African Journal of Basic & Applied Sciences 7 (3): 160-168, 2015
ISSN 2079-2034

DOI: 10.5829/idosi.ajbas.2015.7.3.22301

Massive Structure Challenges in an Internal and External Resource for Virtualized Datacenter

[1]*Ganesan Veerappan and* [2]*C. Suresh Gnana Dhass*

[1]Bharatiyar University, Coimbatore, Tamilnadu, India
[2]Department of CSE, Vivekanandha College of Engineering for Women, Tamil Nadu, India

Abstract: Today the Data Center is the breath of most organization for drives with pumps the data to and from users, storage devices and the World Wide Web. The rising demand for structure and storage computation has determined the growth of large data centers. The huge server farms that run many of present corporate and Internet business data centers whereas structures can include many thousands of servers and can use as much external and internal resources and additional utility to a small constituency. However the massive structures consist of an infrastructures, environments, designs, energy, deployments, managements, security services and reliability/scalability required to drive datacenter efficiently. But each structure results have unique classifications which are devolves to many testing and interesting. In this paper we discuss huge challenges about area of each structural classifies and provide the new generation in virtualized data center. In particularly, emphasis on why and how we move to the enhanced virtualization technologies.

Key words: Virtualization · Datacenter structure · Infrastructure · Management · Energy · Security

INTRODUCTION

Data centers have evolved over the past several decades from single point, concentrated processing centers to dynamic, highly distributed, rapidly changing, virtualized centers that provide myriad services to highly distributed, global user populations. Enterprise, Internet (transaction) and business applications are ever more being moved to large data centers that hold huge server and storage clusters. Current data centers can contain tens of thousands of servers and plans are already being made for data centers holding over a million servers [1]. The traditional, inflexible and hierarchical model of separately provisioned and maintained server, storage and network resources constrains organizations from cost-effectively providing on-demand support for applications and meeting unprecedented service levels.

Data centers run the applications that provide critical information and rich, differentiated content for users. Users now demand an agile, responsive infrastructure that provides exactly the access that they need. From the infrastructure perspective compute, storage and network for dynamic allocation in an automated, orchestrated and logically diversified environment, accommodating a variety of applications. Using orchestration, resources can be pooled within and across multiple data centers to provide an environment that responds dynamically to user needs through virtualization. When designing the data center network, we must consider all communications occurring within the data center itself, between the data center and its users. Energy consumption and sustainability are further domains in data centers where interest and innovation are emerging [2]. In Many ways we could adopting the best practices for energy efficiency in common data center operations.

Next, data center operators must deal with the deployment and planning problems related to estimating a data center's capacity and initial provisioning for new applications [3]. All the segment of data centre is being addressed by virtualization technologies. Physical resources are also being abstracted by virtualization and delivered that in a logical form so the resources are utilized more efficiently and managed through deployment [4]. Manageability means simply knowing the task if a server or other network element is up or down. Particularly, in service provider data centers are supporting multiple customers, the ability to assess service levels on a customer basis are essential to the

Corresponding Author: C. Suresh Gnana Dhass, Department of CSE, Vivekanandha College of Engineering for Women, Tamil Nadu, India.

offering and administration of service-level agreements (SLAs). Security is also an issue, in the sense of shared data center environments and leading to much research on isolating services and building trusted platforms [5, 6]. There is no second thought for security when it comes to optimum network design. No option to Trial-and-error networking, corporate lifeblood network could be compromised as a single vulnerability. In such a emergency settings are acceptable by only specific expertise. The virtual data center network, access point mainly connected to server cluster for reliable access. Scalability must be afforded in every data center. Server load balancing has basic rule and techniques such as reverse proxy caching. All the complications and available solutions increase greatly if the virtual data center is part of a geographically dispersed set.

Data Center Infrastructure: Data centers have grown in popularity as the processing power required by businesses exceeds; what they can maintain within their own corporate infrastructure. Today's data centers involve the rapid growth of compute capacity and increased data volume, demand a high-performance and scalable data center network infrastructure. To achieve a fundamental improvement in both user experience and economics for the next-generation data center, we need to develop universal architecture along with enhanced virtual capability, then only eliminates complexity by collapsing the data center network into a single tier.

Infrastructure Challenges: Nowadays we are facing many infrastructure challenges determining in data center architectures.

A popular trend in data center architecture is the use of large scale, modular data centers composed of shipping containers filled with servers [7], but more radical proposals range from micro data centers placed inside condominium closets [8] to floating barges filled with servers running off of power generated from ocean currents [9]. Data center operator would like to provide service to customers with high availability where an application can continue operating despite data center failures. I propose the use of VPNs to transparently and securely connected pools of data center resources with existing customer infrastructure to both simplify resource management and high availability services.

Infrastructure of Virtual Private Cloud Services in Data Center: Virtual private networks separate the available resources within the datacenter and associate the

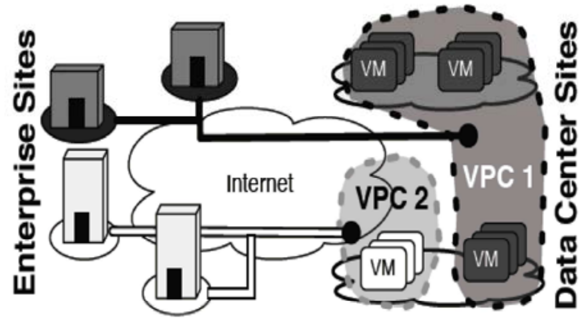

Fig. 1: Resources sharing between the data center & enterprise customer sites.

enterprise customer sites protectively. Figure 1 illustrates such a system. Resources from multiple data center sites are divided up into unique what we call Virtual Private Clouds.

Virtual Private LAN Service (VPLS) can be used to bridge the data center and enterprise customer networks into a single LAN segment. This simplifies the deployment of various applications and gives the enterprise customer in to greater control over network and server resources in the data center.

Storage area network (SAN) and LAN infrastructure are reduces management complexity. As a business adaptation high GbEthernet connectivity in the network access layer, the capacity can be effectively utilized through eliminating redundant LAN and SAN interface connections on servers. Consolidation of storage and Ethernet traffic on common high GbE connections are monitored by server [10]. The new infrastructure can scale from just a few hundred to thousands of server/storage ports, helping customers build highly scalable, high-performance, highly efficient cloud-ready (private, public, or hybrid) data center.

Integrated Virtual Server Network Infrastructure: As new applications and business models emerges in the network design that worked well. To begin with business the work may not be able to support new demands on the IT infrastructure particularly new business requirements. Networks built on fragmented and oversubscribed tree structures which have scaling and consistent performance problems. As added more devices, design, management complexity and costs increases exponentially. An integrated virtual network infrastructure can help to meet this scale of requirements, while mitigating the concerns of cost and complexity. Virtual server reduces the number of physical servers in the data center and provides greater flexibility to meet rapidly changing business needs.

However, server virtualization introduces challenges as well. Virtual machines increase the work density of traffic loads to and from individual machines (because each virtual machine has its own operating system and applications). This increases network link utilization and places additional demands on the network fabric, especially dynamic creation and migration of virtual machines [11].

Datacenter Environment: The environment could be mainly defined in the two stages such as server, network and database environments. This stage has much conservation for full control of data center environment [12].

Server Environment: The Application Server in future data center consolidation will produce two key data centers locally and out site or hot site with a subset of data and systems, as outlined in the following sections. Migration from the i-Series environment: Our core ERP systems are within the migration to the Microsoft based platform, the i-Series platform will become legacy whereas sustained for reference and legacy data storage only. In Server Virtualization the virtual environment will chain up to high support and maintain for improved availability and redundancy requirements. There will be one major upgrade stage for complete in the 2nd quarter and this brings added features and benefits for managing the environment with emphasis on data recovery. The Blade Servers will continue to be alternates of virtual servers while virtualization is not supported.

Network Environment: In WAN environment is a key of un-identity. The fundamental changes to the WAN environment will serve as a failover in the event of present circuit with Century Link has a failure. According to LAN infrastructure for the city will be extended to a private data center as the first phase of an inter-agency project which has designed to improve disaster recovery and business continuity. The second phase will involve extension to a local government agency data center placed outside the country geography. Next switching Technology there are no current major itinerary. Lastly what is required to support disaster recovery, the current infrastructure will remain the same.

Wan Interconnection Environment: In Wireless 802.11 Wireless infrastructures in data center will be evaluated in new environment and can be extended to other sites. Security and bandwidth are major hindrances to further deployment of this system. The Licensed Broadband Spectrum – 4.9 GHz is on hold pending additionally developments in acquiring. The Un-licensed Broadband Spectrum including 5.8 GHz system is still being reviewed for its usage and applications. The Wi-Fi (802.11) is no plans to introduce Wi-Fi in existing project.

Datacenter Design: The technologies also support the new trend towards consolidation and virtualization of government data centers, which is reducing the number of facilities and operating locations. However, the architects are faced lot of challenges when designing the data centers. Centralized servers and applications are keeping secure and accessible from variety of locations. Simply designing a data center to deploy more servers, storage and devices significantly increases network complexity and cost. Government agencies must change the way they view their data center network architecture in order to maximize efficiency gains.

Government Data Center Network Design Considerations: The summarization instead some of the technical considerations for designing a modern data center network that supports a consolidated with centralized server and storage infrastructure, as well as agency applications. The functional data center network design model (Figure 2) shows key design attributes. Each of these attributes is summarized in the sections.

The data center can be viewed by different groups of people interacting to create a highly available and functional end-user requirement. These typically comprise the storage, server, application and network groups. Observing all of the installed devices in the data center, one can see large racks of servers (network servers, blade servers or mainframe systems), different types of storage switches that use Fiber Channel and InfiniBand and a variety of applications (Oracle, SAP, Microsoft.etc.) that utilize these resources to deliver agency requirements.

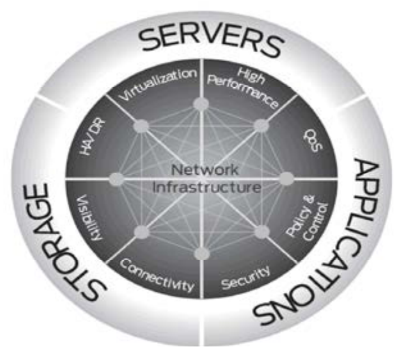

Fig. 2: Data center network functional design model

These three silos are connected through a fast, secure and reliable data center network fabric that forms the fourth silo of systems in the data center. The critical attributes for designing today's data centers availability and superior performance include:

- High Availability Disaster Recovery (HADR).
- Visibility– not only in the network traffic and security events, but also into application traffic.
- Connectivity – ubiquitous connectivity to disparate sets of resources.
- Security – data security and regulatory compliance.
- Policy and Control – centralized policy and control.
- Quality of Service (QoS).
- High Performance – applications, storage, servers and the network.

Design Principles and Plans: Key design principles are derived from operational and technical objectives. The operational objectives are fairly clear—reduce operation expenses, maintain security, adhere to green IT principles and so on. The top-level technical requirements include:

- Leverage shared infrastructures
- Employ virtualization technologies to increase utilization and efficiencies
- Ensure scalability, flexibility, security and application performance over the network

The proposed design which consolidates the data Centers and centralizes the services from Multiple Offices – This principle imposes a variety of technical requirements on the data center network. Centralizing services typically does not improve overall processing time nor data availability, but it often increases overall utilization and allows for more streamlined IT operations. Additionally, centralizing services needs to maintain the unique aspects of legacy distribution whereas different processing instances may belong to different agency entities, such as contracts management or tactical operations. Uniqueness and operational freedom must remain virtually independent.

Energy Management: The survey has identified that computer hardware alone can consume from 33% to 75% of datacenter power [12]. In addition, the power delivery, cooling (HVAC) and other infrastructure subsystems required to support the IT equipment which can constitute up to 50% of power [13]. These characteristics highlights two promising roots for improving datacenter efficiency: (1) reducing the active power signature of compute resources when executing applications, while simultaneously meeting service level agreements (SLAs); and (2) minimizing the energy costs of the HVAC systems used to cool server hardware, by balancing cooling capacity with the dynamic heat generation of workloads. Although management decisions made to maximize the efficiency of IT power usage can variance with load allocations that permit for optimal efficiency of cooling components. There has been significant work focus on power management of compute resources. Methods have been developed to utilize capabilities such as processor voltage/frequency scaling for reduced power profiles of processors and platforms [14]. Storage resources have also offer a strong opportunity to reduce power and thermal usage in enterprise systems [15].

Power Supervising and Benefits: The importance and benefits of heterogeneous compute resources in IT space have been documented from low level processor management to multi-platform management [16, 17]. The Cooling center approach attempt to leverage these contributions, while coordinating them with thermal based management decisions. In datacenter level, power consumption can be reduced by turning servers off and bringing them online based on demand [18]. Other proposed datacenter management approaches have considered temperature-aware workload placement [19]. Lastly, allows it to use for online management that leverages use of coordination, a positive control paradigm for computing systems [20], for improving efficiencies in modern datacenters.

Equipment-Dictated Power and Cooling Requirements: A modular, flexible datacenter must be able to handle equipment and racks with physically in different power and cooling configurations. Most of Sun's equipment is built with front-to-back cooling so that datacenters can be configured with alternating hot and cold aisles. Some of Sun's and other vendors' equipment uses a chimney cooling model, where cold air is brought up from the floor, through the cabinet and exhausted at the top. A datacenter's cooling strategy must be able to accommodate the different types of equipment that will be required by the business [21].

Energy Efficient Best Practices: In traditional days space was a primary consideration in data center design. More recently, the cost of energy and cooling has risen to prominence. Data center managers now must prioritize investment in efficient energy and cooling systems to lower the total cost of operating (TCO) of their facilities [22]. We propose new and adoptable best practices for energy efficiency in future data center operations: such as

- Setup the proper infrastructure and design for save the energy.
- Train the data center for cost and energy efficiency.
- Design to assess multiple factors in optimized manner.
- To optimize the maximum reliable efficiency and productivity.
- Monitor and control data center performance in real time.
- Culture depend operational excellence virtual data center
- Calculate cost and energy usage effectiveness (EUE).
- Maintain the cooling systems; use temperature control and airflow distribution.
- Eliminate the mixing of hot and cold air.
- Use effective air-side or water-side economizers.
- Secure from the natural disaster and maintain the 24x7x365 power management.
- Share and learn from industry partners.
- Make the inter collaboration of resource sharing based energy efficiencies

Data Center Deployment: Next, data center operators must deal with the deployment and planning problems related to estimating a data center's capacity and initial provisioning for new applications. This may require models of an application's resource requirements [23, 24] and an understanding of how they are impacted by different hardware configurations. As data centers attempt to improve resource utilization through server consolidation, it also becomes necessary for data center operators to understand how the placement of applications impacts performance and resource consumption [25].

Physical Deployment Environments: Physical deployment environments vary depending on the kind of application being deployed, the user base of the application, scalability, performance requirements, organizational policies and other factors. A number of infrastructure patterns with similar characteristics can be identified for specific kinds of applications, particularly Internet-based solutions [26]. A virtualization-based technology involves all segments of the data center. Virtualization is the process of abstracting physical resources and presenting them in a logical form so that physical resources can be used in ways that are more efficient and manageable. This article survey the current limitations in the data center with a specific focus on SANs and discusses the virtualization trends in each of the three major segments of the data center - servers, SAN fabric and storage [4].

Limitations with Today's Server Deployments: In most datacenters, a new server is needed for each new application that is deployed. While multi-tasking operating systems are capable of running multiple applications at the same time, limiting applications to one per server provides the advantage of fine tuning the operating system parameters for the specific application as well as preventing scenarios and single application crash shuts down all other applications on the same machine. The problem with the one-server-per-application approach is that servers are often underutilized. According to Microsoft and Sun Microsystems estimates, server utilization rates can range anywhere from 5 % to 15%, which indicates that many computing cycles are left idle. If a company could efficiently consolidate five such servers each running at 10 or lower utilization onto a single server, then the company would only need to spend 1/5th (or 20 percent) of the overall cost. In addition, the resulting 50 percent server utilization after consolidation still leaves ample room for demand spurts.

Server Virtualization: One of the approaches that enable consolidation of servers is Server Virtualization technology. Server Virtualization allows the creation of multiple virtual machines on the same physical machine. Each virtual machine created then has its own set of virtual hardware (e.g., memory, CPU etc.) and is able to host an operating system of its own. This approach facilitates the use of multiple operating system instances and also makes it possible to run different operating systems such as Solaris, Linux and Windows - all on the same physical server. By utilizing this isolated virtual server environment with its own operating system, multiple applications can now be run on a single physical server, thus improving the overall utilization of the server drastically.

In addition, applications and their virtual servers can be dynamically moved from one physical server to another, based on load conditions - thus improving resource allocation and uptime of servers within the data center. While server virtualization addresses the CPU usage, the multiple virtual servers on the physical server are limited to a single view of the SAN regardless of their individual requirements. This limits storage management, QOS and resource accounting capabilities and gives to data protection concerns in the SAN. Vendors such as Emulex are solving this problem by offering Virtual HBA technology that enables each virtual server to have an independent look of the SAN by using virtualized HBAs and thus enabling customers to fully realize the benefits of server virtualization.

Limitations with Today's Storage Fabrics: In the data center parlance, the Fibre Channel storage area network (SAN) is referred to as a fabric. Organizations have arrange base of multiple SAN islands or fabrics. These isolated fabrics are either a result of mergers and acquisitions or a result of internal SAN partitioning policies that were driven by application, technology migration or departmental isolation requirements. While a SAN by itself increases data availability, resource utilization and manageability, the existence of multiple SAN islands within an organization tends to decrease the overall efficiency of managing the SAN fabric. Separate fabrics mean more hardware, more ports, more devices to manage and typically underutilized hardware. The capacity to virtualize the various fabric resources such as the switches, ports and fabric services now enables organizations to gain the flexibility of assigning resources while still maintaining the required fabric isolation.

Limitations with Today's Storage Deployments: In order to maximize the utilization of various disk drives within an array, vendors have implemented storage virtualization at the array level. The individual storage arrays are efficient at managing the storage within the array, capacity utilization falls significantly when considering all the arrays in the entire fabric. A single array that is maximizing utilization and could force the need for a new storage array even though there are other arrays with very low capacity utilization. If a company could consolidate their various storage arrays into a larger pool of shared storage, it would increase storage utilization and thus avoid the costs associated with purchasing additional arrays and related ports.

Network-Based Storage Virtualization: Network-based storage virtualization enables pooling of physical storage from heterogeneous storage arrays into what appears to be a single storage device that is managed from a central console. It also permits storage resources to be altered and updated on the fly without disrupting application performance, thus drastically reducing downtime. Other key advantages of network-based storage virtualization include:

- Easier migration of data between physical storage resources - enabling tiered storage architectures for information lifecycle management
- Support for intelligent applications that solve requirements such as data availability, disaster recovery and point-in-time images for testing and data mining

These intelligent platforms that have taken shape in the form of a switch or a dedicated appliance provide the necessary hardware acceleration for storage virtualization functions and significantly improve the scalability and performance of network-based virtualization solutions. Vendors such as Emulex are addressing the needs of these intelligent platforms by providing intelligent storage processors that are built into various platforms including storage appliances, intelligent switches or front-end arrays. Further, the inherent performance and scalability benefits associated with intelligent storage processors typically tend to eliminate the debate around which architecture to deploy.

Improved Data Center Efficiency through End-to-End Virtualization: The technical developments in server virtualization, fabric virtualization and storage virtualization are set to rush in a new paradigm of enterprise datacenters. The capability to create virtual servers, virtual fabrics and virtual storage on demand, the time taken for provisioning the infrastructure required for data center applications can now be significantly reduced. More importantly, this will enable IT administrators to create an end-to-end resource profile for new applications (CPU, memory, fabric bandwidth, array quality, backup frequency, etc.) and let automated processes instantiate, maintain and reallocate the needed resources even as physical resources are dynamically added, removed or reassigned. In addition to dramatically increasing flexibility in managing the data center resources, end-to-end data center virtualization will also improve the operational efficiency of the data center by lowering the cost of operations through consolidation of resources and by offering increased levels of service assurance enabled by support for high availability and automation.

In all, an end-to-end virtualized data center will provide a utility-class infrastructure for offering a new range of services. Combined with computer and storage grids that enable better deployment at the physical level, data center virtualization will provide a service-oriented architecture.

Data Center Management: Data center provides best-practice guidelines and architectures on all aspects of end-to-end service management to ensure that IT processes are closely aligned with business processes and that IT delivers the correct and appropriate business solutions. Information technology management libraries' (ITML) has two major components are Service Delivery and Service Support. Service Delivery and Service Support cover more of the day-to-day operational processes of IT management [27]. Some of the most common ITML components are:

- Configuration Management.
- Release Management.
- Change Management.
- Incident Management.
- Problem Management.
- Availability Management.

Datacenter Security Services: Network intrusion detection systems deployed at several points: within a single network topology, together with host-based intrusion detection systems and firewalls, can provide a solid, multi-pronged defense against both outside, Internet-based attacks and internal threats, including network misconfiguration, misuse, or negligent practices. The Cisco Intrusion Detection System (IDS) product line provides flexible solutions for data center security. Security is also an major issue, especially for shared data center environments, leading to much research on isolating services and building trusted platforms [28, 29]. Certification, directory, network, encryption and other security components are the challenges to user so they want 100% secure network transaction. When industry struggles with developing, the technology to provide these protective components, technical manager must deal with a daily basis to reduce the network's imminent risk. A complete network security solution includes authentication, authorization, data privacy and perimeter security. Perimeter security is traditionally provided by a firewall, which inspects packets and sessions to determine if they should be transmitted or dropped. In effect, firewalls have become a single point of network access where traffic can be analyzed and controlled according to parameters such as application, address and user for both incoming traffic from remote users and outgoing traffic to the Internet [30].

In general, firewalls are intended to protect resources from several kinds of attacks such as passive eavesdropping/packet sniffing, IP address spoofing, port scans, denial-of-service (DoS) attacks such as synchronize acknowledge attack (SYN flooding), packet injection and application-layer attack.

Data Center Availability/scalability: Availability is usually ensured by the overall network design and implemention in several ways. First, networks are designed with minimum occurrence of service problems and also timely to recover from problems (such as backup recovery policies). Second, the high availability must be considered at each layer of the Open Systems interconnection (OSI) reference model, with redundancy and failover provisions made at the physical, data link (for example, Ethernet), network (for example, IP) and application layers. The most effective solutions are those with constant engineering deliberation tightly integrated whole the data center, rather than those approached with a series of point-solution products and techniques.

Scalability: Scalability must be present in every data center. Server load balancing is the norm and techniques; such as reverse proxy caching are often used to offload servers. Server load balancing, like other aspects of data center design, must also be content-aware, preferably using delayed binding, full URL and cookie inspection and sticky (server) connections as part of the logic of choosing a server for each user request. Specific content may be in high demand and considered hot, Because this content may not be known in advance (such as breaking news stories), advanced data centers should also have the ability to identify and replicate hot content to overflow or backup servers or cache to ensure ability, support that increase demand without compromising performance. Both the complications and available solutions increase greatly if the data center is part of a geographically dispersed set. The ability to provide content at multiple sites allows the server data to be closer to the requesting client, thus providing faster response but also higher availability (due to partial or complete redundancy of the data). But multi data center design must then include considerations of management of the data content, distribution of updates, synchronization of different sources, proper routing of requests, handling of downed servers, additional security and so on.

CONCLUSION

The data centre future will be a simplified, commoditized, virtualized set of computational, storage and network resources that is enabled by ubiquitous, interoperable standards. Such standardization of components and interfaces will enable the transformation, the essential act of a craft industry is manufacturing hand-crafted infrastructure for one application into standard industry for assembling machine-generated infrastructure configurations of many applications. The leap into large scale server virtualization, supported by a homogeneous high speed ethernet infrastructure that unifies the IP and storage area network (SAN), enables a truly converged data center that answers to the challenges of consolidation, increased capacity and lower total cost of ownership.

The Data Center in a Box, cave or container idea has been motivated and argued to be aligned with security and best practices such as defense in depth. It is anticipated that the black box approach can lower the cost and make transparent complexity of deploying secure infrastructure, environment, making it more accessible to the security. The concept of implementing hybrid deployments intelligently combine virtual and physical capabilities in modular design so as to provide better performance, reliability/scalability and physical separation is emphasized. Once testing is completed then deployment might be executed with well established best management practices because a converged data center network will be multi-vendor products working with different stages of the end-to-end network.

REFERENCE

1. Katz, Randy. IEEE spectrum: Tech titans building boom. http://www.spectrum.ieee.org/green-tech/buildings/tech-titans-building-boom.
2. Gmach, D., Y. Chen, C. Hyser, Z. Wang, C. Bash, C Hoover and S. Singhal, 2010. Integrated Management of Application Performance, Power and Cooling in Data Centers, 2010 Network Operations and Management Symposium (NOMS 2010), Osaka, Japan.
3. Zhang, Qi, Cherkasova, Ludmila, Mathews, Guy Greene Wayne and Smirni, 2007. Evgenia. R-capriccio: a capacity planning and anomaly detection tool for enterprise services with live workloads. In Middleware '07: Proceedings of the ACM/IFIP/USENIX 2007 International Conference on Middleware (New York,NY, USA,), Springer-Verlag New York, Inc., pp: 244-265.
4. Bakthavathsalam Ranga, 2007. Next Generation Data Centers - Evolving from Physical to Virtual Infrastructures, Senior Product Marketing Manager, 05.jan. Emulex Corp. www.emulex.com.
5. Garfinkel, Tal, Ben Pfaff, Jim Chow, Mendel Rosenblum and Boneh Dan, 2003. Terra: a virtual machine-based platform for trusted computing. In Proceedings of the nineteenth ACM symposium on Operating systems principles (Bolton Landing, NY, USA,), ACM, pp: 193-206.
6. Gupta, D., L.Cherkasova, R. Gardner and A. Vahdat, 2006. Enforcing performance isolation across virtual machines in xen. In Proceedings of the ACM/IFIP/USENIX 7th International Middleware Conference (Middleware' 2006), Melbourne, Australia.
7. Katz, Randy. IEEE spectrum: Tech titans building boom,http://www.spectrum.ieee.org/green-tech/buildings/tech-titans-building-boom.
8. Church, K., J. Hamilton and A. Greenberg, 2008. On delivering embarassingly distributed cloud services. Hotnets VII ().
9. Clidaras, Jimmy, Stiver, David and Hamburgen, William. Water-Based data center (patent application 20080209234).
10. Dr.-Ing. and Sven Graupner, 2010. Architecture for the Automated Management of Data Center IT Infrastructure, 7.
11. WHITE PAPER - Migrating to the QFabric Data Center Architecture, 2011, Juniper Networks, Inc. www.juniper.net.
12. Greenberg, S., E.Mills, B. Tschudi, P. Rumsey and B. Myatt, 2006. Best practices for data centers: Results from benchmarking 22 data centers, In Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings.
13. Report to congress on server and data center energy efficiency, 2007. Technical report, U.S. Environmental Protection Agency, Energy Star Program.
14. Isci, C., G. Contreras and M. Martonosi, 2006. Live, runtime phase monitoring and prediction on real systems with application to dynamic power management, In Proceedings of the 39th International Symposium on Microarchitecture (MICRO-39).
15. Zhu, Q., Z. Chen, L. Tan, Y. Zhou, K. Keeton and J. Wilkes, 2005. Hibernator: Helping disk arrays sleep through the winter, In Proceedings of the 20th ACM Symposium on Operating Systems Principles (SOSP).

16. Kumar, R., D. Tullsen, P. Ranganathan, N. Jouppi and K. Farkas, 2004. Singleisa heterogeneous multi-core architectures for multithreaded workload performance, In Proceedings of the International Symposium on Computer Architecture (ISCA).
17. Nathuji, R., C. Isci and E. Gorbatov, 2007. Exploiting platform heterogeneity for power efficient data centers. In Proceedings of the IEEE International Conference on Autonomic Computing (ICAC).
18. Chase, J., D. Anderson, P. Thakar, A. Vahdat and R. Doyle, 2001. Managing energy and server resources in hosting centers. In Proceedings of the 18th Symposium on Operating Systems Principles (SOSP).
19. Moore, J., J. Chase, P. Ranganathan and R. Sharma, 2005. Making scheduling cool: Temperature-aware workload placement in data centers, In Proceedings of the USENIX Annual Technical Conference.
20. Raghavendra, R., P. Ranganathan, V. Talwar, Z. Wang and X. Zhu, 2008. No power struggles: A unified multi-level power management architecture for the data center. In Proceedings of the International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS).
21. Nelson Dean, Michael Ryan, Serena DeVito, Ramesh KV,Petr Vlasaty, Brett Rucker and Brian Day, Energ Y Efficient Datacenters the Role of Modularity in Datacenter Design, Part No 820-4688-12.
22. http://www.electronics-cooling.com/articles/2007/feb/a3/
23. Stewart, C. and K. Shen, 2005. Performance modeling and system management for multi-component online services, In Symposium on Networked Systems Design and Implementation.
24. Zhang, Q., L. Cherkasova and E.A. Smirni, 2007. Regression-Based analytic model for dynamic resource provisioning of Multi-Tier applications, In Proc. ICAC.
25. Urgaonkar, B., A. Rosenberg and P. Shenoy, 2007. Application placement on a cluster of servers, In Internernational Journal of Foundations of Computer Science, 18: 1023-1041.
26. http://www.microsoft.com/resources/documentation/msa/idc/all/solution/en-us/default.mspx.
27. Whitepaper, A Blueprint for Better Management from the Desktop to the Data Center.
28. Garfinkel, Tal, Pfaff, Ben, Chow, Jim, Rosenblum, Mendel and Boneh, 2003. Dan. Terra: a virtual machine-based platform for trusted computing. In Proceedings of the nineteenth ACM symposium on Operating systems principles (Bolton Landing, NY, USA,) ACM, pp: 193-206.
29. Gupta, D., L. Cherkasova, R. Gardner and A. Vahdat, 2006. Enforcing performance isolation across virtual machines in xen. In Proceedings of the ACM/IFIP/USENIX 7th International Middleware Conference (Middleware' 2006), Melbourne, Australia.
30. Datacenters, Best Practices for Security and Performance, www.cisco.com/go/offices.