WHITE PAPER

Virtualizing a Windows Active Directoy Domain Infrastructure

vmware®

Table of Contents

Introduction

This white paper focuses on best practices for implementing an enterprise Active Directory infrastructure using virtualization technology. We will discuss how to perform a full migration of your physical AD infrastructure to a virtualized environment. Many of the topics are also relevant if you plan to support a mixed environment.

There are different methods for this type of implementation; however, it is important to understand certain behaviors. As in any virtualized environment, results may vary. Factors that affect the outcome include network speed, hardware performance, system redundancy and virtual machine (VM) density policies (that is, the number and type of virtual machines per server).

Why Virtualize Active Directory?

There are several reasons why you may want to virtualize Windows Active Directory, including hardware consolidation and standardization, improved efficiency in test and development, and better security profiles. Virtualizing Active Directory domain controllers will also allow you to take advantage of the many features of VMware Infrastructure, such as disaster recovery and planning, high availability, and optimized resource utilization.

Hardware Consolidation and Standardization

Virtualization can help administrators consolidate and reduce the amount of hardware in their environments. Such consolidation leads to lower server costs, lower datacenter costs and better resource utilization—in short, fewer bricks and less mortar. A virtual environment offers wide flexibility, as well. It is easy to mount virtual machines on underutilized single-use servers, transforming them into multi-use servers with more efficient utilization rates. It is also easy to move virtual machines to widely disparate legacy boxes without concern for the host's hardware configuration.

Another advantage of VMware is that the guest OS is presented with a standard set of hardware, which can eliminate potential imaging problems. Common issues associated with imaging software, including disparate hardware and SIDs, are all virtually eradicated. Where changing physical hardware might require reactivating Active Directory domain controllers, virtualization makes it easy to maintain a static set of virtual hardware, reducing the need for reactivation.

Test and Development

A virtualized environment facilitates the testing of different configurations without affecting your production environment. For example, you can test and evaluate group policies, practice product upgrades, and experiment with different migration or upgrade strategies. You can also test application development that requires modification to the schema data structures to be sure there are no anomalies, before they are used in a production environment.

Virtualization enables the testing of alternative domain configurations without affecting production. Furthermore, the virtual environment is excellent for testing deployment, helping you to uncover hidden problems before attempting an actual "live" deployment. Using virtual machines, you can also test either existing or new disaster recovery plans.

Security Control

In a virtual environment, fewer physical servers are deployed, so there is less overall security risk at the physical layer, for example with the VMware ESX. (For more details, see the white paper "Security Design of the VMware Infrastructure 3 Architecture," available on www.vmware.com.) New security features in VMware ESX 3.x and VMware VirtualCenter 2.x enable you to apply tighter control over .vmdk files and their data contents.

VMware VirtualCenter offers role-based access control to allow a granularity of administrative delegation that extends and enhances Active Directory's normal Organizational Unit delegation. VirtualCenter also gives smaller organizations with limited physical resources the ability to keep applications separate from Active Directory databases, enhancing security even further.

Migrating to the Virtual World

There are some virtual machine configurations to consider during the transition to a virtual environment.

The Virtual Machine on 64-Bit Windows Server

If using the x64 version of Windows Server 2003 or 2003 R2, one of the primary goals will be to contain the entire Active Directory database within the virtual machine's RAM cache. On 64-bit Windows, employing 16 GB of RAM cache will accommodate a database of approximately 2.5 million users.

Caching the Active Directory database in 64-bit Windows will avoid performance hits related to certain disk operations. For a virtual machine that is a domain controller, adding, modifying, searching, deleting and update operations generally benefit significantly from caching. Write operations will always incur a slight penalty, regardless of whether a domain controller is running on a physical or virtual machine.

There is limited benefit for filling cache on 32-bit Windows for customers with large directories; in fact, in some cases this actually can exhaust kernel resources.

Transitioning from Physical to Virtual

VMware supports the following Microsoft operating systems as virtual machines: Windows 2000, Windows 2003, and R2, all in a native mode or mixed-mode environment. These operating systems are fully functional as domain controllers in an Active Directory infrastructure. As with any other application, there are a number of steps to perform as you migrate Active Directory from the physical to virtual world:

1. Make sure you start with a solid system state backup of your Active Directory database, using your current backup process.
2. Maintain at least one physical Active Directory running all required infrastructure services until you have transitioned all physical servers to virtual machines.

Caution: Never attempt to recover an Active Directory database from a backup copy of an old virtual disk.

3. Consider creating a dedicated virtual switch or virtual machine port group to isolate replication traffic. Creating an isolated environment for the newly created domain controller virtual machines will exempt the virtual machines from contending with any other VMs during database replication.
4. Building a single-processor virtual machine operating as a domain controller should be more than sufficient for performance as a production domain controller. As an infrastructure application, a domain controller tends to use less than 10 percent of CPU resources. Multiprocessor virtual domain controllers generally do not increase their performance linearly. Generally, multiprocessor virtual machines tend to restrict VMware ESX server CPU scheduling flexibility.
5. Create separate, appropriately sized virtual disks for the Active Directory database, logs, and SYSVOL. This is consistent with Microsoft's best practices, and prevents Active Directory from contending with the operating system's boot disk. Using Microsoft's Active Directory Sizer tool can help provide basic sizing estimates of the database and the .vmdk file sizes required for deploying Active Directory. (See the Microsoft Web site for more information on the use of this tool.)
6. Use the Replication Monitor (replmon.exe) from the Support Tools to check, validate and initiate the Knowledge Consistency Checker, guaranteeing all the connectors are in place for proper replication.
7. Give the Active Directory replication process plenty of time. Depending on the size of your Active Directory forest(s), allow 24-48 hours for complete replication.
8. Change the weight and/or priority of the DNS records for virtual machines. This will allow the administrator to direct logon requests so they are more likely to go to virtual machine domain controllers rather than physical boxes. Also, use Performance Monitor to view logon requests and ensure that the DNS records are performing correctly. Active Directory Sites and Services is an alternative method for accomplishing this task, but it requires users to reboot computers to negotiate their assigned subnet.
9. Once you are satisfied with the performance of the virtual machines, decommission the physical domain controllers.

Best Practices for Virtualizing Active Directory

With any Windows OS, there are several steps to ensure that your virtualized Active Directory implementation is a success. These steps include:

1. Controlling clock drift
2. Optimizing network performance
3. Making DNS modifications correctly
4. Replicating database information
5. Providing virtual machine access control
6. Ensuring disaster preparedness and high availability
7. Handling disaster recovery

General Best Practice Guidelines

As you proceed, keep these general best practices in mind:

- Avoid snapshots or REDOs for domain controller virtual machines.
- Do not suspend domain controller virtual machines for long periods.
- Perform consistent and regular system state backups.
- Never attempt to recover an Active Directory database from a backup copy of an old virtual disk.

Planning for Accurate Time

Virtualized machines can easily (and fairly rapidly) drift if they are not receiving constant and consistent time cycles. Windows operating systems keep time based on interrupt timers set by CPU clock cycles. In an VMware ESX host with multiple virtual machines, CPU cycles are not allocated to idle virtual machines. Therefore, to plan for an Active Directory implementation, you must carefully consider the most effective way of providing accurate time to domain controllers.

It is important to understand the relationship between the time source used by clients, member servers and domain controllers. Figure 1 shows the source for the various entities. The PDC Emulator is a Flexible Single Master Operations (FSMO) role that advertises itself as the primary domain controller (PDC) to workstations, member servers and domain controllers running earlier versions of Windows. Only one PDC Emulator in each domain in the forest contains this Flexible Single Master Operations (FSMO) role.

The Domain Controller with the PDC Emulator role for the forest root domain ultimately becomes king, and is the "master" timeserver for the forest--that is the root time server for synchronizing the clocks of all Windows computers in the forest. You can configure the PDC to use an external source to set its time. By modifying the defaults of this domain controller's role to synchronize with an alternative external stratum 1 time source, you can ensure that all other DCs and workstations within the domain are accurate.

Controlling Clock Drift

In a virtualized environment, virtual machines that don't require CPU cycles don't get CPU cycles. For the typical Windows application in a virtual machine, this is not normally a major problem. When virtualizing Microsoft's Active Directory, however, these idle cycles can cause significant time drift for domain controllers. Active Directory is critically dependent on accurate timekeeping, and one of the most important challenges you must address is how to prevent clock drift. In fact, a large part of a successful Active Directory implementation will be in the proper planning of time services.

Microsoft uses Kerberos v5 as the authentication protocol, and time synchronization is critical to the proper functioning of this protocol. The time-stamped authentication tickets generated by Kerberos are based on the workstation's time, and only a five-minute tolerance is allowed. If time drifts significantly, the authentication tickets will not be issued, become outdated or simply expire, resulting in denied authentications or an inability to log into the domain to access network resources.

Accurate timekeeping is also essential for the replication process in a multi-master directory environment. Active Directory replication uses Update Sequence Numbers (USNs) to determine what changes to pull from its replication partners. Domain controllers only look for records newer than their own. If simultaneous changes made on different domain controllers result in identical USN numbers, timestamps are used as one form of tie-breaker.

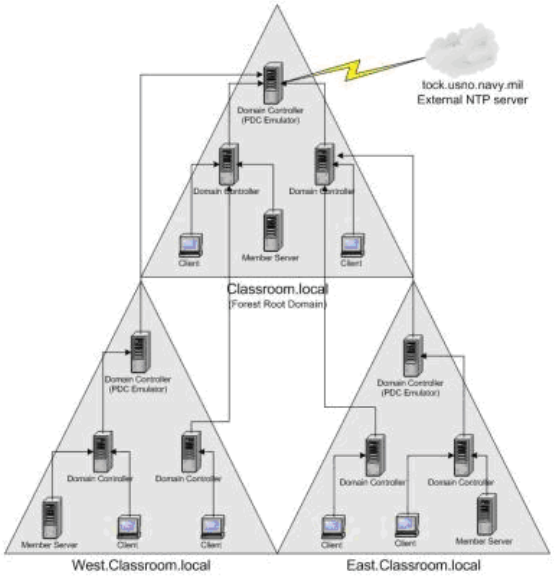

Figure 1. Synchronizing the forest root domain PDC Emulator to an external time source ensures all other DCs will be accurate.

Below is thesuggested way to support this time-sensitive environment in a virtualized infrastructure.

Using Windows Time Service for Synchronization

The first option is to use the Windows Time Service and not VMware Tools synchronization for the forest root PDC Emulator. This requires configuring the forest PDC emulator to use an external time source. The procedure for defining an alternative external time source for this "master" time server is as follows:

1. Modify Registry settings on the PDC Emulator for the forest root domain:

   In this key:

   HKLM\System\CurrentControlSet\Services\W32Time\Parameters\Type

- Change the Type REG_SZ value from NT5DS to NTP.

  This determines from which peers W32Time will accept synchronization. When the REG_SZ value is changed from NT5DS to NTP, the PDC Emulator synchronizes from the list of reliable time servers specified in the NtpServer registry key.

  HKLM\SYSTEM\CurrentControlSet\Services\W32Time\Parameters\NtpServer

- Change the NtpServer value from time.windows.com,0x1 to an external stratum 1 time source—for example, tock.usno.navy.mil,0x1.

This entry specifies a space-delimited list of stratum 1 time servers from which the local computer can obtain reliable time stamps. The list can use either fully-qualified domain names or IP addresses. (If DNS names are used, you must append ,0x1 to the end of each DNS name.)

In this key:

HKLM\System\CurrentControlSet\Services\W32Time\Config

- Change AnnounceFlags REG_DWORD from 10 to 5.

  This entry controls whether the local computer is marked as a reliable time server (which is only possible if the previous registry entry is set to NTP as described above). Change the REG_DWORD value from 10 to 5 here.

2. Stop and restart the time service:

   net stop w32time

   net start w32time

3. Manually force an update:

   w32tm /resync /rediscover

   (Microsoft KB article # 816042 provides detailed instructions for this process.)

Optimizing Network Performance

Network performance is also very important to Active Directory. Slow or unreliable connections to your virtualized domain controllers can affect the authentication and replication process. Determining the level of availability and how you provide that connectivity is equally important.

You can use VirtualCenter 2.x Maps view to verify network infrastructure. This view allows verifying connectivity to appropriate NICs and switches, and validating network paths. By employing two virtual NICs per domain controller virtual machines and by putting domain controllers within the same domain on different virtual switches, you can guarantee that a single NIC failure will not isolate the domain controller virtual machines for that domain. VMware ESX 2.5.x/3.x also allows VLAN port groups to support your existing network infrastructure. VMware ESX 3.x provides several new innovative networking features that offer much greater control over the connectivity provided to your virtualized domain controllers. (Consult the VMware Best Practices guides on networking for information about developing a robust network infrastructure to support your virtual environment.)

Making DNS Modifications

The PDC Emulator FSMO role is very busy in an Active Directory infrastructure. In addition to playing the part of a domain controller and acting as the timekeeper for the domain, the PDC Emulator is responsible for processing password changes for its domain, authenticating failed password requests, and "emulating" a PDC for down-level servers such as NT 4.0 BDCs and clients. In addition, some legacy applications are still written to specifically contact the PDC of the domain.

By modifying the weight and/or priorities of the DNS SRV records, you can relieve the load on the PDC Emulator. Simply direct logon authentications to specific domain controllers or away from the PDC Emulator.

DNS Weight

DNS weight uses a proportional system to distribute the requests among servers. The weight is actually an arbitrary value assigned to DNS SRV records to balance or distribute authentication requests among the domain controllers. By default, the assigned value is 100; reducing this value changes the proportional value relative to other servers so that a server with a lower value receives fewer requests. For example, if a DNS SRV record is lowered to 25 or 50 from a default of 100, it means that server will receive authentication requests 25 or 50 percent of the time in proportion to the others.

DNS Priority

DNS priority allows the administrator to inflate the DNS SRV record to a value so high, artificially, that it would be unlikely to receive a request unless no others are available to respond. By default, the value is set at 0. Setting priority extremely high, say 100 or 200, significantly reduces the chances the server will get the request.

Adjusting Weight and Priority

To adjust the weight and priority in a PDC Emulator, add to the following key:

HKLM\System\CurrentControlSet\Services\Netlogon\Parameters

- Set the LdapSrvWeight DWORD to a decimal value of 25 or 50.
- Set the LdapSrvPriority DWORD decimal value to 100 or 200.

Note that registry changes may require a reboot. These changes can also be performed directly through DNS Manager by simply double-clicking on the record, then adjusting.

Using the weight and priority strategy is an excellent way to wean client requests away from the physical domain controllers and direct them to the virtual machine domain controllers. This will allow you to safely begin the decommissioning process of your physical domain controllers.

Replicating Database Information

In a multi-domain controller or multi-domain environment successful exchange of Active Directory database information with the replication partners is vital. Utilizing tools to determine paths of connectivity and site replication will ensure consistent and accurate databases.

Active Directory Replication Monitor (replmon.exe) provides detailed information about the state of affairs within your Active Directory infrastructure, such as information about connections, updates, replication status, FSMO roles and so forth. Additionally, it can graphically show the relationship between domain controllers. This tool indicates when there is a replication problem (see Figure 3).

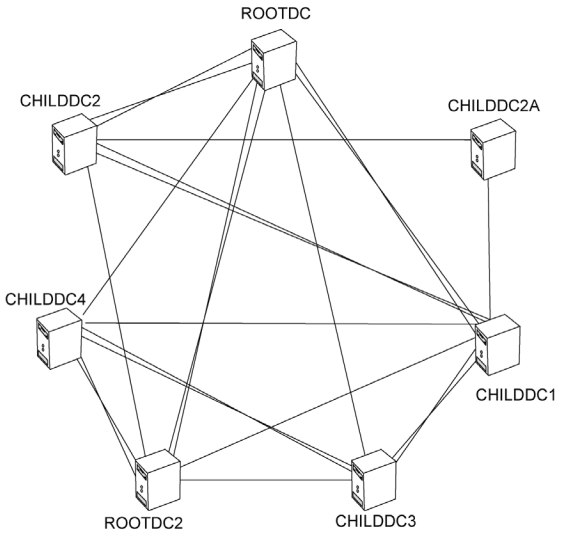

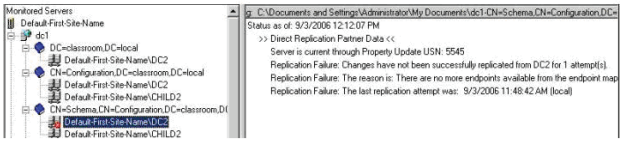

Figure 3. Replication Monitor shows where problems occur.

Providing Virtual Machine Access Control

Providing secured access to domain controller virtual machines is possible using access roles in VirtualCenter. You can give specified users and/or groups the level of privilege they require to maintain the domain controller virtual machines without affecting anything else within the VirtualCenter environment. New role-based access control features provide greater control over the virtual machine access.

Ensuring Disaster Preparedness and High Availability

Since many domain controller virtual machines may (and likely will) be running on a single VMware ESX, eliminating single points of failure and providing a high-availability solution will ensure rapid recovery.

VMware provides solutions for automatically restarting virtual machines. If a VMware ESX goes down, VMware High Availability (HA) can automatically restart a domain controller virtual machine on one of the remaining hosts, preventing loss of Active Directory. Using configuration options, you can prioritize the restart or isolation status for individual virtual machines.

VMware also allows you to specify a priority for restarting virtual machines (see Figure 4). For example, it is important for domain controllers functioning as global catalog servers to be online before your Exchange Server environment initializes. In any event, it is always a best practice to set your domain controller virtual machines as high-priority servers.

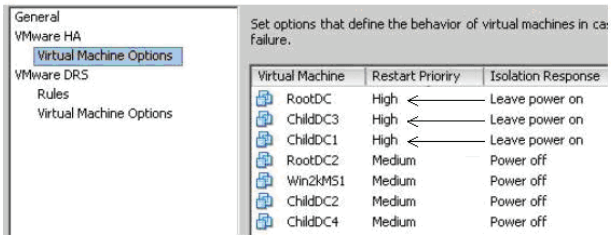

Figure 4. Specifying restart priority for virtual machines.

Additionally, you can implement a script to restart a virtual machine via a loss-of-heartbeat alarm through VirtualCenter. You can accomplish this using a script (available with the VI Perl Toolkit or the VMware Infrastructure SDK 2.0.1, both of which are located on the VMware Web site under Communities. Combined with VMware Distributed Resource Scheduler (DRS), you can ensure that domain controllers from the same domain always reside on different VMware ESX hosts to prevent placing all the domain controllers in one basket. The anti-affinity rules let you specify which domain controllers must stay together and which must be separated.

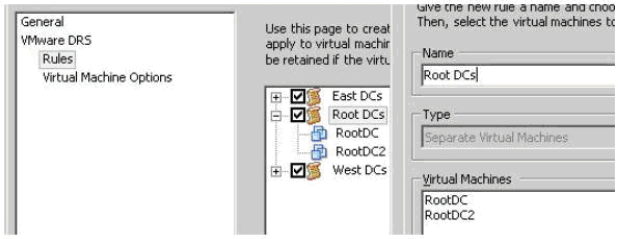

Figure 5. VMware DRS and anti-affinity rules ensure domain controller virtual machines are segregated.

Once again, to ensure database version and USN consistency, do not use snapshots or REDO disk modes for domain controllers. If the physical machine has data or other services that preclude you from avoiding migration altogether, transfer any FSMO roles, disable Global Catalog server, and demote the domain controller to a member server. Then migrate by creating a new virtual machine, promoting it to a domain controller, and transferring the roles to the newly created domain controller.

Handling Disaster Recovery

Developing strategies and plans of attack for disaster recovery is another important facet of virtualizing Active Directory. It is important to have (and practice) a routine for backing up system state data. You can use Windows ntbackup.exe or a third-party solution, but make sure it is capable of doing system state backups. Performing consistent system state backups eliminates hardware incapability when performing a restore, and ensures the integrity of the Active Directory database by committing transactions and updating database IDs. During the system state backups, these database IDs (InvocationIDs) and Update Sequence Numbers (USNs) are properly set and identified. Use Microsoft's KB article# 223346 as guidance on the placement of FSMO roles.

Perform all Active Directory restorations using both the authoritative and non-authoritative technique, as defined by Microsoft's procedures. For more information on backup and recovery procedures for Active Directory, visit the Microsoft Web site.

Because the hardware for all virtual machines is identical in VMware ESX, performing an authoritative or non-authoritative restore is uneventful, regardless of the hardware on which the virtual machine sits.

Never attempt to recover an Active Directory database from a backup copy of an old virtual disk. Figure 6 shows the implications of not properly restoring a virtual machine using Microsoft's method for an authoritative or non-authoritative restore.

On the left, restoring a virtual machine from an old copy of a virtual disk does not update the database IDs or the USNs. As you can see, the database invocation IDs are not synchronized. Unless administratively changed, the DC2 will never accept inbound replication of the changes, since it already has knowledge of the objects that reside in VDC1's Active Directory database. Domain controllers are only interested in changes that are newer than the database they contain.

Using the proper method ensures that missing transactions are restored correctly to the failed database once it synchronizes with its replication partners (see Figure 6).

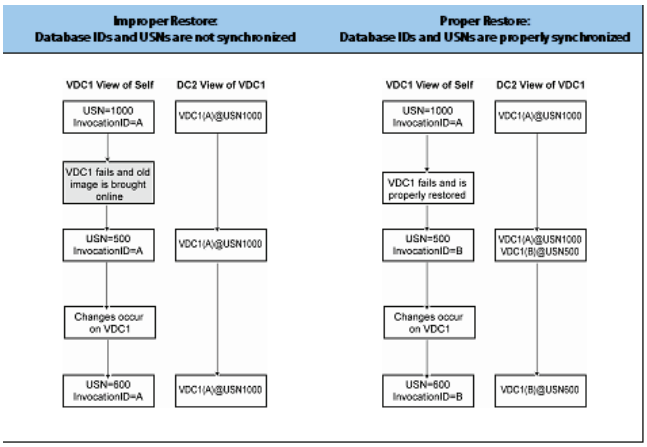

Figure 6. Illustrating an improper and a proper restore.

Note that Microsoft does not support using snapshots or disk imaging of any domain controller. You must perform system state data backups to protect the Active Directory database. Snapshots can potentially corrupt the Active Directory database without even committing or reverting back to the snapshot. (See Knowledge Base article: KB888794 on the Microsoft TechNet Web site.)

Conclusion

There are several excellent reasons for virtualizing Windows Active Directory. Virtualization offers the advantages of hardware consolidation, total cost of ownership reduction, physical machine lifecycle management, mobility and affordable disaster recovery and business continuity solutions. It also provides a convenient environment for test and development, as well as isolation and security.

For success in implementing Active Directory in the virtual environment, you must ensure a successful migration from the physical environment to the virtual environment. Since Active Directory is heavily dependent on a transaction-based datastore, you must guarantee integrity by making sure there is a solid, reliable means of providing accurate time services to the PDC Emulator and other domain controllers throughout the Active Directory forest.

Network performance is another key to success in a virtual Active Directory implementation, since slow or unreliable network connections can make authentication difficult. Modifying DNS weight and priority to reduce load on the primary domain controller assists can help improve network performance.

Because Active Directory depends on reliable replication, ensure continuity by using replmon to monitor your Active Directory. Also, continue regular system state backups, and always restore from a system state backup. Virtual machines make it easy to move domain controllers; use VMware HA and VMware DRS so that no critical domain controllers are on a single host. Practice the art of disaster recovery regularly. Finally, always go back and re-evaluate your strategies; monitor results for improvements and make adjustments when necessary.

About the Authors

Charles Windom

Charles Windom is a Microsoft Solution Architect for VMware. As a consultant, he has designed, implemented and provided support on Microsoft infrastructure and applications to small, medium and enterprise level companies. His special focus has been best practices for design and implementation of systems with Active Directory, Exchange, Microsoft Operations Manager (MOM) and SQL Server. Before joining VMware, Mr. Windom worked for 15 years for such companies as Brocade, Compaq, Siebel, Taos and Hitachi Data Systems, designing and implementing solutions for storage area networks (SAN), disaster recovery and business continuity, email archiving, continuous data protection (CDP), and Microsoft infrastructure and applications. Mr. Windom earned his B.S. in Computer Science and has been Microsoft Certified since 1990. He is currently a Windows Server 2003 Microsoft Certified System Engineer.

Chris Skinner

Chris Skinner is a technical trainer for VMware. As a consultant, he has implemented and trained employees on Microsoft's BackOffice operating systems and applications. Before joining VMware, he spent eight years as an independent consultant and contractor for many Fortune 1000 companies, providing advice and implementation of software network infrastructure solutions. He also spent two and half years working with the United States Air Force, standardizing network operating systems and enterprise applications, including Exchange Server, SMS and HP OpenView. Mr. Skinner earned his B.A. in Business Management and an Executive MBA degree. He has been Microsoft Certified since 1993 and currently holds his MCSE certification.

VMware, Inc. 3401 Hillview Ave Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com

vmware®