – Easier access to distributed data
- Issues - Maintenance of metadata and data consistency

6 Storage Grids and Utilities

- Definition – Transparent sharing of distributed physical storage devices by multiple clients.
- Suppliers
  - Amazon C3 (Storage utility)
- Benefits – Reduced infrastructure costs
- Issues - Performance

Vertical Virtualization Alternatives

- Type 1 Virtual Machine Monitor - Hypervisor running directly on top of a CPU providing an environment for guest operating systems
- Type 2 Virtual Machine Monitor - Hypervisor running on top of a host operating system providing an environment for guest operating systems
- Application Virtual Machine – Platform (CPU, OS) independent environment for running applications
- Virtual Appliance - Pre-configured bundling of application and operating system capabilities into a module that can run on a virtual machine. Provides a means of rapidly deploying applications using OVF standard.

7. Virtual Machine Monitor within Host OS

- Definition – Virtual machine capabilities built on top of a specific operating system. Can be used to partition resources or to host guest operating systems.
- Suppliers
  - Microsoft Virtual Server
  - SWSoft Virtuozzo
  - VMWare Virtual Machine Server
  - Sun Solaris containers
- Benefits – Better utilization for resources
- Issues - Performance

8. Virtual Machine Monitor on CPU

- Definition - Virtual machine capabilities built on top of a CPU not requiring a host operating system
- Suppliers
  - VMWare ESX Server
  - Xen
  - IBM LPAR
- Benefits – Better utilization of resources
- Issues - Functionality

9. Virtual Machines

- Definition - Platform (CPU, OS) independent environment for running applications
- Suppliers
  - Java Virtual Machine
  - Microsoft CLR
- Benefits - Portability
- Issues – Performance

10. Virtual Appliances

- Definition - Pre-configured bundling of application and operating system capabilities into a module that can run on a virtual machine. Provides a means of rapidly deploying applications for utility computing using OVF standard
- Suppliers
  - VMWare Virtual Appliance
  - Amazon Machine Images
- Benefits - Ease of deployment
- Issues – Managing evolving interdependencies across multiple appliances and physical environments

Integrating Horizontal and Vertical Virtualization

Key Standard: Open Virtual Machine Format (OVF)

- Distributed Management Task Force (DMTF) standard
- Open, secure, portable, efficient and extensible format for the packaging and distribution of (collections of) virtual machines
- Its goal is to facilitate the automated, secure management not only of virtual machines, but the appliance as a functional unit
- Uses existing packaging tools to combine one or more virtual machines together with a standards-based XML wrapper
- Gives the virtualization platform a portable package containing all required installation and configuration parameters for the virtual machines
- Allows any virtualization platform that implements the standard to correctly install and run the virtual machines

Future Directions

- Open Virtual Machine Format (OVF) provides a portable standard for describing virtual appliances
- Using OVF, virtual appliances and machines can be installed on a wide range of platforms, virtual and non-virtual.
- This enable the secure rapid deployment of applications in a distributed heterogeneous environment to support utility computing, software as a service, and grids
- Distributed virtual appliances and machines have applications in net-centric operations
- It will be necessary to create standard representations to support the configuration of a network of virtual appliances and machines