

How AppSense Performance Manager helps to reduce your carbon footprint



AppSense | Saving the environment since 2002

- Reduce server count
- Reduce your carbon footprint
- Reduce energy costs by up to 40%
- Maintain Quality of Service



Fact file

- Energy costs may increase from 10% of the IT budget today to over 50% in the next few years (Gartner)
- On average, servers still use about 30% of their peak electricity consumption while sitting idle (Forrester)
- 50% of today's data centers have insufficient power and cooling capacity to meet the demands of high-density equipment (Gartner)
- The total power and cooling bill for servers in the US stands at \$14 billion per annum (IDC)

Power consumption is becoming a critical factor for datacenter budgets. The IT department is the second largest consumer of power in organizations, secondly only to lighting and HVAC.

The trend towards high-density, high-power consumption equipment such as blade servers makes the cost of housing, powering and cooling the datacenter more relevant than ever. Not only does this impact the bottom line in terms of power costs, it also creates a larger carbon footprint for your organization. Ratification of the Kyoto Protocol means that limiting greenhouse gas emissions will become part of most organization's regulatory compliance initiatives, and IT as a major power consumer will be a prime target for compliance.

As more business processes are computerized, the number of servers in the typical datacenter is increasing. Meanwhile, electricity costs are rising such that power costs are likely soon to exceed the cost of hardware. It is estimated that only half of the electricity consumed in a typical data centre is actually used by the servers, with 25% going to cooling, 12% to air movement and 13% to electricity transformation and lighting.

It is an irony that most servers typically run at only 5-10% of their resource utilization, yet the difference in power consumption at higher utilization is negligible.

Server consolidation

By optimizing server resources such as CPU, Memory and Disk it is possible to consolidate your datacenter and get the maximum value for the kW hours of power expended. This can lead to savings of hundreds of thousands of dollars per year in power and cooling costs and a significant reduction in carbon footprint.

For example, consolidating 100 physical servers by 40% could save more than 120,000 kWh each year, cutting over 50 tonnes of CO2 emissions. This is as much as an area of mature forest the size of 12 football fields can remove from the atmosphere in the same period.

Using technology for server consolidation

With such a low average resource utilization rate among servers today, technology that can maximize the use of these resources provide opportunity to reduce server count. AppSense has been enabling organizations to reduce server count and take large measurable steps to reducing greenhouse gas emissions for many years.

AppSense Performance Manager dynamically optimizes server resources such as CPU, Memory and Disk to ensure maximum utilization at all times. This in turn provides opportunities for a reduction in hardware, maintenance, cooling and power costs and a tangible reduction in carbon footprint. Many AppSense customers have realized cost savings in the region of 40% by using AppSense Performance Manager to automatically manage their server workload.

Be green while maintaining quality of service

By dynamically optimizing server resources, not only can servers support more users and applications, but the quality of service provided to these users is also optimized. Consistent and predictable application response times are achieved using AppSense Performance Manager through the automatic tailoring of server utilization based on user demand.

About AppSense

We are the leading provider of user virtualization technology to enterprise organizations. User virtualization is a way of managing user-specific information independent of the desktop, and applying this information into any desktop (local install, virtualized, published, streamed etc) on-demand. This enables IT to standardize the desktop build, automate desktop and application delivery, and migrate users to new desktops – all while ensuring the user experience is seamless, personal, predictable and easily manageable.

How AppSense Performance Manager helps to reduce your carbon footprint



AppSense.com
iwanttoknowmore@AppSense.com

Did you know?

It is widely assumed that a typical server draws on average about 300 watts of energy. Based on this number, it will therefore consume approximately 2600 kW per year. According to government figures, 1,960kWh produces 1 tonne of CO₂. This means that a single server generates 1.32 tonnes of CO₂ per year – equivalent to the CO₂ produced by a single passenger flying from New York to Los Angeles!

How AppSense Performance Manager helps to reduce carbon footprint

By defining user and application-based policies, critical resources such as CPU, Memory and Disk can be allocated fairly and prioritized appropriately. Using these business rules, AppSense Performance Manager is able to react to changes in the environment by re-allocating system resources in response to changing demand. This has the effect of maximizing hardware resources and enabling more users to work comfortably on the server. Less servers means less power consumption, less cost and less carbon.

AppSense Performance Manager features:

System Resource Entitlement

Policies are defined for CPU, Memory and Disk management by defining resource reservations and limits for users or applications. Application states may also be included to provide precise control over applications delivered to physical and virtual desktops as well as shared-use environments such as server based computing.

Disk Resource Management

Prevent I/O Request Packet (IRP) bottlenecks from impacting mission critical applications. Disk Resource Management prioritizes the IRPs in accordance with business policy, ensuring disk availability to specific applications by preventing less important processes from creating bottlenecks.

Application Discovery Mode

Application discovery mode gathers the information required to create application groups by quickly scanning target devices. All applications and processes along with property information such as network path are detailed in a comprehensive report. Application groups are created by selecting applications and processes from the report.

CPU Smart Scheduler™

CPU Smart Scheduler™ allocates CPU resource in accordance to business policy by allocating a relative share to the user or application. For instance, if an application is assigned a share factor twice that of a second application, the former will receive higher priority access to the CPU when there is contention.

Thread Throttling™

CPU thread throttling policies automatically trigger when the system is heavily loaded and apply gradual throttling to any runaway threads within each process, preventing rogue processes from consuming excessive CPU resource and reducing the quality of service for all other users on the hardware.

Physical Memory Control

Automatically trim working sets based on application events and states, such as application startup, idle, minimized and in the background. This effectively releases fast access memory (RAM) back to the operating system thereby enabling a significant increase in user density or application instances.

Virtual Memory Optimization

By automatically analyzing and optimizing the way in which Dynamic Link Libraries (DLLs) are loaded by applications, virtual memory overheads and system paging can be

significantly reduced. Optimized DLLs are stored in a separate cache and loaded dynamically, leaving the original applications intact.

Virtual Memory Limits

User memory limits can be applied to restrict the amount of virtual memory utilized. Users can be warned, and then prevented from launching additional applications, when virtual memory utilization reaches critical levels. Application memory limits can also be applied to individual applications giving greater control over virtual memory consumption and on a per application basis.

CPU Application Limits

Administrators can also define hard CPU limits, to restrict an application's access to processor resources. For instance, if an application is limited to 70% then it will never be allowed to use more than 70% of the CPU resources.

CPU Reservations

Define CPU reservations to provide mission critical applications with a guaranteed minimum resource allocation. For example, if an application is allocated a reservation of 20% it will continue to get priority access to the CPU while it is using 20% or less of the CPU resources.

CPU Affinity Assignment

Guarantee processing power goes where it's most needed. On multiprocessor systems, policies can be assigned which bind specific users and applications to a CPU. This allows mission critical applications to run exclusively on a dedicated CPU.

Statistical Analysis and Reporting

Report on CPU, Memory and Disk usage at the process level on a per user or application basis. Tabular reports and graphs are used to report on defined events to show resource consumption and optimization. When used in conjunction with AppSense Management Center, alerts can be raised when configurable events are triggered.

AppSense Configuration Templates

Take full advantage of pre-built corporate policy best practice by importing AppSense Configuration Templates. AppSense Performance Manager is able to import an unlimited number of resource configurations and use these in combination. A selection of Configuration Templates, such as "BoostOffice" to prioritize resources to the MS Office application set, is available from www.myAppSense.com. This template library is maintained and updated frequently.