

At a glance:  
Containing costs  
Server inventory  
Determining CPU, disk I/O,  
and network utilisation  
Placing virtual machine  
guests to virtual machine  
hosts

# Essential tools for planning your virtual infrastructure

JAY SAULS AND BALDWIN NG

Virtualisation is not a new technology, but it's now exploding into many new areas, far beyond the original platform virtualisation developed by IBM in the 1960s. A major reason for this rapid

growth is the potential for virtualisation technologies to help contain costs in the data centre and across IT operations. With energy costs rising and the appetite for more IT services by different parts of your organisation ever increasing, you may be struggling to keep up with the server growth while holding down operations costs. Virtualisation can help. And Microsoft and other software makers have been developing new technologies, including server, desktop and application virtualisation, that will make virtual resources more widely available.

In this article we look at a toolset that can help you with your virtualisation projects – the Microsoft Assessment and Planning Toolkit (MAP), which is a member of the Microsoft Solution Accelerators family. What are Solution Accelerators? Simply put, they are automation tools and guidance you need for almost every IT project. Solution Accelerators are free, authoritative resources that help you proactively plan, deliver and operate your IT infrastructure – from desktops to servers, from physical to virtual.

Designed in collaboration with indus-



Figure 1 The Microsoft Assessment and Planning Toolkit 3.1

try experts and tested in real customer environments, Solution Accelerators offer just the right level of prescriptive guidance and automation to achieve high performance and build a cross-product IT infrastructure – all without having to reinvent the wheel. These resources focus on key concerns of companies of all sizes, such as infrastructure assessment and planning, virtualisation, communication and collaboration, and security and compliance.

**Tool highlights**

The newly released MAP Toolkit 3.1 (see Figure 1) is a network-wide infrastructure assessment tool designed to help you quickly gain a baseline understanding of your servers and desktops, including whether they are ready for upgrade or migration to technologies such as Hyper-V, Microsoft Application Virtualization, Windows Server 2008, Windows Vista, Office 2007 and SQL Server.

As Figure 2 shows, MAP autogenerated readiness-assessment reports and proposals, saving you a lot of work you'd otherwise have to do yourself. MAP Toolkit 3.1 supports both server and desktop scenarios, including evaluating desktop machines for upgrade to Windows Vista and planning for the rollout of Microsoft Application Virtualization.

For server-virtualisation planning, three key features of the tool can help you determine which of your physical machines are good candidates: server inventory, performance metric gathering and virtualisation placement recommendations. Let's walk through a scenario that shows how you can use these features to build a detailed picture of your current environment and plan for the transition to virtualisation.

**Find out what you've got**

Many large enterprises have some form of server-management infrastructure in place, such as System Center Operations Manager. Some environments are less actively managed. If you're working in an environment

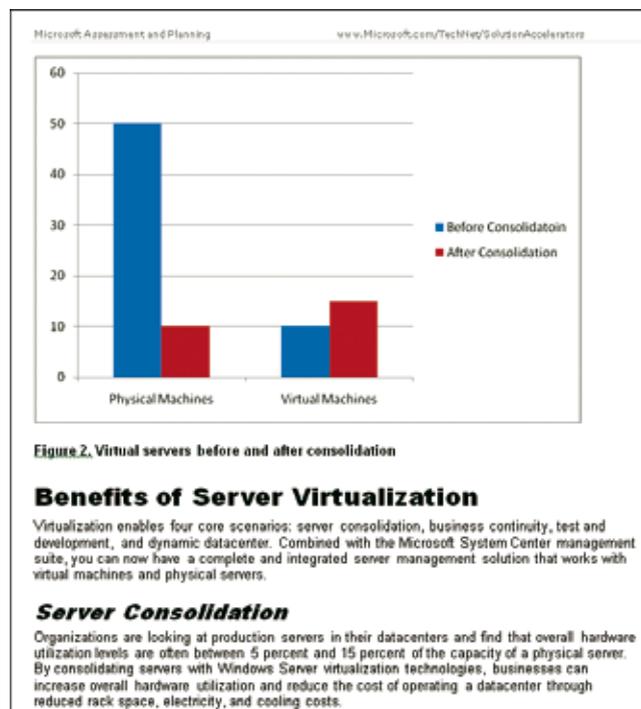


Figure 2 Sample auto-generated migration proposal

where you're really not sure what servers are in place or what they're running, the Server Inventory feature can quickly help you paint a clear picture. Server Inventory is able to discover machines in your environment using Active Directory or IP address range scan, or by simply importing a list of machines from a flat file. As Figure 3 shows, you can also use multiple discovery methods together (Active Directory and IP address scan, for example) to find machines that should be in Active Directory but aren't.

As machines are discovered, MAP Toolkit 3.1 attempts to connect to each one and retrieve more than 50 classes of data via Windows Management Instrumentation (WMI). Several different kinds of reports can be generated based on this data. For server consolidation with virtualisation, we're most interested in the reports created when you select Migrate specific roles and services to Windows Server 2008.

Even if you're not doing a migration to Windows Server 2008 right now, the two reports this scenario creates are valuable in determining the kind of hardware in your environment (including processor type and speed, disk sizes and memory installed), as well as what software roles you have assigned to each server (DNS servers, domain controllers and Terminal Services).

These reports are provided to you as Microsoft Excel® workbooks, so it's easy to use the sorting and filtering tools in Excel to

identify exactly which machines in your environment would make good candidates for virtualisation.

The server inventory reports provide data on several attributes you can consider when choosing servers to virtualise. As an example, you may want to ensure that the computers you're virtualising are running an OS supported by the Hyper-V virtualisation platform. In the WS2008HardwareAssessment Excel workbook created by the server-inventory process, under the Server Inventory tab, you can examine the Current Operating System and Service Pack columns to find any systems that may need to be upgraded.

**Determine current utilisations**

The next step in finding virtualisation candidates is to measure how heavily utilised the physical machines are. The Performance Metrics Wizard in MAP Toolkit 3.1 (see Figure 4) records performance counters from a list of machines you've identified as virtualisation candidates. CPU utilisation, disk I/O utilisation and network utilisation are all captured for later review.

This part of the tool was designed to watch machines over a fairly long period of time (up to a month), which gives you some flexibility in choosing how you use it. If you already have a good idea of when the highest utilisations in your environment occur (for example, you may have a weekly batch-processing window), you can set up the wiz-

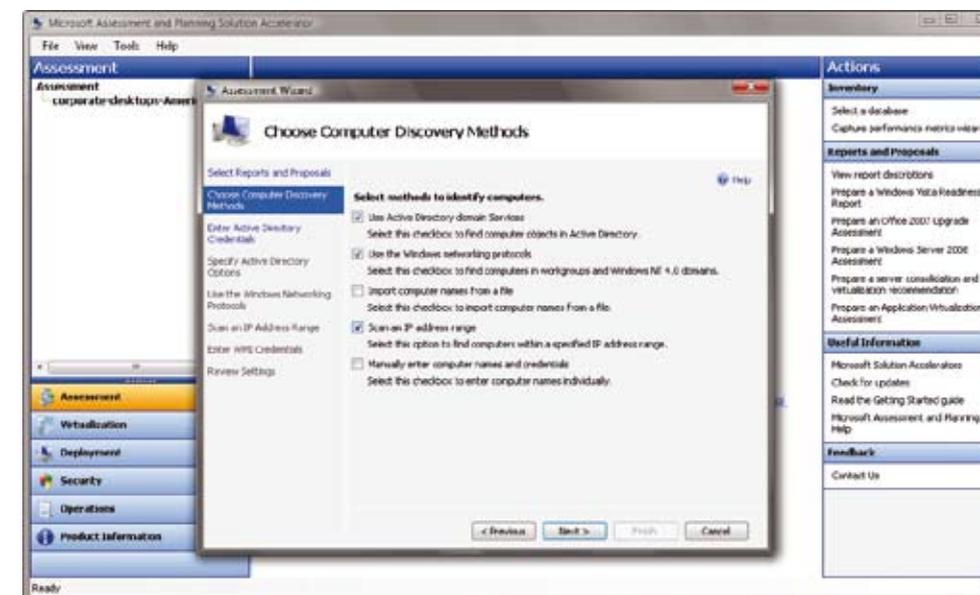


Figure 3 Choosing the computer discovery method

ard to run for only a few hours around that time. However, if you're unsure when your highest utilisations occur or if you'd like to observe many machines that have different times of peak use, you can set up the wizard to run for several days or weeks to record data from your computers. This lets the tool automatically identify peak utilisations from a larger set of data.

Using the information from your earlier inventory, access the sorting and filtering capabilities in Excel to create a list of machines you'd like to monitor. The Performance Metrics Wizard will accept this list as a flat file, with one machine name per line. Copy the list of machine names from Excel into a new, blank text file in Notepad, then save it. When prompted, import this file into the wizard and set the end time of the performance metrics gathering period. The wizard will go into data-collection mode and gather performance counters from the specified machines every five minutes.

When the Performance Metrics Wizard completes, it will create a detailed Excel report listing all the machines observed and all the performance metrics for different categories of utilisation: CPU, physical disk, logical disk and network. Using this information, you can further refine your list of candidates to consider for virtualisation.

You can target machines with extremely low utilisation and look to consolidate several of them as virtual machines (VMs) on a

single host, or you can combine this information with the inventory data to find machines that have mission-critical services that need high availability. Virtualisation can be used to provide enhanced disaster-recovery capabilities for high-availability environments. After your analysis is done, you can move on to the final step of generating virtual machine placement recommendations.

### Server placement recommendations

The Server Virtualization and Consolidation Wizard allows you to specify the configuration of a hypothetical virtualisation host machine in terms of CPU, network, disk subsystem and memory. Using the performance data collected by the Performance Metrics Wizard, the Server Virtualization and Consolidation Wizard computes a good mapping of VM guests to VM hosts, ensuring that each host is well utilised without risking overutilisation during peak usage.

When the Server Virtualization and Consolidation Wizard completes, it produces a detailed Excel spreadsheet with the proposed mapping of guest VMs to physical hosts. For each VM and host, the report shows projected disk I/O, network I/O and CPU utilisation. As guest VMs are added to physical host machines during the fitting process, new host machines are allocated as previous hosts reach capacity. For each new host, a message is displayed that indicates why a new host had to be allocated.



Figure 4 The Performance Metrics Wizard

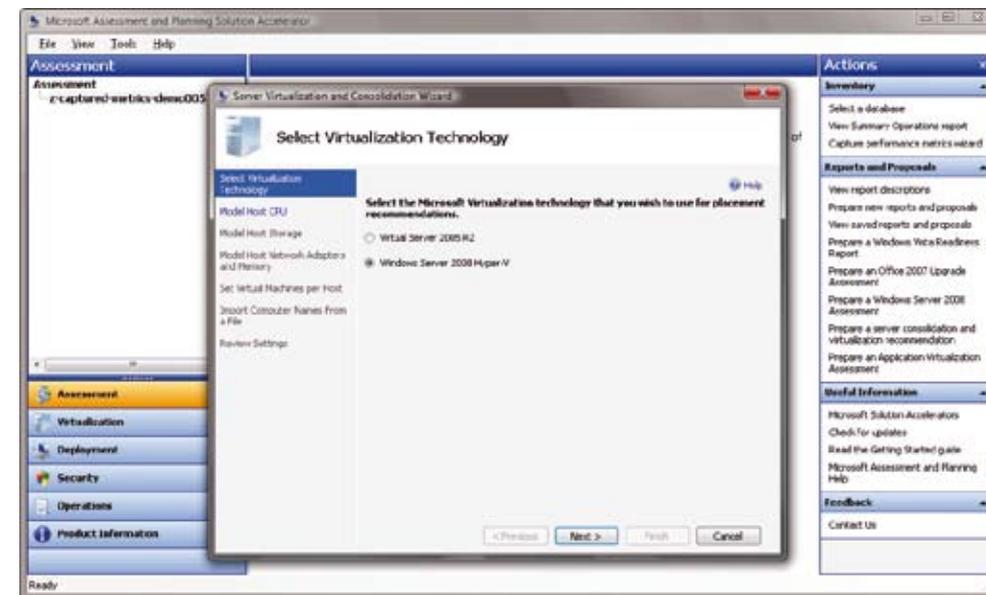


Figure 5 Selecting the virtualisation technology to use for placement recommendations

If several of your hosts get the same message, you might consider rerunning the Placement Wizard with a different host configuration to try to increase the number of guests assigned to each host. For example, if several of your hosts display the message that they reached the limit on memory, you could try rerunning the wizard and specifying a larger amount of RAM for the host configuration. This should increase the number of guests that can be assigned to each host.

### Virtual machine configuration

Each page in the Server Virtualisation and Consolidation Wizard allows you to configure a major component of the hypothetical computer configuration that will be used for modelling the capacity of the VM host computer. Models are provided for both Virtual Server 2005 R2 and Hyper-V, and you'll be asked which virtualisation technology you want to use on the first page of the wizard (see Figure 5).

Models for several families of CPUs from both AMD and Intel are provided on the CPU configuration page. If a model for your exact CPU isn't available, choose the closest match. You can configure the count of physical processors, cores and hyper-threads per core on this page. Note that if you choose Hyper-V as your virtualisation platform, only 64-bit CPUs will be available to choose from (see Figure 6).

The third page in the wizard lets you con-

figure a disk subsystem in terms of number of disk drives and RAID technology to compute both total disk space available on the host and a value for the throughput of the disk subsystem. The throughput value is provided automatically based on the interface technology and the disk Reliability and Performance Monitor (RPM) speed chosen.

In the MAP Toolkit 3.1, direct support for modelling storage area network (SAN) and network attached storage (NAS) disk solutions is not provided. For a rough approximation, choose the SCSI disk technology and configure a RAID level that's similar to your SAN configuration, with disks that are

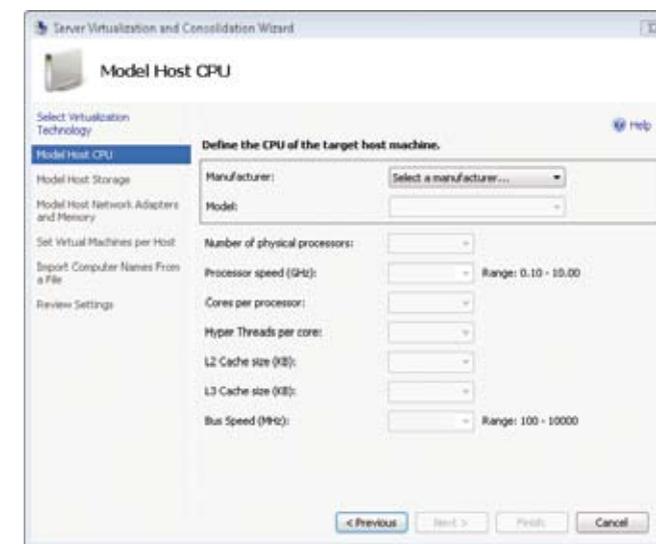


Figure 6 Host CPU configuration

of similar speed. This approximation should be close enough in order to help you do your initial planning.

#### Placement constraints

As virtualised guest computers are mapped to physical hosts, six factors are taken into consideration to evaluate whether or not a guest will fit within the remaining capacity of a host computer:

- CPU utilisation
- Network I/O utilisation
- Disk subsystem I/O utilisation
- Memory
- Disk space
- Guest OS support on the specified virtualisation technology

CPU, network and disk subsystem utilisation are fairly straightforward. Measured utilisations on the physical virtualisation guest candidate are converted to an equivalent utilisation on the virtualisation host computer, which is then added to the existing utilisation of the host. If the total utilisation of all guests assigned to that host for that resource is less than the threshold limit, then the guest is assigned to the host. If not, the guest won't be placed and other guests are tried to see if one can be found to fit.

As an example, suppose you have gathered performance metrics about a computer with a 1.2 GHz Pentium 4 CPU that showed a 32 percent CPU utilisation. If you configure a virtualisation host with a quad-core Intel Xeon 5300 series CPU, the original measured CPU utilisation from the potential guest computer will be extrapolated to an equivalent projected utilisation on the host computer. This new utilisation will be lower than the original measured value owing to the newer CPU architecture, different CPU clock speed and extra cores.

Conversely, if you configured a host computer with a less powerful CPU than the one in the guest, the projected utilisation will be higher once extrapolated to the host. In either case, the new extrapolated CPU utilisation will be used in calculating the overall usage of the host computer as guests are assigned to it. It's worth briefly discussing just how the utilisation of memory and disk space is decided. The total amount of physical

memory installed on each potential virtualisation guest is reserved on the host machine. Additionally, a small amount of memory (512MB) is reserved for the host itself. When the total memory on the host is exhausted owing to the assignment of guests and the host overhead, no more virtualisation guests may be placed on that host computer.

Disk space is handled similarly: the total amount of physical disk space on each potential virtualisation guest is reserved on the host machine. Additionally, 50GB of disk space is reserved for use by the virtualisation host. Once the total disk usage of all the guests assigned to the host plus the space reserved for the host itself reaches the size specified in the host configuration, no more guests can be assigned.

Based on the virtualisation platform chosen, checks will be made to ensure that the OS running on the guest computer is actually supported by the virtualisation platform. For example, 64-bit OSs are not supported as guests on Virtual Server 2005 R2, so they will not be placed if you choose that as your virtualisation technology.

#### Considering your options

By running the Server Virtualization and Consolidation Wizard multiple times, it's easy to compare the results of different hardware choices for the virtualisation host machine. Each time you run the wizard, you can choose different CPU, disk and memory options to find the solution that's going to be best for your environment.

Pay attention to the Last Placement Failed Reason column on the Consolidation Recommendations tab of the server consolidation Excel workbook. These reasons will give you clues as to why more guest machines could not be placed on a particular host. If several hosts are displaying the same reason, such as "Memory utilisation of guest exceeds available memory on host", increasing the amount of memory on the host configuration may reduce the total number of hosts that are necessary.

#### Requirements

One of the advantages of MAP Toolkit 3.1 is that it does not require installation of any software on the computers being invento-

ried. All operations are conducted remotely. However, there are a few factors that must be taken into account in order for everything to work properly.

In order for the Inventory Wizard and the Performance Metrics Wizard to complete their work, a few TCP ports must be open on the target computer and accessible from the computer where the MAP Toolkit is installed. The Getting Started Guide goes into detail about the requirements; in short, the ports for authentication, DCOM and WMI must be open on the computers being inventoried. If there are any firewalls installed on the network between the computer running MAP Toolkit 3.1 and the inventory targets, these firewalls must be configured to allow this traffic to pass as well.

MAP Toolkit 3.1 is designed to inventory environments of up to 100,000 computers and to collect performance metrics from up to 3,000 computers at a time. By default, this tool will install a private instance of SQL Server 2005 Express Edition. Due to the 4GB database file size limitation in SQL Server Express, this configuration only allows you to inventory about 20,000 computers before hitting the file-size limit enforced by SQL Server Express. In order to work with more computers during inventory, there are some defaults that must be modified.

To work with larger environments, you'll have to use a full version of SQL Server, such as SQL Server 2005 Standard Edition. Before you install the MAP Toolkit, install a new instance of SQL Server 2005 and name the instance "MAPS." Then run the MAP Toolkit installer as usual on the same computer. The MAP Toolkit requires a local instance of SQL Server, and the installer will automatically detect a local "MAPS" instance of the database server and install into it, instead of installing SQL Server Express Edition.

For even larger environments, you can use the scale-out feature, which allows multiple helper computers that have the MAP Toolkit installed to execute a large inventory cooperatively and store all the data gathered into a single central database. This configuration also requires the use of SQL Server Standard Edition to avoid the 4GB file size limit.

See the Getting Started Guide for details on how to install this tool in this configura-

## MAP Toolkit has saved hundreds of hours of IT planning time for Hyper-V deployments

tion and to edit the database connection strings appropriately for this mode. The scale-out feature can also be used to put SQL Server on a machine other than where the MAP Toolkit is installed.

#### Wrap up

Solution Accelerators such as the Microsoft Assessment and Planning Toolkit are all about doing the heavy lifting for your IT projects so you do not have to. The MAP Toolkit has saved many users hundreds of hours of IT planning time for Hyper-V deployments. Are you ready to give this a try?

If you would like, you can download the Microsoft Assessment and Planning Toolkit free at <http://microsoft.com/map>, or you can get other Virtualization Solution Accelerators at <http://microsoft.com/vsa>. We also encourage you to visit the MAP Toolkit Team Blog at <http://blogs.technet.com/mapblog> to find more tips and view demo videos. ■

---

**JAY SAULS** is a senior program manager in the Solution Accelerators group at Microsoft, where he has worked on designing and implementing tools such as Windows Vista Hardware Assessment and the Microsoft Assessment and Planning Toolkit. Previously Jay worked on other Microsoft product teams, including Microsoft Commerce Server, Microsoft BizTalk and System Center Capacity Planner.

**BALDWIN NG** is the senior product manager for the Microsoft Solution Accelerators Core Engineering Team. His current projects include the Microsoft Assessment and Planning Toolkit and Virtualization Solution Accelerators. Baldwin has more than 16 years of product-management and development experience in IT, e-commerce and wireless Internet services working with Microsoft, Corbis and Alltel. He frequently blogs at [blogs.technet.com/mapblog](http://blogs.technet.com/mapblog)