F5 White Paper

Virtual Clustered Multiprocessing (vCMP)

Clustered Multiprocessing (CMP) technology introduced by F5 in 2008 enabled organizations to consolidate physical, purpose-built resources into a single virtual entity that provides near 1:1 scaling of performance by simply adding or upgrading resource blades. Virtual Clustered Multiprocessing (vCMP), the industry's first purpose-built hypervisor, completes this functionality by allowing administrators to completely segment those resources into independent, virtual ADCs.

by KJ (Ken) Salchow, Jr.
Sr. Manager, Technical Marketing and Syndication

Contents

Introduction 3

Consolidation Considerations 3

How Conventional Solutions Fall Short 4

Multi-Tenancy 4

Virtual Appliances 4

Virtual Clustered Multiprocessing 6

True, Purpose-Built Hypervisor 6

Deep Integration 6

The Payoff 7

Conclusion 9

Introduction

Data center consolidation and virtualization have changed the way organizations look at CapEx and OpEx. Gone are the days when adding new capacity or applications was simply accomplished by buying "more." Today, CIOs and architects are looking to maximize the return on investment in hardware and software through virtualization technologies that enable them to squeeze every ounce of computing power from their existing data centers.

This is most apparent in the world of application servers, but the potential benefits for other devices, firewalls, routers, and Application Delivery Controllers (ADCs) cannot be ignored. Consequently, most vendors offer strategies around multi-tenancy or virtual appliances in one form or another to provide the same kind of flexibility for their solutions that OS virtualization offers in the server world.

While both multi-tenancy and virtual appliances improve organizations' deployment flexibility and their ability to get maximum ROI from both CapEx and short-term OpEx, these strategies have failed to provide the same kind of high-reliability, high-performance solutions as traditional purpose-built systems.

Until now. F5® Virtual Clustered Multiprocessing™ (vCMP) technology, coupled with Clustered Multiprocessing™ (CMP) technology, application delivery software, purpose-built hardware, and virtual edition (VE) solutions, finally gives organizations a complete, end-to-end virtualization strategy for application delivery.

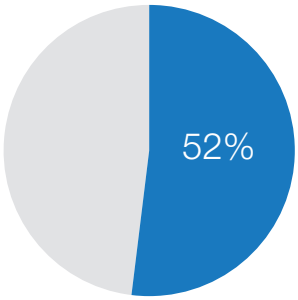

52% of surveyed customers consider consolidation to be an important feature for managing their applications and networks.

Source: TechValidate Survey of 104 F5 BIG-IP users
TVID: 8B7-806-7ED

Consolidation Considerations

As organizations began consolidating their data centers, they quickly realized that to succeed, they had to eliminate all excess; simply moving remote systems into a single data center quickly consumed all the space, power, and cooling available. Additionally, many of those systems used less capacity than was provided by the hardware they were deployed on. Many organizations also routinely over-provisioned their server hardware to accommodate future growth or unpredictable demand. At an individual application level this makes complete sense; however, this strategy replicated hundreds of times across hundreds of applications resulted in massive quantities of unused resources simply wasting away in the data center. Data center appliances also tended to be over-provisioned for the same reasons.

Reclaiming these stagnant resources to minimize the amount of rack space, power, and cooling needed emerged as a critical goal of data center consolidation.

However, it is still important to account for future growth and unpredictable load. On paper, commercial virtualization solutions seem to fit the bill, but each has weaknesses. These solutions come in two general forms: multi-tenancy and virtual appliances.

How Conventional Solutions Fall Short

Multi-Tenancy

Many appliance vendors use multi-tenancy to segment their solutions and provide unique management and operation for disparate groups. Through administrative partitioning, an organization can configure a device to service multiple customers or business units without those customers realizing that others are also using the same physical solution. While this provides the appearance of separation between tenants, the reality is that all of them share the same hardware resources. Therefore, if any one customer misconfigures their portion of the system, or causes excessive use of those resources, it can negatively affect other users. Multi-tenant solutions provide advanced controls to reduce this possibility (such as processor, memory, and bandwidth limits); however, the fact that it still remains possible makes these solutions unfavorable in many cases.

Multi-tenancy has other limitations: a single hardware failure affects multiple customers; customers must use the same versions of software, which limits flexibility; and while individual customers only see their unique portion of the system configuration, the overall configuration of the appliance includes all the configurations, making it complex and difficult to manage. The result is that while multi-tenancy achieves many of the goals of consolidation and can reduce CapEx, it tends to increase OpEx over the long term.

Virtual Appliances

Virtual appliances are also used to address the requirements of data center consolidation. A virtual appliance takes the software that once ran on dedicated, purpose-built hardware and ports it into a virtual machine that typically runs on a general-purpose hypervisor in the same manner that application server virtual machines do. They can run on commodity hardware and be moved in the event of failure or resource exhaustion. This provides some significant benefits not addressed

by multi-tenancy. First, each virtual instance is completely independent from any other virtual instance on the same host—whether it's another virtual appliance or an application instance. While the virtual appliance still shares the same hardware with other virtual appliances, the separation is much more distinct and controlled than in multi-tenancy. Second, if the underlying hardware fails or another instance causes excessive resource utilization, the virtual appliance can be moved to another hardware host fairly simply and with little interruption. This helps address the issue of hardware failures affecting multiple customers and resource limitations that multi-tenancy generally cannot.

Unfortunately, when it comes to virtual appliances that do significant amounts of heavy processing, many organizations are finding that they can't provide the same level of service as their physical counterparts, especially in consolidated environments where increased traffic accompanies the consolidation effort. The dedicated hardware that many of these appliances were originally deployed on was highly specialized and designed to boost performance beyond what general-purpose systems are capable of. For example, when appliances had to process real-time network traffic, the software was tuned for specific network interface hardware; but in a virtualized deployment, it can only be tuned to the virtualized interface, as the physical interface is never certain. In addition, processing real-time data encrypted by SSL/TLS requires dedicated, special-purpose hardware to keep from adding latency to the network, particularly as SSL keys change from 1024-bit to 2048-bit (requiring five to seven times the processing power). While it is certainly possible to deploy SSL processing hardware on commodity servers, it must be deployed on every machine that the virtual appliance might move to, which reduces the CapEx savings of the consolidation effort.

Even without these concerns, running virtual appliances on the same hardware/hypervisor as virtualized application instances can be dangerous as their need for processor, network interface, and memory often starve other applications. This results in virtual appliances often being deployed as the only virtual machine on an individual, physical piece of hardware much like traditional non-virtualized versions, but with the overhead of a hypervisor as well.

While both multi-tenancy and virtual appliances provide additional value over traditional, single-use appliances, both have drawbacks that limit their overall benefit. Multi-tenancy lacks complete isolation, fault tolerance, and ease of migration; and virtual appliances lack the performance of dedicated, purpose-built hardware. The ideal solution would provide all these features.

Virtual Clustered Multiprocessing

Clustered Multiprocessing (CMP) technology introduced by F5 in 2008 enabled the consolidation of physical, purpose-built resources into a single virtual entity that provides the most scalable Application Delivery Controller (ADC) to date. It's the only solution that provides near 1:1 scaling of performance by simply adding or upgrading resource blades.

Virtual Clustered Multiprocessing (vCMP) is the industry's first purpose-built hypervisor—it allows the complete segmentation of those purpose-built, scalable resources into independent, virtual ADCs.

True, Purpose-Built Hypervisor

All hardware isn't equal—organizations need purpose-built hardware for highly reliable, high-performance solutions. The same can be said for hypervisors. Most general-purpose hypervisors are designed to handle the broadest set of possible physical hardware configurations and guest operating system requirements. This flexibility is what makes them powerful enough to manage the complexities of today's modern data center that runs the gamut of hardware and software combinations.

ADCs, however, are not general-purpose computing platforms. While virtualized ADCs running on general-purpose hypervisors and commodity hardware have immense value for application-specific profiles and intelligence, they cannot provide the scale, high availability, and performance required for core ADC traffic management. These require purpose-built solutions.

The vCMP hypervisor is specifically designed for F5 hardware, which was designed to meet application delivery requirements. vCMP was also built to host F5's application delivery software, not to host general-purpose software. It is an application delivery hypervisor uniquely tuned and purpose-built for that purpose and none other. Unlike other vendors' solutions, which deploy general-purpose hypervisors and attempt to tune them for best-effort performance as an ADC hypervisor, vCMP is designed from the ground up to be an ADC hypervisor.

Deep Integration

The deep integration between hardware, hypervisor, and ADC software provides immense benefits over other solutions, most notably in efficiency and reliability.

Because F5 has complete control over the underlying hardware and the overlaying software, the vCMP hypervisor is more efficient in design and operation than general-purpose hypervisors. vCMP needs to implement support only for F5-specific hardware, and it needs to support only the processes required by F5 software. The result is a much more streamlined and efficient hypervisor design.

This more efficient hypervisor also improves reliability. Its streamlined nature makes the vCMP hypervisor easier to maintain and easier to troubleshoot. In addition, because F5 controls the entire stack, changes in supported hardware or software that require modification of the hypervisor are also controlled entirely by F5, rather than third-party developers or hardware manufacturers.

The result of this deep integration is a more stable and controlled platform that simply performs.

The Payoff

The payoff of a purpose-built hypervisor that's deeply integrated with the underlying hardware and guest software is the most powerful virtualized ADC solution available today. With vCMP, organizations can independently operate virtual instances without sacrificing interoperability with existing equipment, purpose-built hardware, or orchestration solutions.

With vCMP, administrators can run multiple instances of TMOS®, each isolated from the others. Unlike some implementations, because vCMP is a true hypervisor, the guest ADCs are completely isolated—so they can run entirely different versions of ADC software. This means that test and development staff can create new virtual ADC instances to test new versions of software without any effect on existing deployments. Or, competing business units can choose if/when they upgrade their virtual instances to meet their unique business requirements. All they have to do is provision a new instance, apply their existing configuration, and then test the upgrade process and results. Any problems can be addressed by simply removing the instance and starting over. Alternatively, administrators can upgrade individual instances in place without having to upgrade all instances.

Because each guest is its own complete ADC, individual business units or other customers have complete control over their deployment, the ability to further segment their deployment using administrative controls, and the ability to manage independent logs and configurations. However, a failure or misstep cannot affect any other virtual instance. Rebooting the instance, runaway processes, and flat-out misconfigurations are isolated from all other instances.

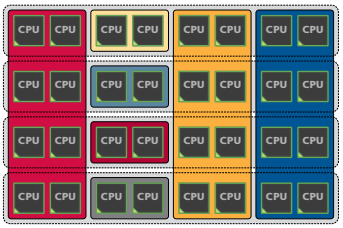

When vCMP is deployed on a multi-blade VIPRION® system, administrators can configure the vCMP instances of ADCs in multiple ways. The instances may exist on single blades, or they may span multiple blades. In addition, since VIPRION chassis allow new blades to be added on the fly, vCMP instances can be configured to automatically expand across the new blade resources without interruption. The primary benefits of vCMP are:

- Instances can run different versions
- Network isolation
- Resource isolation
- Fault isolation

All exist within a single VIPRION chassis (with the benefits that provides), and each instance is completely independent from the others. On-demand scale with complete control over resource allocation is the best of both worlds.

The deep integration of vCMP also enables it to work seamlessly with existing functionality. For instance, CMP allows new compute resources to be added incrementally and become instantly available to the ADC. When vCMP is in operation, those new resources can be automatically allocated to existing virtual instances without any interruption, reboot, or reconfiguration. On the other side of the stack, when configuring vCMP guest allocation, the hypervisor can directly assign IP addresses for management and VLAN tags along with the resource allocation restrictions. Creating a new ADC instance can be done in a matter of minutes, and a new administrator can log in and start their configuration. Other vendors' virtual ADC solutions require reboot of virtual instances before new resources are available, and each instance must be manually configured before being ready for further configuration. vCMP allows virtual instances full access to new network interfaces, VLANs, and even entirely new resource blades instantly and without interruption.

Flexible allocation allows administrators to designate CPU resources (and blades on chassis models) to guests upon creation. Dynamic scaling allows reallocation of CPU resources, without disruption. This makes it possible to redistribute resources to better align with the need for business agility in addressing growth and scale, as well as support additional or new application delivery services that may require more CPU resources. Administrators can size guests according to what's required for each deployment—and modify when those requirements change.

Because the hypervisor is purpose-built, it can also be integrated into existing data center orchestration solutions. This allows organizations to dynamically create and provision new ADC instances as part of their existing orchestration solutions. Using iControl®, an F5 management control plane API, administrators can spin guests up and down programmatically using a variety of cloud management platforms and frameworks, as well as custom solutions. Programmatic control of ADC instances enables elastic application network services and reduces operational overhead. This integration is identical to, and utilizes the same methods as, integrating with individual ADC devices (physical, virtual editions, or vCMP instances). In addition, general-purpose hypervisors have very limited network options and can't support the virtual MAC addresses required by VRRP (RFC 3768) and VLAN groups (proxy ARP bridge). vCMP can act as any standard network device; it can act as a router, a bridge, or a proxy ARP device as needed. This simply isn't possible with general-purpose hypervisor use.

Conclusion

Today's consolidated data centers require flexibility and agility. CIOs and architects are looking for solutions that provide the best return on investment despite rapidly changing requirements and business needs. Yet traditional solutions for virtualizing and consolidating hardware devices don't meet these requirements. Multi-tenancy solutions fail to provide sufficient flexibility by not allowing isolation of instances, and although they often reduce CapEx, they actually increase long-term OpEx. Virtual appliances provide flexibility but at the cost of performance and scale.

True agility is only gained with complete flexibility in resource deployment. The ability to consolidate without giving up individual resource isolation is critical. The flexibility to utilize virtualization without forfeiting the performance of dedicated, purpose-built hardware is necessary. vCMP, combined with CMP, F5's purpose-built hardware, virtual editions (VEs), and F5's industry-leading TMOS-based solutions provides the most complete set of options for application delivery with the greatest flexibility and agility. This allows customers to choose the right solution for their unique needs while helping to drive maximum return on their investment in consolidation and virtualization.

F5 now delivers true agility through limitless growth as well as flexible allocation—all without sacrificing performance, reliability, or flexibility.

F5 Networks, Inc. 401 Elliott Avenue West, Seattle, WA 98119 888-882-4447 www.f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com