Dynamic memory is an enhancement to Hyper-V R2 which pools all the memory available on a physical host and dynamically distributes it to virtual machines running on that host as necessary. That means based on changes in workload, virtual machines will be able to receive new memory allocations *without a service interruption* through Dynamic Memory Balancing. In short, Dynamic Memory is exactly what it's named.

Let's dive in an explain how all this works starting with the new Dynamic Memory settings. Here are the new settings available on a per virtual machine basis. Here's a screenshot:

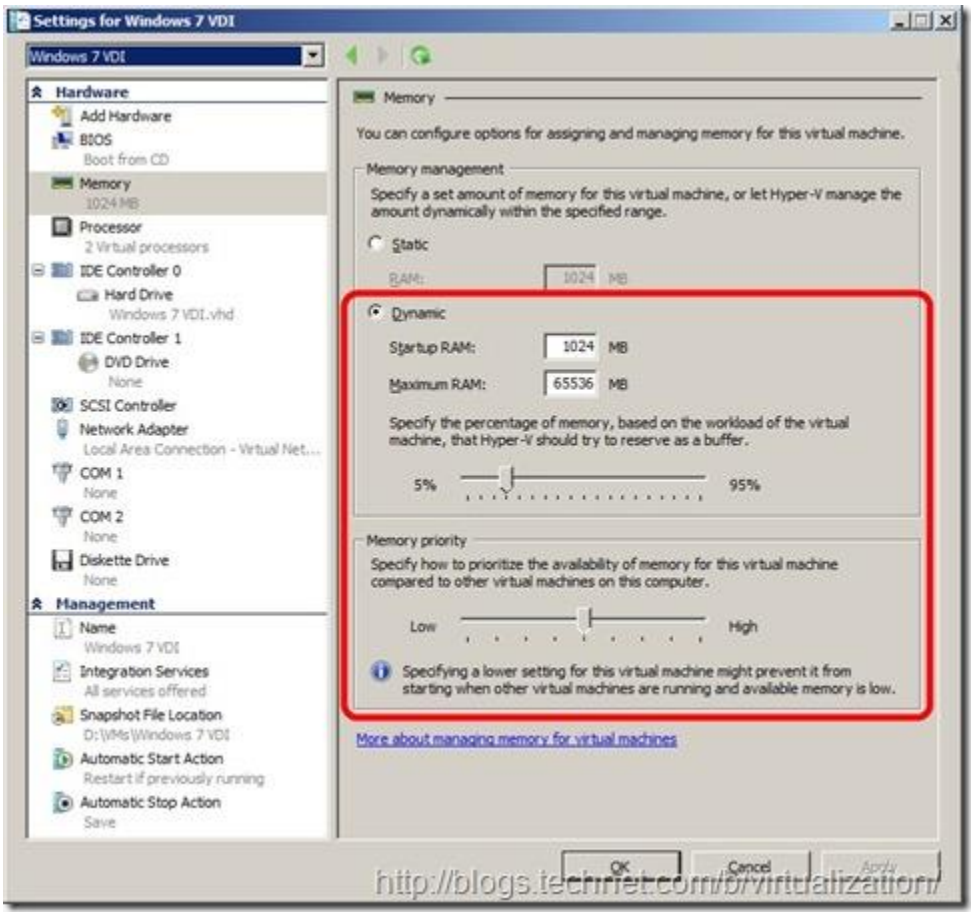

Dynamic Memory In Depth

With Hyper-V (V1 & R2), memory is *statically* assigned to a virtual machine. Meaning you assign memory to a virtual machine and when that virtual machine is turned on, Hyper-V allocates and provides that memory to the virtual machine. That memory is held while the virtual machine is running or paused. When the virtual machine is saved or shut down, that memory is released. Below is a screenshot for assigning memory to a virtual machine in Hyper-V V1/R2:

With Hyper-V Dynamic Memory there are two values: **Startup RAM** and **Maximum RAM** and it looks like this:

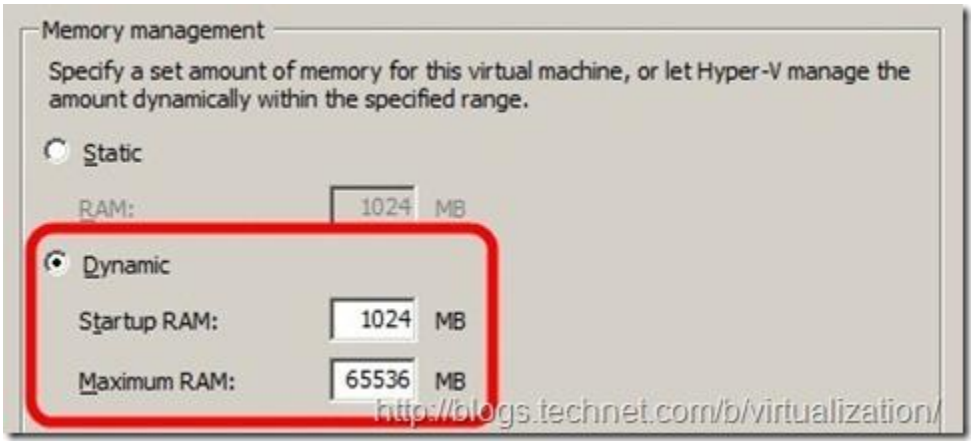

Startup RAM is the *initial/startup* amount of memory assigned to a virtual machine. When a virtual machine is started this is the amount of memory the virtual machine will be allocated. In this example, the virtual machine will start with 1 GB.

The **Maximum RAM** setting is the maximum amount of memory that the guest operating system can grow to, up to 64 GB of memory (provided the guest OS supports that much memory). Based on the settings above, here's an example of what the memory allocation *could* look like over a workday...

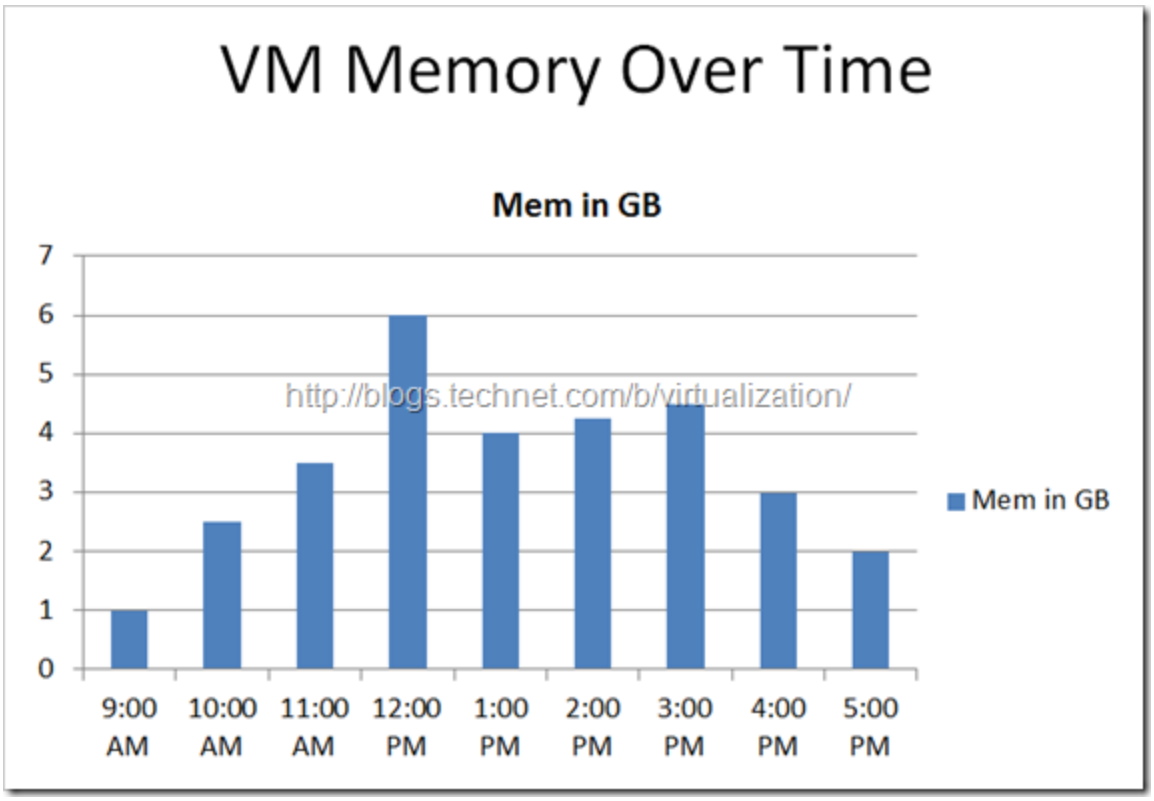

As you can see, the workload is dynamically allocated memory based on demand.

Next, let's look at the Memory Buffer.

Memory Buffer: In one of the earlier blogs posts in this series, we discussed the complexity of capacity planning in terms of memory. To summarize, there is no "one size fits all" answer for every workload as deployments can vary based on scale and performance requirements. However, one consistent bit of feedback was that customers always felt more comfortable by providing additional memory headroom 'just in case.'

We completely agree.

The point being you want to avoid a situation where a workload needs memory and Hyper-V has to start looking for it. You want some set aside memory as buffer for these situations, especially for bursty workloads.

The Dynamic Memory buffer property specifies the amount of memory available in a virtual machine for file cache purposes (e.g. SuperFetch) or as free memory. The range of values are from 5 to 95. A target memory buffer is specified in percentages of free memory and is based on current runtime memory usage. A target memory buffer percentage of 20% means that in a VM where 1 GB is used, 250 MB will be 'free' (or available) ideally for a total amount of 1.25 GB in the virtual machine. By default, Hyper-V Dynamic Memory uses a default buffer allocation of 20%. If you find this percentage is too conservative or not conservative enough, you can adjust this setting on the fly while the virtual machine is running without downtime.

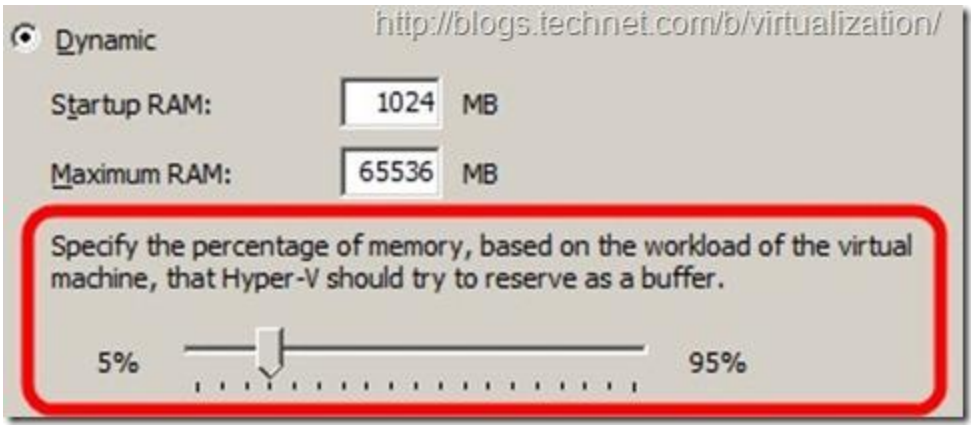

This takes us to the last Dynamic Memory setting, Memory Priority.

Memory Priority: By default, all virtual machines are created equal in terms of memory prioritization. However, it's very likely you'll want to prioritize memory allocation based on workload. For example, I can see a scenario where one would give domain controllers greater memory priority than a departmental print server. Memory Priority is a per virtual machine setting which indicates the relative priority of the virtual machine's memory needs measured against the needs of other virtual machines. The default is set to 'medium'. If you find that you need to adjust this setting, you can adjust this setting on the fly while the virtual machine is running without downtime.