GIGANET SYSTEMS

Precision • Performance • Repeatability

How To Test VMotion Technologies - The Network Emulator

Application Note

Introduction

Data Center virtualization is fast becoming a key technology today. While it isn't really a new concept, the level of attention being paid to it has certainly increased in line with trends in a couple of key areas like Disaster Recovery and Cloud Computing.

For their part, Disaster Recovery and business continuity solutions require geographical dispersion of the data centers hosting the business data and applications. Un-interrupted availability of these applications and systems requires that data and applications move across these data centers in an automatic and reliable manner. Accomplishing these goals typically requires replication of data and applications and naturally increasing the deployed server and storage hardware.

The physical space, power, and cooling requirements for duplicating typically underutilized dedicated servers would seem to make these solutions cost prohibitive for all but those with the deepest pockets. Compounding this challenge, the additional equipment and facilities not only increases the general management burden, but also significantly increases the complexity of the systems' architecture.

VMotion and **SVMotion** technologies address these challenges. These key virtualization tools provide the ability to transparently move Virtual Machines (VMotion) and Storage systems (SVMotion)) across data centers and allow organizations to get more value out of their existing infrastructure through the consolidation of existing under-utilized physical servers into fewer servers and facilities running at much higher capacities. Not only do organizations benefit from greatly reduced hardware, power, cooling, and overhead costs, but virtualization also greatly simplifies the infrastructure, architecture, and operation of the environment, further reducing costs to the organization.

Of course, VMotion and SVMotion are complex technologies and their use places additional burdens on the Ethernet Networks and Fibre Channel SANs interconnecting the data centers. Successful migration of live Virtual Machines (VMs) requires predictable operation of networking and computing infrastructure in conjunction with management software under both normal and error conditions. System downtime or failures can be disastrous both economically and from a business information perspective.

Before an organization can realize the benefits afforded by these technologies, the must work flawlessly. Their products and services must be thoroughly tested under real-world conditions to ensure that the end users' trust in them is justified. Fortunately, GigaNet Systems has a simple solution.

Emulate and Test

In an ideal world, everything would be plug-and-play straight out of the box. New products and services could be activated with the simple click of a button and everything would work exactly as advertised. Unfortunately, this is just not how it works today. More often than not, new equipment or services must interact with existing equipment or services and the results are often surprising. Clearly, the risks inherent in introducing an untested or unproven product or service into a live or existing network are just too high. **Pre-deployment validation of system operation and performance is a must.**

The seemingly obvious solution is to test using the actual live network where the product or service is to be used. Unfortunately, this is highly impractical for the reasons listed above, as well as for the fact that conditions in a live network are unpredictable. Interpreting results in this environment is challenging at best and, with no control over the conditions for any given test case, repeatable testing is impossible.

What if you could build a test bed that would act like a real-world network but provide complete control and repeatability? You can, with a network emulator – the right network emulator.

The right network emulator provides the test user with the ability to mimic the available bandwidth, delay and other impairment conditions of a real-world network, and evaluate the performance, stability, or functionality of a product or service as it would function once deployed. Moreover, it does so with precision and performance to ensure that all traffic and impairments are processed EXACTLY as intended. For these reasons, there is only one choice – a hardware-based network emulator.

Hardware-based network emulators use custom-designed and dedicated Integrated Circuits (ICs) to process impairments and traffic. Unlike low-end emulators, they possess the processing power to deal with higher volumes of traffic and multiple impairments simultaneously.

GigaNet Systems' VirtualNet line of 1G/10G Ethernet Network Emulators and VirtualSAN line of 4x/8x Fibre Channel SAN Emulators are hardware-based in-line (pass-thru) network emulators using dedicated FPGAs to deliver full line-rate performance for all traffic at all frame sizes, AND regardless of the impairment settings applied by the user. With the ability to replicate real-world network conditions in a 100% controlled manner, users are guaranteed that no other variable has been introduced and results are precise and, most importantly, repeatable.

With a VirtualNet or VirtualSAN Emulator from GigaNet Systems, test engineers can confidently focus on any number of Virtualization specific test areas including among others:

1. Selection of proper optimization techniques
2. Performing SLA compliance verification
3. Validation of system response to additional latency introduced by VM traffic routing
4. Identification and evaluation of various storage strategies prior to deployment.

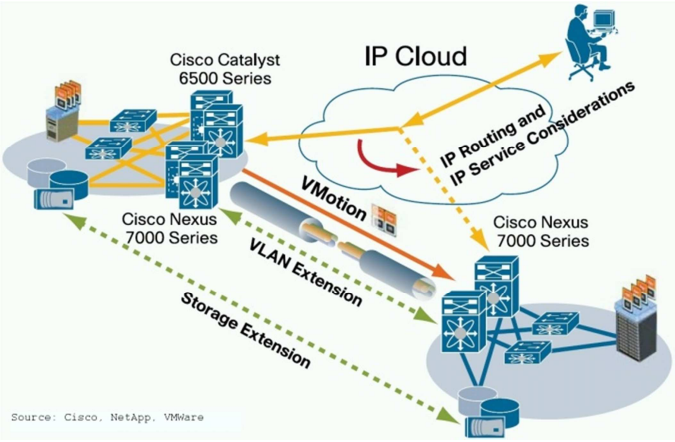

VM Transition Time

While transitioning between Data Centers, a VM is temporarily unavailable to the user. In order to ensure compliance with SLAs, VM migration success rate and transition times must be known. Since VM migration times are highly dependent on network latency, available bandwidth, and bit error rates (resulting in frame drops), tuning of algorithms and equipment choices can have a drastic impact on transition times. GigaNet emulators will help optimize the system and SLAs.

VM Traffic Routing

After migration, a VM maintains its existing IP address. Traffic to the migrated virtual machine traversing a Layer 3 network needs to be re-routed to the new data center. Existing application sessions may have to continue to be routed through the original data center due to specific or existing IP service requirements, such as firewalls. In addition to moving large data sets during the VM transition, the network infrastructure must also cope with the persistent load of forwarding traffic between data centers. Each of these scenarios introduces latency and the impact must be characterized and addressed. GigaNet emulators provide users with the ability to validate system response to latency and many other commonly occurring network impairments.

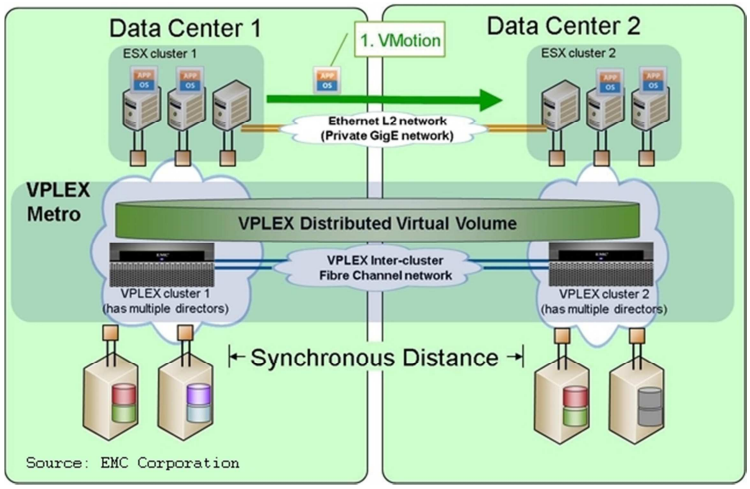

VM Storage Extension

When a VM transitions to a new data center, the storage data may remain at the original site. In this scenario, network latency and available bandwidth has a direct impact on application performance. Where there is a large distance between the data centers, the I/O access latency between them may be prohibitively large. Furthermore, Fibre Channel can suffer from a buffer credit starvation problem as the latency (distance) between the data centers is increased. With geographically dispersed data centers, sophisticated approaches such as caching or symmetric replication (SRDF) may be required. Using GigaNet VirtualNet/VirtualSAN provides testers with the ability to confidently select the right storage strategies for their needs – based on the exact conditions they will face.

Key Emulator Requirements for Virtual Testing

To properly evaluate the areas above requires an emulator not only with a complete range of impairment capabilities, but also one that provides extreme granularity and functionality.

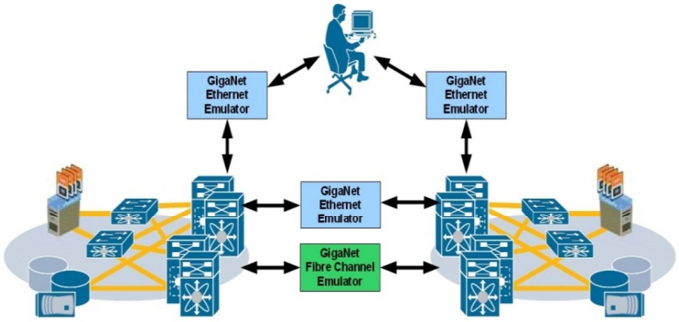

GigaNet Systems' line of Ethernet Network and SAN Emulators' comprehensive list of impairments (and their importance to (S)VMotion testing) includes (among many others):

Delay The importance of latency cannot be underestimated. From the characterization of transition times to the impact of data re-routing to the simulation of geographically dispersed data centers, latency plays a critical role in the development, selection, and deployment of VMotion and SVMotion architectures, products and services. VirtualNet and VirtualSAN give users the power to precisely control delay - providing up to 10s of delay in increments as little as 6.4µs.

Jitter (Packet / Delay / PCR) Jitter can have a profound effect on system or service performance depending on the mix of components involved. Developing a clear picture of which system algorithms are able to transparently address variable delay and which systems prefer a static delay will help guide the product-architecture mix decision. Building on the precision of the delay functionality, delay jitter will range between the Min. and Max. Delay as set by the user.

Bandwidth Control Bandwidth controls on the network emulator are used to set the maximum amount of traffic that may be sent via the impairment path. As with delay (latency), the amount of time it takes to transfer the Virtual Machine will depend heavily on the bandwidth available between the two locations. Use Bandwidth Controls to restrict traffic and increase Out-of-Sync issues to characterize system response. Restricting bandwidth can also help define and validate SLAs as the longer transmission times between VMs will impact the availability of the new VM. Of course, GigaNet emulators also include traffic shaping so as to avoid unexpected traffic bursts when attempting to restrict bandwidth.

Frame Drop Frame Drop simply emulates frame/packet loss and allow the user to drop selected traffic (randomly or targeted) at a specified rate, duration, and distribution. Dropping packets can be used to simulate errors and evaluate everything from how the system's error correction mechanisms work to characterizing the performance hit on overall system throughput or availability. Thorough testing will ensure maximum system performance under adverse conditions as equipment, settings, and configurations are individually optimized.

Reorder Frame reorder impairment emulates situations where frames/packets arrive out of order at the destination from situations such as routing table updates, route flaps, or because of TCP retransmission. From a test perspective, users want to know what happens when re-ordered packets arrive too far out of order. How does the system handle it? Does it delete/drop it? Does it recover it in its original place? What is considered "too many" re-ordered packets? Selection for frames to trigger the reorder event can be done either randomly or via user-defined criteria specified as an equation or both. In addition, user can specify the duration of the reorder impairment.

Corruption Frame Corruption introduces bit errors into the user data traffic to emulate data corruption during transmission. From a VM Testing perspective, this function would be used to corrupt a packet and determine whether or not the system is able to detect the error and discard it as expected. If not, then the integrity of the machine and data must be questioned. Corruption impairment controls allow the user to target specific frames, set the corruption duration, the Bit Error Rate (BER), error burst length, and corruption envelope (i.e. which part of the frame is to be corrupted).

Additionally, GigaNet Systems' emulators allow the user to identify specific checksums to be recalculated after corruption. This feature proves especially beneficial when testing a system supporting encrypted data. Testers can use VirtualNet or VirtualSAN to corrupt the encrypted data and then re-compute the standard Layer 2 network checksums. From a security perspective, the system under test should still be able to identify the corrupt packet and discard.

*Note: Checksum Correction is a **very** intensive function and Software-based emulators see significant performance degradation. Only VirtualNet offers Checksum Correction at both 1G and 10G speeds.*

Line BER This impairment emulates the bit errors that occur in normal networks due to cable / fiber defects, sub-optimal optical power, or low-rate errors that can occur due to random events. When evaluating the hardware to be used in a Virtual environment, Line BER becomes a critical test as users seek to evaluate the actual performance of the product(s) under consideration. How does the hardware respond to adverse conditions? Does it recover? Does it interoperate well with the other systems in my architecture? Using VirtualNet or VirtualSAN Line BER controls, users can answer these questions by setting BER Duration, Rate, Burst Length, and Mode (how to corrupt or change a targeted bit).

Loss of Signal (LOS) As its name implies, this particular VirtualNet and VirtualSAN impairment emulates a Loss of Signal (LOS) or fiber cut conditions. This feature is used to test system response and recovery to the LOS condition. Typical Data Center architectures will include a redundant path for availability and security. This feature enables the user to simulate a Loss of Signal and validate that the hardware switchover to the redundant link happens and that the signal converges. Both of these must happen in a finite period of time to ensure no Timeout occurs. Of course, everything must also be transparent to the higher layers.

Individually, each of the features described above provide some solid capabilities for evaluating Virtual Machine and Storage Data Center products and systems. Together, however, they provide a comprehensive and powerful toolkit for the virtual service provider to understand the complexities and nuances of their services when it comes to such things as:

- Evaluating network infrastructure equipment
- Determining trade-offs between competing technologies such as caching and Symmetric Remote Data Facility (SRDF)
- Optimizing Network and Storage Architectures
- Establishing Availability Metrics

Conclusion

Data center virtualization is a complex application. Poorly executed solutions will result in commercial failure and can even cause irreparable damage to the hardware vendor or data center service provider's reputation. Successful deployment requires the transparent interoperability of the multiple components that form an end-to-end network. Not only does that mean testing, but it requires testing with the right product – a hardware-based network emulator.

GigaNet Systems portfolio of Ethernet and SAN Emulators are the right products. No longer do you have to perform testing "in a vacuum" which only provides best-case performance numbers.

GigaNet Systems network emulators allow you to perform precise real-world testing of virtualization technologies

With VirtualNet and/or VirtualSAN in the test environment, engineers can introduce bit-errors, delay, bandwidth controls, frame drop, corrupt frame data, modify specific protocol fields in selected frames, and much more to observe the resulting system response. This testing allows for validation of system performance and robustness under both everyday usage scenarios as well as specific corner cases to prevent costly post-deployment operational and performance issues.

These network emulators are in-line devices making them an easy addition to new or existing test setups and have a comprehensive feature set for performing detailed and dynamic real-world testing of products and systems supporting virtual environments.

Whether pre- or post-deployment, the goal of any testing is to establish and maintain confidence in the product or service being delivered. GigaNet's VirtualNet and VirtualSAN deliver the industry's greatest performance and precision to ensure that your tests, and your results, are repeatable and reliable – so that your product or service will perform ***EXACTLY*** as expected when deployed.

Better Testing → Better Reliability → Better Products