Data Center Transformation from "Always On" to "Always Available"

Eliminating Wasted Energy by Automatically Powering Down Servers.

Introduction

This project will validate management software and supporting hardware which is integrated into various data center components to manage the power-state of large quantities of internet servers. Currently, servers are always powered on even during low or no-use periods when they still consume nearly as much electricity as a fully loaded working server. The management software will monitor server use and turn servers on and off as needed.

The system will also support existing customer service level agreements (SLAs) so that availability of computing services is not compromised. The system of algorithms used will create timely responses to a complete range of server demand scenarios.

The system can be used across multiple data centers for global control of energy use. With global utilization status and trends, customers will be able to use the systems intelligence to track and report on estimated energy consumption and carbon dioxide (CO_2) emissions.

Benefits for Our Industry and Our Nation

In a typical data center, server utilization is 10%–15% with a peak load of up to 90%. Typical consumption of energy during a server idle period is 60% to 80% of the power consumption of a fully loaded server. Servers with low utilization rates use a high percentage of power to support a small number of transactions, resulting in a needlessly high energy cost per transaction. With such constant low load-levels, data centers waste up to 70% of their power to keep all of their servers running for small periods of peak load. Using automated power-management systems, a switch from "always on" to "always available" is expected to save up to 50% of server energy use in data centers with large server farms.

Applications in Our Nation's Industry

The growth of internet applications has created a growing population of large "server farms" where hundreds or even thousands of servers are commissioned to serve the exact same, but scalable service for internet users. There must be enough servers to handle both minimum and maximum loads at acceptable

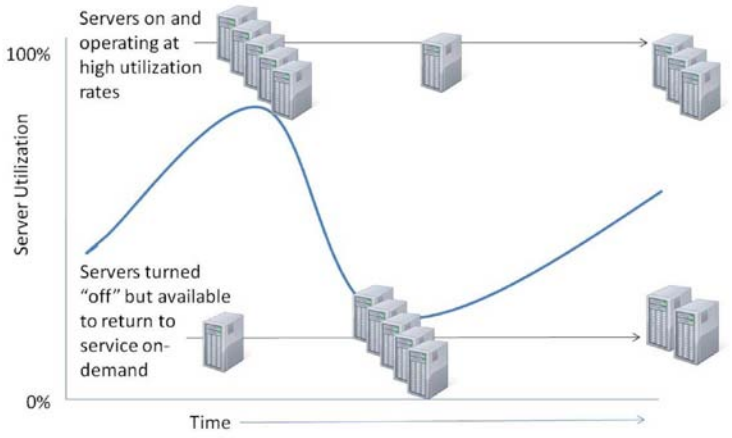

Turning servers on and off according to activity periods.

Illustration courtesy of U.S. Department of Energy's Industrial Technologies Program.

SLAs for response and downtime avoidance. By solving this problem, this project will have greatest impact for these industry representatives:

- Developers and distributors of data center/IT management software
- Electric utilities with demand-side management and renewable energy services
- Owners and operators of large or multiple ICT facilities with many homogenous application servers
- Standards organizations for computer electronics power management protocols and services

Project Description

This project will apply and evaluate the performance of software for controlling servers in data centers with a goal of minimizing server energy consumption. This software will assure data center energy efficiency by transparently and safely powering up/ down servers based on application demand. With its innovative algorithms and non-disruptive operations, this software has shown a reduction in energy usage as much as 70%, with an average reduction of 50%.

This software automatically turns on servers as processing demand increases in a "just in time" manner. The system is designed to scale to support from fewer than 1000 servers up to tens of thousands of servers.

The software continually analyzes application load, power consumption, environmental conditions, utility line status and external events to determine how many servers are needed to support application load at any given moment. It is also designed to work with sever virtualization technologies which further increase the energy efficiency of computing assets.

Barrier

Data center operators lack the information and control needed to monitor and manage power consumption within their facilities. Furthermore, such information is not easily paired with utility charges or level of service information which would inform decision makers about use of assets.

This integrated solution informs the operators, allows both automated and manual control of energy consumption, provides insight into energy and emissions data, and can integrate with electricity demand-side programs. This creates the feedback loop which is missing from these facilities today.

Pathways

This project will finalize the integration of software from two organizations followed by complete testing and verification at several data centers. Project pathways include:

- Selection and installation of data center test sites.
- Final design and integration of products and patented technologies
- Testing and general availability release to test sites

Milestones

- First software release (Completed)
- Final product release
- Beta customer trial and evaluation

Commercialization

This project will integrate software of two key participating organizations into a suite of integrated near commercial ready software. The partners in this project will work together to deploy the integrated suite while training multiple staff members of the licensed installation and operations partner. The end result of this project will be commercial ready field tested software developed by partners with an existing data center customer.

Project Partners

Power Assure, Inc.
Santa Clara, CA
Principal Investigator: Clemens Pfeiffer
E-mail: clemens@powerassure.com

Palo Alto Research Center (PARC) Incorporated
Palo Alto, CA

For additional information, please contact

Gideon Varga
Technology Manager
U.S. Department of Energy
Industrial Technologies Program
Phone: (202) 586-0082
E-mail: Gideon.Varga@ee.doe.gov

Information & Communications Technology

Power Supply Chain

Field Experiment

U.S. DEPARTMENT OF ENERGY
Energy Efficiency & Renewable Energy

EERE Information Center
1-877-EERE-INFO (1-877-337-3463)
eere.energy.gov/informationcenter

DOE/EE-0494 • May 2011
Printed with a renewable-source ink on paper containing at least 50% wastepaper, including 10% post consumer waste.