Monitoring in a Virtualized Environment

Midhun Chandran, Jayant Walvekar

Persistent Systems Limited
Virtualization Competency
402 E, Bhageerath, Senapati Bapat Road
Pune-411016, Maharashtra, India
{midhun_chandran, jayant_walvekar}@persistent.co.in

***Abstract*—Monitoring solutions for virtualized infrastructure (VI) should evolve to collect, analyze and provide configuration recommendation based on a broader range of operational metrics. A virtualized infrastructure is a complex interaction of hardware (servers, network and storage), hosting variety of multi-tier application with specific service level requirements and governed by their security and compliance policies. Most existing solutions of today monitor and analyze only a subset of these interactions. The analysis and recommendation obtained tend to optimize only particular aspects of the infrastructure and can potentially introduce violations for the others. A virtualized infrastructure is dynamic in nature, providing immense opportunities to automate configuration changes to virtual machines, networks and storage. It delivers the capability to administer the whole of infrastructure as a large resource pool shared by multiple workloads. Monitoring solutions that look at only few aspects end up forcing administrators to create silos within the infrastructure that are specially designed to ensure that business service requirements are met for the specific applications running there. A monitoring solution that can collect and analyze multiple aspects for assisting in decision making and process automation can deliver greater efficiency to the virtualized infrastructure.**

In this paper we argue the importance of having a monitoring solution that provides a holistic view of the virtualized infrastructure. We discuss the need for solutions to be capable of monitoring and analyzing a broader set of metrics such as health of infrastructure components; performance of operating environment such as hypervisors, operating systems and application running on them; capacity utilization indicators for server, networks and storages; information available with configuration and change management database containing policies including security and compliance policies. We also take a look at what these broader set of metrics are and who would be interested in them. The paper further proposes a monitoring framework for collecting and analyzing the above mentioned aspects of a virtual infrastructure to develop a more complete solution.

Keywords—Virtualization, Monitoring, Analytics, Performance, Capacity, System Health, Security, Compliance

1. INTRODUCTION

Monitoring is essential to ensuring the availability, security and usability of IT infrastructure. There are always challenges in keeping pace with new innovations in infrastructure technologies. Virtualization as a technology and its rapid adoption is such an example which has thrown open new problems and challenges in monitoring. IT departments in organizations are in need of new tools to monitor the additional complexity introduced by the technology.

Earlier it was the application, operating systems, and physical infrastructure, including storage and networking that were the primary objects of monitoring. A virtualized infrastructure introduces a new layer that needs to be monitored – the hypervisor. With virtualization, the operating systems are working with virtual resources made available to them by the hypervisor. While the hypervisor adds a new dimension to monitoring, the greater challenges lie elsewhere.

The power of virtualization lies in its dynamic nature. It has the potential to combine the individual compute capacity available on physical nodes into one large compute unit which can then be partitioned for use by all virtualized workloads - the promise of reduced costs and increased efficiencies that virtualization represents is a direct result of these dynamics. And, it is precisely this dynamic nature that makes monitoring a real challenge in a virtualized infrastructure.

In this paper we discuss these challenges and the need for two major paradigm shifts in the way monitoring solutions have been designed a. Work with the global view of the infrastructure b. Consolidate and analyze a broader set of metrics. Use of traditional monitoring tools can make management of a virtual infrastructure difficult for IT teams. The lack of visibility across various infrastructure components makes problem identification and resolution a tedious and lengthy process. Monitoring tools specifically designed for virtual infrastructure have started to emerge and even the existing ones have begun to adapt to the needs of virtualization [1]. We look at some of the gaps that exist today and propose solutions that we believe will address them.

2. VI MONITORING REQUIREMENTS

Most often than not the different solutions that build up an IT infrastructure comes from a variety of software vendors. The hardware consisting of the servers, networking infrastructure and the storage infrastructure come from different vendors. Each of these infrastructure components comes with their own management and monitoring tools. The virtualization layer comes with its own management software. It is likely that an enterprise's IT infrastructure may even consists of heterogeneous hypervisors each with independent management and monitoring software. Then there are tools that

DOI: 10.5176_2010-2283_1.1.26

specialize in specific capabilities like capacity management. There could be multiple such tools in the same environment each independently analyzing capacity for servers, networks and storages. IT enterprises also use a variety of tools for security monitoring on machines and network. Considering that the monitoring scenario is already complicated, it will help to take a fresh look at the capabilities that different human actors in an enterprise expect from a monitoring solution.

- **System Administrators**: There are at least three specialized areas in a virtualized IT infrastructure typically administered by different individuals or groups: virtual infrastructure (VI) administrators, storage administrators and network administrators. Virtualization brings interesting new interactions between these three domains [2]. For instance, two VMs on different servers can conflict with each other because their virtual disks are on the same storage mount. In another instance network latencies may increase significantly if two VMs having heavy network traffic between them are placed on separate servers. In a virtual environment it is often difficult to isolate a problem condition if all three dimensions are not evaluated in a combined manner. Administrators desire to see a comprehensive reporting of the performance and health of these systems in an integrated manner. Since their roles require them to ensure the infrastructure delivers on their Service Level Agreements (SLAs), they expect the monitoring system to be capable of detecting and alerting them to any hardware or software failures, performance degradation or capacity imbalance or overload. In addition, the expectation from the tools is to suggest placement of VMs on servers and storages that are compliant with the network policies.
- **Application Owners**: Many application owners are skeptical installing and configuring their application in a virtualized environment. First, they are often unsure how their application will behave in a virtualized environment. They need a monitoring tool to be able to study the behavior of the application on a physical server and be able to predict its behavior in a virtualized infrastructure [3,4]; many virtual infrastructure capacity planning tools are trying to do it although in a limited way. Application owners have their fears compounded by the fact that virtual machines hosting various application tiers or even different applications now share the same physical infrastructure and therefore interact and potentially conflict in ways that are difficult to anticipate beforehand. The application owners need to assure themselves that their end users are getting the required performance from the application to be productive, and that they are able to deliver on response time and availability SLAs [5,6,7]. As a result, there is a greater need for application availability and response monitoring to be in place to detect any unavailability or service deterioration issues as early as possible. Additionally, the monitoring tools need to ascertain if the observed problems are an application issue or an infrastructure issue and help in assigning the responsibility to the right owner.
- **Capacity Management and Planning**: With a dynamic virtualized environment, capacity management becomes more of a day-to-day activity compared to physical infrastructure, where procurement delays can impede flexibility. A virtualized environment can be very adaptable to changes in workload, to the evolution of applications, and can be flexible in providing windows for maintenance activities [2]. All of the popular hypervisor today have the capability to migrate a virtual machine live to another physical server. Therefore, the analytics built on the monitored data should provide the capacity manager with insight into the placement of virtual machines by examining historical patterns and current requirements. Similarly, capacity planners now need to have a global picture of the entire infrastructure, rather than capacity assigned to individual applications. They need to know in advance, what will be a constraining resource in the near future by looking at the overall usage trends across the infrastructure.
- **Facilities Management**: Facilities managers need to understand the power usage trends and predictive analysis for the future. They need recommendations about how additional efficiency and cost reduction can be extracted from the infrastructure. Additionally, they would like to see a mechanism in place to consolidate workloads to a minimum number of physical servers at off-peak hours. With virtualization the organization IT essentially manages a centralized infrastructure shared by multiple business units or departments. Therefore, a mechanism to be able to account for the infrastructure usage back to the departments should be in place to recover the current investments and prepare justification for the future.
- **IT Head**: The IT head needs high-level reporting and analysis on resource usage, license audits, capacity planning, cost of ownership, chargeback, VI performance, trends and predictive analysis, security threats, compliance, ROI, and TCOs, in addition to any perceived risks or threats. These details are required for the IT managers to plan and strategize.
- **IT Security Head**: The security head needs to see reporting from a compliance perspective and additional analysis to identify threats, security violations [8,9,10]. This need to be done with the help of security policy monitoring (configuration, patch management, VM sprawl, access control etc.), infrastructure security monitoring, and compliance monitoring.

3. METRICS TO BE MONITORED

Figure 1 summarizes the different perspectives that various entities in an organization have towards monitoring. And in a virtualized infrastructure all these perspective become more dependent on each other in the sense that changes to one should consider the impact on the others. In this section we will look into the details of what metrics are needed to satisfy the monitoring needs of each of these perspectives.

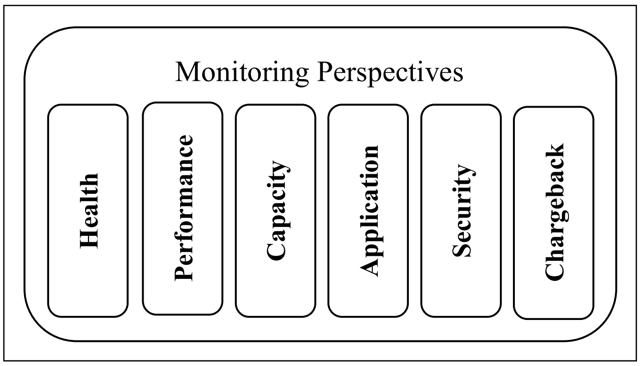

Fig. 1. IT Infrastructure Monitoring Perspectives

- **Health Monitoring** – Monitoring health/status of the complete infrastructure which requires monitoring of physical server hardware status, hypervisor status, virtual machine status, physical and virtual network switches and routers, and storage systems.
- **Performance Monitoring** – Basic performance monitoring looks at the CPU, memory, storage and network performance metrics from the VM guest OS as well as from the hypervisor. These metrics typically get monitored even in non-virtualized environments. The virtualization-specific metrics could be for specific entities that are introduced by various virtualization technologies, e.g., the cluster and datacenter concepts in VMware. The behavior of other virtualization features can also be measured as metrics such as how frequently VM migrations are occurring or when other high availability or scalability features are engaged. Then, there are specialized applications built using virtualization like desktop virtualization (VDI). Monitoring for such solution require more parameters to be collected from the virtual machine as well as the hypervisor layer, such as the service time for VM provisioned in response to an end user desktop connection request.
- **Capacity Monitoring** Today's organizations are truly dynamic and its resource utilization/requirements are continuously evolving. So, continuous planning of various resources such as servers, desktops, network, and storage is required. This requirement demands periodic audits of physical, as well as virtual resources. The capacity monitoring requires end-to-end continuous capacity monitoring of the following key metrics:
  - **Server Utilization**: Peak/Average server resource utilization – memory /CPU/resource, server bottlenecks and correlation with workload metrics like number of users/VMs
  - **Memory Usage**: Memory utilization on each server, capacity bottlenecks and relationship with number of users/VMs
  - **Network Usage**: Peak/Average network utilization, capacity/bandwidth bottlenecks and relationship with number of users/VMs
  - **Storage Utilization**: Overall storage capacity metrics, VM/Virtual disk utilization, I/O performance metrics, snapshot monitoring and correlation with number of users/VMs
- **Security and Compliance Monitoring** Virtualization introduces a new set of security risks due to VM sprawl, and offline VMs. The hypervisor is itself a new threat target [11]. Live migration of VMs can potentially conflict with the way access control is managed and policies are applied. IT security and compliance monitoring becomes critical for securing the virtualized environment. Security and compliance monitoring requires end-to-end VI activity monitoring for:
  - **VM Sprawl**: Metrics to monitor the VM activities as they get cloned or copied or migrated within or across network, or even to a different storage location.
  - **Configuration metrics**: Virtual server configuration monitoring to ensure that they are compliant with standards and hardening guidelines, VM configuration monitoring for software licensing policy enforcement. VI Events which help enforce/detect violations of IT policy. This includes individual security, organization security policy monitoring.
  - **Access Control**: Access control monitoring and reports for role-based access control enforcement.
  - **Compliance monitoring**: Metrics to validate/audit IT setups and processes for standards and regulations such as HIPAA, SOX, GLBA.
- **Monitoring For Billing and Chargeback** – In a virtualized environment, where the infrastructure gets centralized, it is important to measure resource usage by different business units, groups, and users. This information can be used to distribute/amortize, and in some cases, recover the cost correctly across the organization through a proper chargeback mechanism. The chargeback could be based on dynamic parameters such as resource usage and/or fixed parameters. To compute the correct chargeback information in a dynamic virtualized environment, it is important to monitor virtual, as well as physical resource usage and allocations, and be able to normalize the same across the infrastructure. Chargeback monitoring requires end-to-end VI activity monitoring for:
  - **Standard metrics** – All chargeable resource metrics like CPU usage, memory usage, storage usage, network usage metric
  - **Key VI events**: VI Events for virtual resource life cycle events like start date and end date of VM creation and allocation
  - **Configuration monitoring**: VM configuration in terms of assigned resources and reservations and also applications installed to account for software licensing costs.
  - **VM usage metrics**: VM uptime, number of VMs can vary depending on how the charging model is employed in the organization
- **Application Monitoring** – The need for application monitoring is very important in a virtualized

environment. Particularly because the application may have problems even if the VM or the physical server on which it is running looks perfectly normal. Application monitoring is required to monitor the basic health of application servers with the help of application specific response time and throughput metrics. The analytics on this data should be able to correlate the application-observed metrics to all layers of the infrastructure to be able to perform a root-cause analysis in the event of something going wrong. Application performance monitoring using the capture of network traffic is an interesting development in this area.

4. MONITORING FRAMEWORK

In the previous sections we discussed why there is a need in a virtualized infrastructure for monitoring solutions to look at broader aspects. In this section we introduce a framework, components of which we have been building, which focuses on centralizing the data collection process so that all data is available from one data source for higher level analytics to operate with. While defining the architecture for the framework, the three key objectives we considered are:

- Ability to collect monitoring data from variety of sources cutting across performance, capacity, compliance and security goals
- Analytical Processing capability to co-relate data collected from these sources to deliver better results in comparison to the existing solutions
- Make the data and the analysis available for other application using APIs. This is in addition to having its own reporting and notifications capability

4.1 Architecture

Figure 2 shows the high level architecture. We have focused on keeping the architecture open such that it is easy to extend it by plugging in new modules at any level. At the bottom is the monitoring data collector layer that interfaces with various monitoring data sources. It is intended to collect data from all the available sources. The virtual infrastructure represents the hardware (server, network and storage) and the software components (hypervisors and management software). The other components that the monitoring data collector fetches data from are the configuration database and existing tools for monitoring application performance.

On top of the monitoring layer is the analytics layer. The analytics layer processes the data captured and produces a variety of analytical results. The main advantage of being able to collect data from a variety of source is to ensure that the analytical results obtained are accurate and actionable. The emphasis here is on the fact that the analysis should not come up with a placement recommendation that is not appropriate with respect to some security or network policy. Most often the analytical result has to be analyzed manually to ensure that aspects that are not included in the analytics are also validated. This can often lead to delays and errors. However, a general approach forward for monitoring solutions should be for the analytics to automatically feed in the recommendation to the VI infrastructure and automate problem resolution.

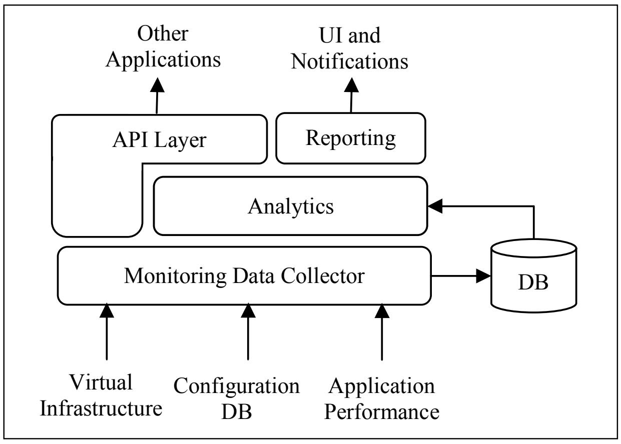

Fig. 2. Monitoring Framework Architecture

The framework provides two external interfaces. One is the user interface that displays reports based on the analytical modules integrated with the framework. Other capabilities of the user interface let administrators to tune the parameters that govern the analytical processing. For instance administrator can set thresholds to be used in computations. Notification and alerts may be displayed on the user interface or be delivered to configured end points.

The other interface is the API interface that will let other third party solutions to build upon the capability of this framework. The API interface will provide access to raw data as well as results of analytic computations. Our focus is more on the API interface than the graphical user interface as we believe that it will provide more flexibility for others to quickly build application on top.

All layers of the framework follow a pluggable architecture with interfaces defined such that modules can be built independently for extending its capabilities, for example, an independent module can be written to add data collection capability to support a new hypervisor or bring in additional analytical capabilities.

4.2 Monitoring Data Collector

Figure 3 shows a little more details about the monitoring data collector layer. It is designed to be a pluggable framework. Interfaces get defined for collection from various types of data sources and their relevant metrics. Pluggable modules implementing those interfaces can then be written to collect the data from respective data source and fed into the collector database.

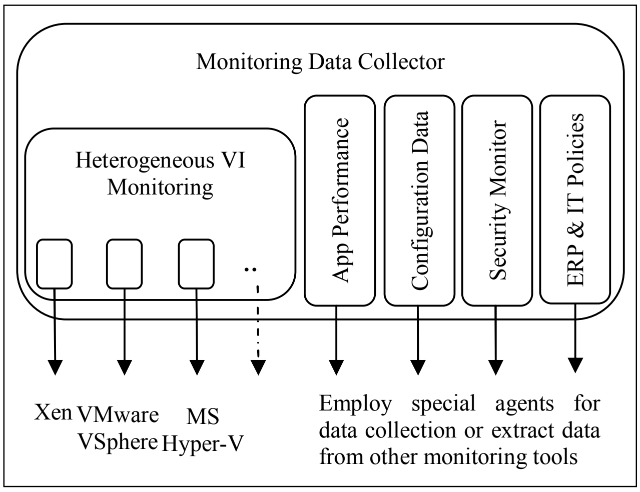
Monitoring Data Collector
Heterogeneous VI Monitoring
App Performance
Configuration Data
Security Monitor
ERP & IT Policies
Xen
VMware VSphere
MS Hyper-V
Employ special agents for data collection or extract data from other monitoring tools

Fig. 3. Monitoring Data Collector Details

Being a virtualized infrastructure the basic information collected by the data collector is from the hypervisor. The heterogeneous virtual infrastructure monitoring module defines interfaces for capturing real time data. This module itself has a pluggable architecture such that data collectors for different hypervisors can be built independently and integrated. Most hypervisors today have APIs and SDKs available that can provide the required data without having to write any specific agents. There is still some challenge to implement them because each hypervisor have their own object models and APIs and there if often some effort need to map these across hypervisors especially in case of how and what performance metrics are collected and how frequently they are sampled. There is no standardization yet on the APIs (DMTF VMAN is yet to be widely implemented) provided by hypervisors.

The information needed from the VI infrastructure is configuration information such as CPU, memory, storage and network capacity of the physical servers and the allocations to VMs. The real time information required will be the compute resource utilization metrics for the physical host and VMs i.e. CPU memory disk and network. These utilization metrics should be as absolute values and not as percentages. They should then be normalized such that a VM's utilization footprint on another server can be derived. Association of VMs with physical servers is one other useful information that is available using the APIs. VMs on a server get connected to specific virtual networks and are place on shared storage. The analysis required to optimize VM placements on servers need these additional details.

The other important piece of information that the data collection layers collects is from the configuration and change management databases. The expected information that is relevant here for analysis is the services that are associated with specific VMs in the infrastructure and a grouping of VMs based on the applications that they support. For example, consider a case of an ERP application running on a virtualized infrastructure catering to a user population size of 500 users. The ERP application would be running in production on multiple VMs with instances of web servers, middleware servers and database servers. In addition there will be development and test setups that replicate the production environment at a smaller scale. For providing useful analytics it is important to know how the VMs instances are grouped i.e. into production, development and tests at the same time the interaction between the various tiers. For instance, if it can be known beforehand that there is significant network traffic between the web server and the middleware servers then they could be collocated so that their communication latencies are lowered and do not load the physical interface. Of course, this is subject to the condition that all the other capacity requirements also allow them to be placed on the same server node. Recently there has been an emergence of tools that are capable of automatically mapping application over the infrastructure, referred to as application dependency mapping, by analyzing the network traffic.

The third major source of information that needs to be fed into is the real time application performance data. This data is required to validate that the end user performance is also meeting the service level requirements. One way to get this data is to interface with tools that are capable of measuring application performance by capturing the data flowing on the network. Port mirroring or SPAN as it otherwise called is done to route all traffic to an appliance that captures the network data. Usually these tools listen to all the network traffic going in and out of the application, even between the tiers to provide a tier-by-tier metrics for user visible performance characteristics like response time in addition to other metrics like throughput and workload characteristics. They make use of protocol analyzers to get into details about the transactions that have been performed on the application. For example, a http protocol analyzer can given specific actions executed by the users on a web based ecommerce site or a MySQL traffic give information of the queries executed against the database. They also provide additional details like response times of these transactions and the amount of data transferred while executing the transactions.

The other data sources are typical in the form a database that could either be queried directly or provide APIs for extracting the details. The modules developed will integrate with products specializing in these areas to collect the information.

4.3 Monitoring Analytics

A broader set of metrics as discussed in the previous section helps the analytics layer on top to have a more complete picture of the infrastructure and enables better decision making. Performance management analytics are the most widely required from a daily operations view. The performance management analytics are geared towards ensuring that infrastructure is well balanced on workload and the applications running on the infrastructure are responsive to its users. Current tools either look at application performance metrics to identify a fault or look at infrastructure metrics. Therefore, the root cause analysis done to identify a performance issues in the application cannot be drilled up to the infrastructure. The visibility is often limited to the application tier that is not performing well. It will be very useful to be able to accurately identify whether degradation in application performance is an application problem or an infrastructure issue. This will help assign responsibilities to the right individual and reduce turnaround time. Some of the application performance issues that are identified as infrastructure issues can then potentially be fixed automatically. The availability of storage and network configuration data along with security and compliance rules will help ensure that automating the infrastructure re-configuration does not violate these policies.

Capacity management analytics utilize real time information as well as the historical data about workloads and their footprint on the infrastructure. Using sophisticated time series analysis some level of predictive capability can be built to anticipate issues with the infrastructure over the short term as well as long term. Trend analysis based on historical data can provide information for capacity planning activity by giving warnings in advance on resources that are likely to be bottlenecks in future. Policy data can now be introduced into these analytics to build necessary buffer for fitting in the overheads created by the policy constraints.

The requirements for analytics in virtual infrastructure are really vast and have interesting and useful applications. Getting all the required data at one place for these analytics to be built is the goal of the monitoring framework.

5. CONCLUSION

We discussed the monitoring challenges in a virtualized environment and why it is essential to have a broad based monitoring solution. We introduced the open monitoring framework that is intended to overcome some of these challenges. The development of the consolidated framework for monitoring is currently a work in progress with modules being added to enhance its capabilities. We firmly believe that power of virtualization can truly be harnessed if monitoring solutions are able to bring in significant automation into identifying and fixing infrastructure issues. The challenge of today is that monitoring solutions specialize only on few monitoring aspects and hence the analysis cannot be validated for aspects that are not covered, thus requiring manual intervention for approving the configuration recommendations.

REFERENCES

[1] David Williams, Debra Curtis. Magic Quadrant for IT Event Correlation and Analysis. Gartner RAS Core Research Note, 2009.

[2] Borja Sotomayor, Rub´en S. Montero, Ignacio M. Llorente, and Ian Foster. An Open Source Solution for Virtual Infrastructure Management in Private and Hybrid Clouds. IEEE Internet computing, special issue on cloud computing, 2009

[3] B. Sotomayor, K. Keahey, and I. Foster, "Overhead matters: A model for virtual resource management," in VTDC '06: Proceedings of the 1st International Workshop on Virtualization Technology in Distributed Computing. IEEE Computer Society, 2006, p. 5.

[4] Padma Apparao, Ravi Iyer, Xiaomin Zhang, Don Newell, Tom Adelmeyer. Characterization & analysis of a server consolidation benchmark. ACM/Usenix International Conference On Virtual Execution Environments, 2008.

[5] K. Keahey, I. Foster, T. Freeman, and X. Zhang, "Virtual workspaces: Achieving quality of service and quality of life on the grid," Scientific Programming, vol. 13, no. 4, pp. 265–276, 2005.

[6] Danilo Ardagna, Raffaela Mirandola, Marco Trubian, Li Zhang. Run-time resource management in SOA virtualized environments. 1st international workshop on Quality of service-oriented software systems, 2009

[7] I. Cunha, J. Almeida, V. Almeida, and M. Santos. Self-adaptive capacity management for multi-tier virtualized environments. In Integrated Network Management, pages 129--138, 2007.

[8] Mihai Christodorescu, Reiner Sailer, Douglas Lee Schales, Daniele Sgandurra, Diego Zamboni. Cloud security is not (just) virtualization security: a short paper. ACM workshop on Cloud computing security, 2009.

[9] Bernhard Jansen, Hari-Govind V. Ramasamy, Matthias Schunter. Policy enforcement and compliance proofs for Xen virtual machines. ACM/Usenix International Conference On Virtual Execution Environments, 2008.

[10] Frank Siebenlist, Challenges and opportunities for virtualized security in the clouds. Symposium on Access Control Models and Technologies, 2009

[11] Bryan D. Payne, Martim Carbone, Wenke Lee. Secure and flexible monitoring of virtual machines. Computer Security Applications Conference, Annual, 0:385--397, 2007.

Midhun Chandran, Architect, Virtualization Practice, Persistent Systems

Midhun is an experienced software architect working with software companies building management, monitoring and automation products for virtualized environment. He has over 10 years of experience in the software industry and has a strong background in performance engineering of scalable software applications.

Midhun has a M.S is Software Systems degree from BITS, Pilani.

Jayant Walvekar, Associate Vice President, Practice Head for Virtualization, Persistent Systems

Jayant has over 17 years experience in the software industry, and is responsible for establishing virtualization practice at Persistent. As the Associate VP of Persistent he is responsible for executing projects in virtualization area, building virtualization practice, and providing consulting and engineering services to software companies (ISVs) to enhance their products to adopt virtualization platforms.

Prior to Persistent, Jayant worked for Cognos Incorporated, Canada and Infosys, Bangalore. Jayant has earned his Bachelor's in Computer Engineering from Walchand College of Engineering, Sangli.