



Reducing Cost with AppSense:

Increasing the number of users per Server

When server resources such as CPU, memory and disk are left unmanaged on a Terminal Server, there is a good chance the number of users that server can support is not optimal. The number of users serviced per machine directly impacts the bottom line, as hard costs such as hardware, power, cooling, maintenance and software are directly impacted.

By analyzing the per-user and per-application demands on hardware resources and applying business rules to share available resources fairly across the user community, the overall efficiency of the server farm can be optimized. This optimization can have a tremendous effect on the number of users supported per server, with obvious downstream effects to cost.

AppSense Performance Manager optimizes server resources such as CPU, Memory (Virtual and Physical) and Disk usage at an application, user, process and even the thread level, resulting in maximum utilization of hardware resources at all times. With this increase in server utilization, the number of concurrent users accessing the machine can be significantly increased, with no degradation to the end user's quality of service.

Third party tests have shown that on average, 40% of hardware costs can be saved using AppSense Performance manager to make more efficient use of Terminal Server resources.

Based on Gartner figures, using AppSense Performance Manager to reduce the number of servers by 40% will save an organization on average \$8,531 per server over 3 years. AppSense Performance Manager has return on investment period of 5 months, which means an organization can save \$6 for every \$1 spent.

Grange Insurance increases user count by 100% with AppSense

“ I was amazed by how much we saved. We were looking at spending \$30,000 on hardware, storage and support to roll out an additional three Citrix servers. For a third of that we implemented AppSense Performance Manager. Even our most resource intensive applications can fit double the number of users on each Citrix server - which is a 100% improvement on server utilization too. ”

Adam Wilson, Citrix Administrator,
Grange Insurance.

AppSense Performance Manager prescribed for healthy servers

“ With the introduction of AppSense Performance Manager we have seen a 100% increase in our server capabilities in terms of concurrent connections. This equates to between 60 and 80 users per blade with no degradation in response time. The result has been a projected 20% reduction in the number of servers required to complete the project. ”

Greg Davis, Senior Enterprise Network Analyst, Banner Health

KEY FEATURES

- > Resource Policy Management
- > CPU Scheduling and Disk Management
- > Physical and Virtual Memory Optimization
- > Can be used on Terminal Server / Citrix XenApp, VMware View, Citrix XenDesktop as well as physical PCs

KEY BENEFITS

- > Increase server capacity and consolidate hardware
- > Extend hardware lifecycle
- > Reduce power consumption and carbon emissions
- > No need for multiple point solutions, one technology for all platforms

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

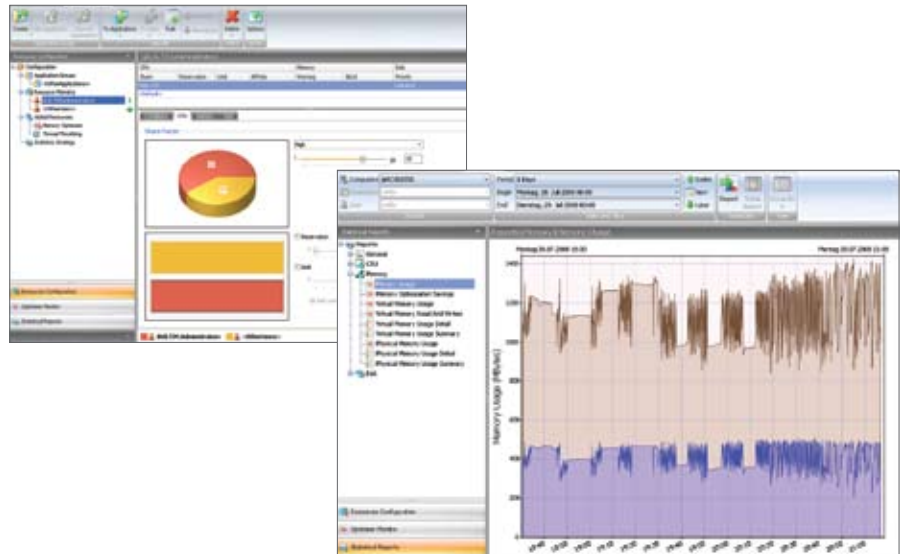
Define CPU reservations to provide mission critical applications with a guaranteed minimum resource allocation. For example, if an application is allocated a reservation of 20% it will continue to get priority access to the CPU while it is using 20% or less of the CPU resources.

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.



AppSense Management Suite is used in server based computing environments such as Microsoft Terminal Services and Citrix XenApp, and is also used in hosted virtual desktop environments and local PCs to ensure users receive a consistent, predictable and responsive working environment.



To learn more about AppSense Management Suite, please visit www.appsense.com

