HITACHI
Inspire the Next

Microsoft Hyper-V Live Migration over Distance

WebTech Q&A Session – January 20, 2011

1. **The environment you were addressing in your presentation reflected the Hitachi Data Systems USP VM storage system and other networking components. Will this solution apply to other Hitachi Data Systems controllers as well? How about other storage vendors? Or other networking components?**
   Yes, we're currently validating support for both the Hitachi Adaptable Modular Storage 2000 family and the latest enterprise class system, Hitachi Virtual Storage Platform. While other storage vendors may offer a somewhat similar configuration, only HDS storage systems can leverage the power and functionality of the Hitachi Storage Cluster for Microsoft® Hyper-V™ technology – a significant industry differentiator. Other networking components might also be utilized, but the Brocade equipment used in the Hyper-V Live Migration over Distance Solution offer unique capabilities that provide the best solution for the infrastructure presented.

2. **We've been talking about the test and lab environments. Can you give us insight into production experiences?**
   Several Hitachi customers are using Hitachi Storage Cluster for Microsoft Hyper-V in production. One example is MAAF, a leading public insurance company for private individuals and professionals in France. According to a MAAF executive, "We have invested in Microsoft Hyper-V virtualization to optimize our enterprise technology footprint. Combining this investment with Hitachi Storage Cluster for Microsoft Hyper-V has amplified the value of our investment."

3. **What purchases for storage allocation are required? What about when LUN are shared between clients?**
   Shared storage between cluster nodes is one of the core requirements for Microsoft Failover Clusters. There are no special requirements beyond this.

4. **Does the scope include testing over non-dedicated fiber? What speed connection was tested?**
   Yes, we utilized spools of raw fiber between the locations, connected through Ciena 4200 DWDMs. Speed is a function of the laser strength and the optics on either end of the fiber run, just like when using "dark fiber" between two company data center locations.

5. **As the volumes of simultaneous users is scaled upward from 200, what are the impacts to the failover?**
   The number of simultaneous users doesn't directly impact a live migration (which is a different operation from an unplanned virtual machine or node failover). The factors involved include physical and virtual machine RAM configurations, processor configurations and speed and, most importantly, network bandwidth between the two locations. The process of live migration transfers the memory pages of a running VM across the Ethernet network. Therefore, user load is typically not a factor (beyond the amount of changes currently in memory during a live migration.