



NEC High-Availability Fault Tolerant-Virtual Server Solution

NEC Solutions (America), Inc.
2890 Scott Boulevard
Santa Clara, CA 95050
www.necsam.com
888-632-8701

Table of Contents

<i>Abstract</i>	2
<i>Introduction</i>	3
<i>Business Benefits of the FT-VS Solution</i>	4
Server Consolidation	4
Re-hosting Legacy Applications	4
Reduce Software Development and Testing Cycles	4
Lower Capital Expenses	5
Server Proliferation	5
Standard vs. Cluster Application Costs	5
Lower Operating Expenses	6
Administration Cost.....	6
Facilities Cost	6
Support Cost.....	6
<i>Fault Tolerant Servers by NEC Solutions America</i>	7
Introducing the Express5800/340Hb-R	7
Availability	8
Hardware Design	8
CPU Modules.....	8
I/O modules	9
Full Technology Stack	9
Highly Stable Operating System.....	9
Hardened Device Drivers	9
<i>Microsoft Virtual Server 2005</i>	10
Virtual Server Architecture	10
<i>Key Features: Maximize Management Efficiency</i>	11
ESMPRO Management Suite	12
Built-in Emergency Remote Management: MWA/RomPilot	12
Support for Plug-in Modules	13
Integration with Microsoft Management Tools	13
Standards-based Management Support	13
<i>Summary</i>	14

Abstract

This document introduces a high availability Fault Tolerant – Virtual Server (FT-VS) solution from NEC Solutions America for IT organizations that want to use virtual machine technologies, but are concerned with servers failing.

Introduction

IT organizations today are faced with pressure to lower facilities costs, limit the overall number of servers, and reduce administration and staffing required for their data centers. Any number of reasons, ranging from the one server-one application paradigm to the overall weak global economy in the past few years, can take the blame. To mitigate the cost cutting pressures of IT organizations, one solution that has been growing in recent years is software virtualization solutions. Gartner predicts that partitioning and mixed-workload technologies, including virtualization solutions, will triple between 2003 and 2008. By 2008, enterprises that do not leverage virtualization technologies will spend 25 percent *more* annually for hardware, software, labor, and space for Intel servers.¹

A single physical server is no longer limited to running one operating system at a time. Software virtualization solutions allow one physical server to simultaneously run different operating systems and related applications each in their own isolated software partition. An application running in a virtual machine (VM) – essentially a computer within a computer, implemented in software – behaves the same way as an application running on a physical system. It writes information to the file system in the same manner, and interacts with the network just as if it were sending packets to a physical network card.

The concern in IT organizations when virtualization solutions are running on *one* server, especially in a production environment with business-critical applications, is the risk of losing *all* VMs if the host server – the machine with virtual server software installed – fails. Enterprises seeking to use virtual machines have a new solution to alleviate their fears of server failure issues. NEC Solutions America offers the NEC High-Availability Fault Tolerant (FT) – Virtual Server solution. This solution is comprised of the Intel-based NEC Express5800/340Hb-R server bundled with Microsoft Virtual Server 2005. By delivering the classic 99.999% availability, NEC's High-Availability FT server is ideally suited for use with virtual machine products.

¹ Gartner Research, November 2003.

Business Benefits of the FT-VS Solution

The business benefits of the FT – VS Solution come from selecting a server consolidation method for a specific workload, re-hosting legacy applications, or reducing software development and testing cycles. The solution also has several ways of significantly contributing toward lowering both capital and operating expenses.

Server Consolidation

IT organizations use server consolidation to reduce the number of physical servers under management for an optimized IT infrastructure. The solution consisting of the Express5800/340Hb-R as a host machine with Microsoft Virtual Server 2005 is useful for a variety of consolidation scenarios, but is not a panacea for server consolidation. Microsoft recommends consolidation by workload; specific workloads demand specific approaches. Microsoft is currently promoting Virtual Server 2005 as the recommended approach for server consolidation of infrastructure services, disaster recovery environments and department or branch office services.

Re-hosting Legacy Applications

Virtual Server 2005 on the Express5800/340Hb-R allows IT organizations to run legacy applications in their native software environment in virtual machines without rewriting application logic, reconfiguring networks, or retraining end users. This gives the organization the option to carefully plan to refresh out-of-warranty infrastructure systems first, and then either upgrade or rewrite out-of-service applications on their own schedule.

Reduce Software Development and Testing Cycles

Virtual Server 2005 on the Express5800/340Hb-R reduces software development and testing cycles by shortening the time needed to provision, install, build, test, and restore a new machine. The start of development and testing projects traditionally requires servers to be configured from scratch. The process is very time-consuming. Virtual Server 2005 has the powerful ability to take the testing/development environment and virtualize it completely. By first building ahead of time and then selecting from a library of virtual hard drive images, a developer/tester can now avoid the whole repetitive process of doing individual installations to achieve a desired environment. What once took days can now be shortened to a few hours.

Developers and testers can also take advantage of efficiencies associated with the disk features in Virtual Server 2005 to reduce software development and testing cycles. Utilizing differencing and undo disk features, changes made in a virtual machine can always be rolled back to the previous state. No longer is rebuilding an entire environment necessary to achieve this same effect.

Lower Capital Expenses

Server Proliferation

One of the major causes of the recent necessity for IT organizations to lower capital expenses has been server proliferation. The problem became most acute once servers were inexpensive enough for department managers to purchase, often added ad-hoc to accommodate new users or applications, with the IT organization having little control over their deployment. Add to this the one server, one application paradigm: Application developers typically built applications with the assumption that there would only be one application running on a server, and consequently recommended -or even required- deployment on a dedicated server.

A good indication of the extent of the proliferation problem can be found in examining server utilization rates. Low utilization rates are consistent with the scenario described above of organizations having too many servers that were not added with enough centralized planning coordinated by an IT department. Some studies have suggested that individual server utilization for Windows is generally less than 25 percent.² Microsoft reports that for the installed base of Windows NT 4, the average utilization per server is about 15 percent.³ Contrast this to mainframes. Gartner reports that primarily because of effective virtualization on mainframes, utilization rates easily exceed the 80 percent range.⁴ Microsoft Virtual Server 2005 presents a remedy to the proliferation problem in the Windows environment: now there is the ability to take applications on older, under-utilized servers and move them to a single, newer machine, such as the NEC Express5800/340Hb-R.

Standard vs. Cluster Application Costs

The NEC High-Availability FT – Virtual Server solution, by requiring only single copies of all standard applications on the system, offers a substantial opportunity for lower capital expenses over alternative solutions of linking conventional systems into a cluster network. Running a single version of Virtual Server 2005 on the Express5800/340Hb-R is less expensive than running multiple 4-CPU server clusters each hosting a different operating system and application stack.

For a cluster network, not only are *multiple* copies of applications required, but also the software applications must be customized in order to initiate recovery from a crash and handle shared processing. This custom code can be complex and expensive, depending on the sophistication of the application. It is extremely difficult to code for every possible failover scenario; programmers must select the most likely scenarios and remain available to handle unanticipated failures.

² *Solution Accelerator for Domain Server Consolidation and Migration: Windows NT 4.0 to Windows 2003.*

³ Drive Time. July 2004 Sales Edition. Microsoft Virtual Server.

⁴ Gartner Research, November 2003.

Lower Operating Expenses

Companies who are considering the move to a Virtual Server environment may be able to justify the financial investment purely in terms of operating expenses eliminated by the new hardware and software. Expenses such as service contracts, operating system patches, upgrades and maintenance must be included in any financial justification to show the full impact of the purchase. The reduction in operating expenses alone could be used to pay for the new equipment, even if there is no money left in the capital budget. Generally server hardware purchased via leases are paid out of a company's operating budget; therefore, if the new equipment eliminates more expenses than the cost of the monthly lease payments, the company is better off as a result of this investment.

Operating leases are also very useful if the company has a hiring or headcount freeze – where generally there is money in the operating budget that cannot be used. Purchasing a Virtual Server environment with an operating lease in effect trades off equipment for headcount by reducing the administrative and management work required by the remaining people.

Administration Cost

The unique modular design of the NEC Express5800/340Hb-R allows service by non-technical staff without incurring server downtime, application interruption, or data loss. Taking advantage of using these less expensive personnel can greatly contribute toward keeping IT administration costs down. Should a component require replacement, a non-technical staff member can easily remove the module and replace it with another.

Another advantage of the NEC Express5800/340Hb-R is that monitoring, testing, and software upgrades can be performed remotely by a single, centralized IT staff, eliminating the duplication of staff in remote office environments.

Facilities Cost

Another consequence of the server proliferation problem described earlier in this paper is high facilities costs. An overabundance of servers means greater space, power and cooling demands for the IT organization. Taking advantage of Virtual Server's ability to consolidate servers, and thus lower the overall server count in a data center, will positively contribute toward lowering facilities costs.

Support Cost

If there is a support issue for a Microsoft application in a guest operating system created by a non-Microsoft virtual machine, Microsoft will not offer any support. Microsoft will require the problem be first reproduced on a stand-alone server. Contact for support to the non-Microsoft virtual machine vendor would be a first step to resolving the problem, but if this is not sufficient, it will undoubtedly take extra company time and resources to reproduce the problem on a stand-alone server. A data center with the NEC High-Availability FT – Virtual Server solution has the potential for everything – from the virtual server product, applications, and operating system – to have support come from one company, Microsoft. A call to just one company for support means lower support

costs involved for resolving the problem when compared to having to resort to the alternative method explained above.

Fault Tolerant Servers by NEC Solutions America

NEC Solutions America launched a line of fault tolerant servers (FT servers) in 2001. Working with Microsoft and Intel, NEC was able to bring the benefits of fault tolerant systems running proprietary UNIX operating systems to the Windows-Intel server platform. Companies that could not afford the hundreds of thousands of dollars it cost for proprietary fault tolerant servers can now much more easily afford FT servers running the Windows operating system.

Introducing the Express5800/340Hb-R

NEC is extending their fault tolerant offering from the NEC Express5800/320Lb (2-way server) to the higher ends of the market with a new 4-way server, the Express5800/340Hb-R. This server can be configured with:

- Maximum of 8 physical or 4 logical (4 x CPUs output) CPUs
- 12GB of logical memory
- 2 x 10/100/1000 redundant NICs
- 8 full sized PCI slots (6 available)
- Redundant video cards
- CD/External floppy drive
- Redundant serial/USB and SCSI ports

This configuration provides the performance to host a Microsoft Virtual Server architecture. With four logical CPUs the server has the computing power to support multiple operating systems and application stacks in a fully fault tolerant environment without sacrificing application performance.

Like its predecessor, the NEC Express5800/340Hb-R delivers up to 99.999% continuous availability, averaging less than five minutes of unplanned downtime per year, including time spent repairing failures, installing upgrades, and performing general maintenance. This is one of the highest levels of availability in the IT industry. It is this general availability and ease of maintenance that make the NEC Express5800/340Hb-R ideal for being the host server for Microsoft Virtual Server 2005.

Availability

Table 1 compares availability levels, clearly showing the dominant position of the Express 5800/340Hb-R in the Fault Tolerant category.

	Availability	Average Annual Downtime
Fault Tolerant Continuous Availability (CA)	99.999%	5 Minutes
Cluster High Availability (HA)	99.9%	8 Hours 45 Minutes
Stand Alone Server w/RAID	99.5%	43 Hours 23 Minutes

Table 1: Availability comparison.

Source: IDC

Hardware Design

The hardware design is the key to minimizing downtime. The Express5800/340Hb-R utilizes blade-like packaging, separating processing/memory, and I/O onto separate customer replaceable, hot-swappable modules. If any hardware component within a module fails, the processing continues uninterrupted, without the time delays associated with cluster-based solutions. The failed module can be replaced while the application continues to run at 100% on the redundant module without the downtime, data loss, or application interruption associated with traditional server architectures. The result is a server designed to handle failed hardware components within the system and still remain fully operational without any loss of data or performance.

CPU Modules

The foundation of fault tolerance is the hardware redundancy and instant failover capacity. The 24-7 design separates the PCI I/O from processing and adds additional fault tolerant detection logic (see Figure 1). The system contains two CPU motherboards for Dual Modular Redundancy (DMR). The CPU motherboards are run in lockstep. Fault-detection and isolation logic (implemented as a custom ASIC) compares CPU output from all motherboards. Any comparison problem indicates an error. DMR systems rely on fault-detection logic on each motherboard to determine which board is in error.

The area in blue represents a general purpose server and shows all the single points of failure inherent in that server design. The area in orange is the FT server from NEC and shows the modular design, the relationship of the modules, and the fact that there is no single point of failure in the server.

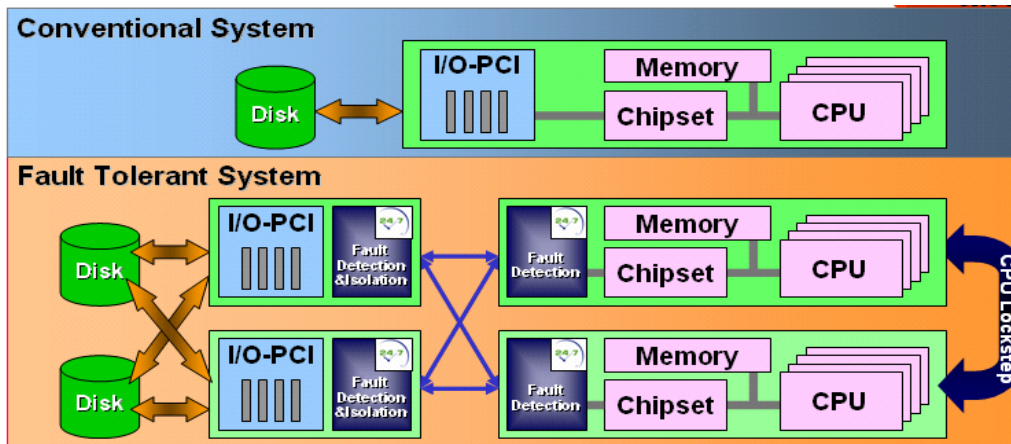


Figure 1: Full Architecture Stack

I/O modules

The base system has two independent PCI busses. All critical PCI adapters are also duplicated (SCSI controllers, Ethernet NICs, Remote Management and Fibre Channel Host Bus Adapters (HBAs)). If a PCI bus or adapter fails, the system will transparently continue operation on the second adapter, with no loss of data or LAN connections.

Full Technology Stack

The other contributing factor to minimizing downtime is the Express5800/340Hb-R delivering the full technology stack. The key components, explained below, include: fault tolerant hardware (covered in the previous section), a highly stable operating system, and hardened software device drivers.

Highly Stable Operating System

The Express5800/340Hb-R uses the Microsoft Windows Server 2003 Enterprise Edition. Microsoft designed it to fully support fault tolerant servers including the lockstepping of the CPUs and fully redundant modules. Windows-based fault tolerant servers must pass the same rigorous Windows Hardware Compatibility Tests (HCT) as general purpose servers, ensuring that the applications running on them will behave no differently than any typical, stand-alone server. The NEC FT servers are fully WHQL certified for all Windows applications.

Hardened Device Drivers

The Express5800/340Hb-R includes hardened software device drivers for each device within the system and for a selected set of fully tested PCI adapters. Each of these drivers is certified for compliance and digitally signed for authenticity using Microsoft procedures and tools. The benefits of using hardened device drivers are to:

- Detect and prevent writing beyond the physical memory allocated
- Perform monitoring and maintenance activities by containing additional instrumentation and diagnostic capabilities
- Provide support for duplex and hot-swap operations

Microsoft Virtual Server 2005

Virtual Server 2005 runs as an application on a host machine operating system. The software transforms physical systems into a pool of logical computing resources. Operating systems and applications are isolated in multiple virtual machines that reside on a single piece of hardware.

Microsoft Virtual Server 2005 features robust storage, networking, and management features in an easy-to-use package that includes a simple, seven-step installation and a Web-based management console. Because Virtual Server 2005 is a Windows-based solution, the operating system is already familiar to administrators, and little additional training is required for implementation. This results in lower cost and faster deployment for the IT organization.

Virtual Server is available in two editions: Microsoft Virtual Server 2005 Enterprise Edition, and Microsoft Virtual Server 2005 Standard Edition. Features across the two Editions are the same, scalability is the only difference. Microsoft Virtual Server 2005 Standard Edition will support up to four logical processors (supported by both the Express5800/320Lb and Express5800/340Hb-R servers by NEC) and Microsoft Virtual Server 2005 Enterprise Edition will support up to 32 physical processors.

Virtual Server Architecture

Figure 2 below illustrates the basic architecture of Microsoft's virtual machine technology.

Starting from the bottom of the logical stack:

- The host operating system – Windows Server 2003 – manages the host system
- Virtual Server provides a Virtual Machine Monitor virtualization layer that manages virtual machines, providing the software infrastructure for hardware emulation
- Each virtual machine consists of a set of virtualized devices (that is, network cards, hard drives, and so on); the virtual hardware for each virtual machine
- A guest operating system and applications run in the virtual machine – unaware, for example, that the network interface card (NIC) that interacts with Virtual Server is only a software simulation of a physical device. When a guest operating system is running, the special-purpose Virtual Machine Monitor kernel takes mediated control over the CPU and hardware during virtual machine operations, creating an isolated environment in which the guest operating system and applications run close to the hardware at the highest possible performance

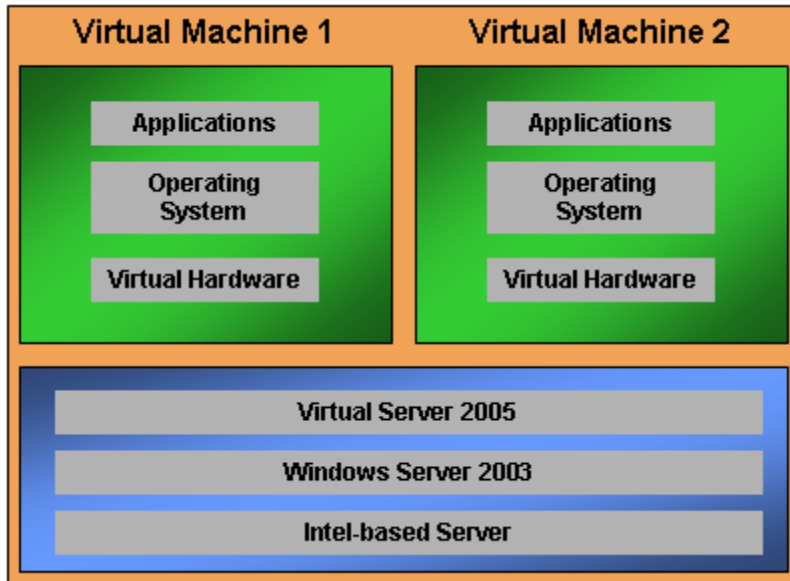


Figure 2: Virtual Server Architecture

Key Features: Maximize Management Efficiency

Both the NEC Express5800/340Hb-R and Microsoft Virtual Server 2005 include several manageability features that contribute to the solution's overall emphasis on maximizing management efficiency of IT organizations.

Three key features for the NEC Express5800/340Hb-R are:

- ESMPRO[®] Management Suite to support monitoring and management of the fault-tolerant components of the system
- Management Workstation Application (MWA)/RomPilot to provide built-in emergency remote management
- Plug-in module support to work with popular Enterprise/Network frameworks

Two key features for Microsoft Virtual Server 2005 are:

- Integration with other Microsoft management tools
- Standards-based management support

ESMPRO Management Suite

The Express5800/340Hb-R includes the ESMPRO Management Suite to support monitoring and management of the fault tolerant components of the system. The ESMPRO suite consists of an industry-standard, SNMP-based server agent and LAN-based manager application that allow for remote administration and maintenance. See Figure 3 for understanding ESMPRO in relation to the operating system, HAL, and user applications.

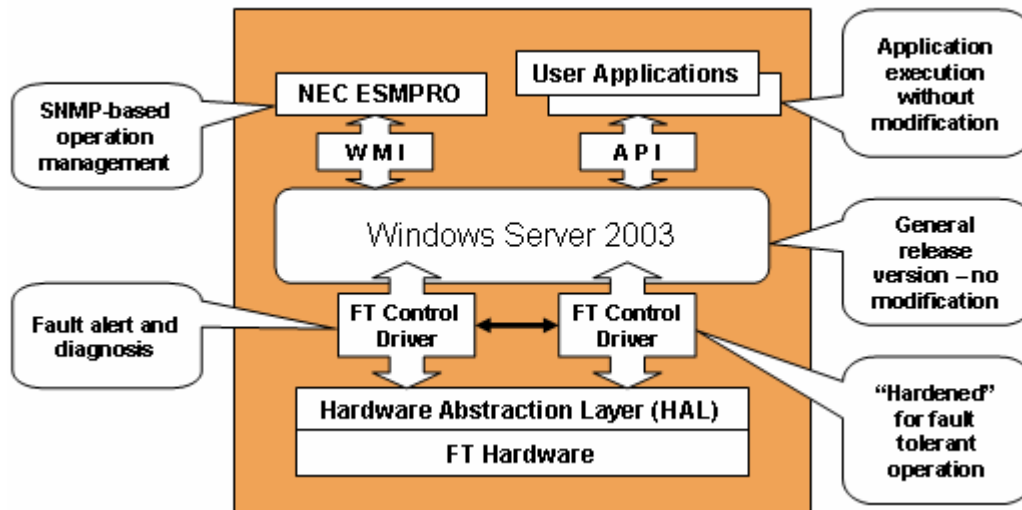


Figure 3: ESMPRO and FT Architecture

The ESMPRO tools allow an administrator to graphically display the system configuration, review operating parameters and thresholds of components, and take a redundant module offline for diagnosis, replacement, or upgrading.

Built-in Emergency Remote Management: MWA/RomPilot

When exceptional conditions arise such as OS failure, LAN failure, server hardware failure or server power loss, the server agents cannot report alerts or send information back to ESMPRO's management console. At the console, the icon for the affected server changes to an unknown state, and there is no information concerning the nature of the problem. To address this condition, NEC includes MWA/RomPilot, for integrated emergency remote management. It is independent of the system CPU or operating system.

MWA is the management software installed on the management console that allows an administrator to access the server over the network, or dial directly into the server system using a standard modem. RomPilot is an extended feature of the system BIOS that provides a pre-boot remote management facility. This includes remote monitoring, access and control, and provides the ability to download DOS utilities (such as maintenance utilities and diagnostics) to run on the remote system.

Support for Plug-in Modules

ESMPRO is designed for easy integration with leading enterprise and network managers such as:

- CA Unicenter TNG/Framework
- IBM/Tivoli Enterprise
- HP OpenView Network Node Manager
- HP BullSoft OpenMaster

ESMPRO uses standard MIB interfaces and these can be captured and utilized by these enterprise management applications to provide a unified device management strategy and tool set.

Integration with Microsoft Management Tools

There are no specialized tools required for Virtual Server that only manage virtual machines. Virtual Server 2005 provides physical machine interoperability with Windows Server System management products, so IT organizations can leverage the same existing tools to manage across both physical and virtual machines. The management tools include:

- Microsoft Operations Manager (MOM) 2005 Management Pack (MP) for Virtual Server provides extensive guest-host mapping for event and performance management
- Virtual Server Migration Toolkit (VSMT) for migrating from physical to virtual (P2V) and virtual to virtual (V2V)
- Systems Management Server (SMS) and Automated Deployment Services (ADS) SMS 2003 Service Pack (SP) 1 and ADS 2.0 will have specific components created for Virtual Server

Standards-based Management Support

Virtual Server 2005 makes use of industry-standard technologies like Hypertext Transfer Protocol (HTTP), Remote Desktop Protocol (Terminal Services), XML, and Performance Monitor to handle common management tasks. Virtual machine configurations are stored in XML files, and server monitoring and configuration is done through an HTTP/HTML interface.

Virtual Server makes it possible for users to administer their servers through any Hypertext Markup Language (HTML) browser on any platform, unlike some competitive products that need their own remote client executable. It is easy to integrate management of the Virtual Server 2005 solution into any enterprise management toolset used in an IT organization.

Summary

Enterprises seeking to use virtual machine technologies have a new solution to alleviate their fears of server failure issues. The NEC High Availability FT – Virtual Server solution is comprised of the Intel-based NEC Express5800/340Hb-R server bundled with Microsoft’s Virtual Server 2005 software. By delivering the classic “five nines” availability, the Express5800/340Hb-R server is ideally suited for use with Microsoft’s virtual machine product.

The main business benefits of the FT – Virtual Server Solution come from selecting a server consolidation method for a specific workload, re-hosting legacy applications, or reducing software development and testing cycles. The solution also has several ways of significantly contributing toward lowering both capital and operating expenses.

In order to contribute to maximizing management efficiency, the NEC Express5800/340Hb-R server has integrated server management features that allow centralized management to perform monitoring, testing, and software upgrades even at remote office environments. Virtual Server 2005 offers integration with other Microsoft management tools (for example, MOM and SMS) and industry-standard technologies such as HTTP, Remote Desktop Protocol, XML, and Performance Monitor.

IT organizations seeking to use virtual machine technologies now have a new solution to consider.

About NEC Solutions (America), Inc.

NEC Solutions (America), Inc. is a premier provider of integrated solutions for the Connected Enterprise in North America. As an affiliate of NEC Corporation (NASDAQ: NIPNY) (FTSE: 6701q.1), NEC Solutions America taps into a global resource network to help clients leverage technology to achieve a competitive edge. From mobile enterprise computing systems, biometric security solutions, business intelligence, projector and plasma display solutions, business services management and IT professional services, the expertise is delivered with the personal attention needed to address individual situations. With headquarters in Rancho Cordova, California, NEC Solutions America maintains research, marketing, sales and support facilities throughout the United States. Information regarding NEC Solutions America can be found at www.necsam.com.

NEC and ESMPRO are registered trademarks of NEC Corporation and/or one or more of its subsidiaries. Intel is a registered trademark of Intel Corporation. Microsoft and Windows are registered trademarks of Microsoft Corporation. RomPilot is a trademark of Phoenix Technologies Ltd. All other trademarks are property of their respective owners. ©2004 NEC Solutions (America), Inc. All rights reserved.

This document is for informational purposes only. NEC Solutions (America), Inc. shall not be liable for technical or editorial errors or omissions contained herein. Information, prices, and specifications are subject to change without notice.

NEC Solutions (America), Inc.
2890 Scott Boulevard
Santa Clara, CA 95050
www.necsam.com