

ORACLE GRID COMPUTING GLOSSARY

ORACLE GRID COMPUTING

Grid computing is the coordinated use of a large number of servers and storage acting as one computer. With grid computing, businesses no longer need to worry about spikes in demand and the cost of maintaining excess capacity. Computing power is now available whenever it's needed. Grids are built with low-cost modular components, so you can start small and preserve your investment as your business grows.

Like any technology, Grid Computing has a specialized lexicon of terms, acronyms, and concepts. To help you avoid confusion and better understand Grid Computing and what it means to your organization, this glossary includes more than two dozen grid specific words and their definitions.

Abstraction

Changing the interface of an object and exposing it in a more useful form.

Abstraction may be done by providing a layer of software which makes an object or collection of objects easier to manage by exposing a different object with a different interface.

Adaptive Enterprise

A marketing term used extensively by HP. An adaptive enterprise matches supply and demand of IT resources in real time, so the business can quickly manage and capitalize on change. An adaptive enterprise can synchronize its IT resources with the demands of its business environment; similar to IBM's On-Demand Business.

Automation

The replacement of manual operations by computerized methods. Automation can both improve efficiency and reduce errors brought on by manual intervention.

Autonomic Computing

The capability of computer systems and networks to automatically configure themselves to changing conditions and heal themselves in the event of failure.

Autonomy implies that less human intervention is required for operation under such conditions.

Blades

A computing system that includes processors and memory on a single board, but where other resources such as power, cooling, network access, and storage services are shared. Blades are designed to be easily installed and removed and are typically smaller than rack-optimized servers.

Capacity on Demand

Processing power that is available as needed in a timely manner without disrupting other business priorities. Capacity on demand frequently involves additional capacity installed but not available for use until needed.

Clustering

Connecting two or more computers together in such a way that they appear to be a single computing resource. Clustering is used for parallel processing, load balancing, and fault tolerance. Clustering is a popular strategy for implementing grid computing, since it is relatively easy to add new CPUs simply by adding a new server or blade to the overall cluster. Clusters are typically transparent to users and applications.

Cluster Interconnect

A very high-speed connection allowing nodes in a cluster to communicate.

Consolidate

One of the three steps recommended by Oracle to help data centers move to a Grid Computing architecture. By consolidating resources such as servers, storage, network, and data, a company can reduce overall data center management costs while simplifying provisioning of the overall capacity for various business drivers.

Data Center

A facility that provides a suitable environment (power, cooling, network connectivity, management services) for housing information technology equipment (servers, storage) and providing IT services and support to customers.

Data Provisioning

Making data available when and where it's needed and as it becomes available.

Enterprise Grid Alliance (EGA)

The Enterprise Grid Alliance is a consortium of leading vendors and customers focused on developing Enterprise Grid solutions. The EGA, an open, independent, vendor-neutral organization, was formed in order to identify obstacles facing enterprises in adopting grid technologies, and to promote open, interoperable solutions to these problems. <http://www.gridalliance.org>

Enterprise Grid

A collection of networked grid components usually managed by a single enterprise, such as a business entity responsible for managing a pool of resources, a set of services, and the assignment of resources to services to meet business goals. Services range from traditional commercial enterprise applications, such as ERP or CRM, to newer, distributed applications or services; boundaries are defined by management responsibility and control.

Enterprise grids are typically differentiated by management practices and technology which:

- Make management service or application centric, rather than component centric;
- Enable the pooling and sharing of networked resources;
- Enable agility through rapid and automated service provisioning.

Fibre Channel

A high-speed network technology typically used to build Storage Area Networks (SAN). The Fibre Channel Protocol serializes SCSI commands into Fibre Channel frames. Although Fibre Channel can be used as a general-purpose network carrying ATM, IP, and other protocols, it has been primarily used for transporting SCSI traffic from servers to disk arrays. Fibre Channel supports single mode and multimode fiber connections as well as coaxial cable and twisted pair.

www.fibrechannel.org

Gigabit and 10-Gigabit Ethernet

An Ethernet technology that raises transmission speed to either 1Gbps or 10Gbps and is compatible with existing Ethernet networks. Gigabit Ethernet is being deployed in large numbers in both corporate and public data networks, and an even faster 10-Gigabit Ethernet standard is nearing completion.

Global Grid Forum

The standards body for defining standard specifications for global grids.

www.gridforum.org

Globus Alliance

A group that conducts research and development for academic grids. The alliance, creators of the Globus Toolkit, includes representatives from Argonne National Laboratory, the University of Southern California Information Sciences Institute, the University of Chicago, the University of Edinburgh, and the Swedish Center for Parallel Computers. www.globus.org

Globus Toolkit

A kit designed by the Globus Alliance to provide a set of tools based on standard grid APIs. Its latest development version, GT3, is based on standards currently being drafted by the Global Grid Forum.

Grid Computing

A computing architecture that provides computing resources using many computers acting as one virtual computing resource. On the client side, grid computing provides a pool of shared resources, allowing complete transparency as to where and how a task is performed. On the server side, grid computing enables enterprises to provision resources to respond to client requests. Grid computing environments are often typified by:

- **Distributed, shared pools of resources** to achieve greater performance, scaling, resilience, and utilization
- **Flexibility**, as components are regularly re-provisioned in line with business goals
- Application or service architectures, such as **Service Oriented Architectures (SOAs)**, which are disaggregated/distributed
- **Consolidation** into a smaller number of larger resource pools for easier

provisioning and greater utilization

- **Automation** of day-to-day manual maintenance—software installation, patching, upgrading, workload balancing, etc.
- **Standardization** of components and/or their interfaces, configurations, processes, and applications (see "Virtualization").

InfiniBand

InfiniBand is the next generation switch fabric architecture that enables increased network bandwidth (10Gbps), redundant connectivity support, shared resources, lower CPU utilization, and use of industry standard (vs. proprietary) components. InfiniBand technology can be used to connect servers with remote storage, networking devices, and other servers. It can also be used inside servers for inter-processor communication (IPC) in clusters. www.infinibandta.org

LDAP (Lightweight Directory Access Protocol)

A protocol used to access information directories. LDAP is designed to provide a common method for fast hierarchical searching e-mail addresses on the Internet. LDAP makes it possible for almost any application that is running on nearly any computer platform to obtain directory information, such as e-mail addresses, login names and passwords.

N1

Sun's architecture for the next -generation data center. The architecture is designed to make the entire data center behave as a single, unified system. N1 is designed to reduce management complexity and cost, increase data center resource utilization, improve infrastructure responsiveness and agility, and ensure investment protection.

NAS (Network-Attached Storage)

NAS provides storage resources to servers and workstations over a shared or dedicated network. This network is typically Ethernet running the Internet Protocol (IP). NAS operates with industry-standard protocols such as Network File System (NFS), Common Internet File System (CIFS), and SCSI over IP (iSCSI) and uses the server's network interface card (NIC) for communication.

Node

A network processing location in a grid. A node can be a computer, a set of clustered blades, or some other device, such as a printer.

On Demand Business

IBM marketing term that denotes a company whose business processes—integrated end-to-end across the company and with key partners, suppliers and customers—can respond with flexibility and speed to any customer demand, opportunity, or external threat. On-demand businesses are responsive, variable, focused, and resilient.

Pooling

Combining separate computing resources into a single logical group.

Provisioning

Providing or allocating the requested computing resource.

RAID (Redundant Array of Independent Disks)

A method of accessing multiple individual disks as if the array were one larger disk, spreading data access out over these multiple disks. RAID improves access time and reduces the risk of losing all data if one drive fails.

Resource

Any component in a computing system or network used to perform a specific computing task (storage, servers, data, database server, application server, etc.).

SAN (Storage Area Network)

A high-speed sub-network of shared storage devices. In large enterprises, a SAN connects multiple servers to a centralized pool of disk storage. Compared to managing hundreds of servers, each with their own disks, SANs reduce system administration overhead. By treating all the company's storage as a single resource, disk maintenance and routine backups are easier to schedule and control. In some SANs, the disks themselves can copy data to other disks for backup without any processing overhead at the host computers.

Server

A computer resource in a network that is shared by multiple users. The term can refer either to a specific piece of hardware or a software process, such as “database server” or “Web server.”

Server Provisioning

The allocation of servers for use for application software and users.

Service-Oriented Architecture (SOA)

Principles and practices for sharing, reusing, and orchestrating business logic represented as services or components. SOA is the most optimized and cost-effective way to deliver applications deployed on grid-based infrastructures.

SLA (Service Level Agreement)

The agreement between a user and a computing service provider to determine the type, capacity, and quality of service. SLAs are used by vendors and customers as well as internally by IT shops and their end users. They can specify, for example, bandwidth availability, response times for routine and ad hoc queries, response time for problem resolution (network down, machine failure, etc.), and steps to be taken in the event of problems with penalties for non-compliance.

SMP (Symmetric Multiprocessing)

A computer architecture that utilizes multiple CPUs to complete individual processes simultaneously. Any idle processor can be assigned any task, and additional CPUs can be added to improve performance and handle increased loads.

Software Provisioning

Installing, configuring, and making software available for applications and users.

Storage Provisioning

The allocation of storage for use for applications (including databases) and users.

Utility Computing

A pay-as-you-go model of computing. Instead of paying for computer resources to handle the peak load at all times, you pay only for the computing you use; analogous to an electric utility.

Utility Data Center

An infrastructure solution proposed by HP that allows virtualization of computing resources for the data center. The Utility Data Center includes servers, storage, and networking products that are integrated and deployed by intelligent management software that allows them to be shared and dynamically re-provisioned to accommodate changing workloads.

Virtualization

Allows interacting with a resource using an abstract mechanism so that the underlying physical resource can be replaced with another one of similar capability without affecting the resource consumer. Virtualization balances supply and demand by providing a transparent, aggregated computing resource.

VPN (Virtual Private Network)

A private network that is configured within a public network in order to take advantage of the economies of scale and management facilities of large networks. VPNs are widely used by enterprises to create Wide Area Networks that span large geographic areas, to provide site-to-site connections to branch offices and to allow mobile users to dial up their company LANs.

Web Service

A Web service is a software system designed to support interoperable machine-to-machine interaction over a network. It has an interface described in a machine-processable format (specifically WSDL). Other systems interact with the Web service in a manner prescribed by its description using SOAP-messages, typically conveyed using HTTP with an XML serialization in conjunction with other Web-related standards.

Workload Balancing

Distributing the workload across multiple systems to optimize system usage and response time for the user.

For more information on Grid Computing, visit www.oracle.com/grid



Grid Computing Glossary
July, 2005

Oracle Corporation
World Headquarters
500 Oracle Parkway
Redwood Shores, CA 94065
U.S.A.

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
www.oracle.com

Copyright © 2005, Oracle. All rights reserved.

This document is provided for information purposes only and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Other names may be trademarks of their respective owners.