Overcoming the Capacity Management Challenge in a Virtualized Infrastructure

The essential principles of capacity management should be applied to virtualized infrastructures. However, the level of complexity requires the judicious application of automated methods. We discuss a tractable capacity management solution to the complexities of virtualized environments.

The Virtualized Infrastructure Capacity Management Challenge

The virtualization of large portions of enterprise computing has once again raised the challenge of performing effective capacity management. In recent years the proliferation of over-provisioned systems resulted in the marginalization of the capacity management practice. Now, the cost-saving motive for virtualization necessitates effective capacity management of the entire virtual infrastructure. However, IT managers soon discover that the complexity of capacity management becomes significantly magnified in a virtualized environment for two reasons:

1. Increase of systems requiring capacity management
2. Multi-layer capacity management

Complexity of Increased Numbers

Because of the explosion in the number of servers in an enterprise environment and the substantial effort required for capacity management for each server, many infrastructure managers have elected to perform capacity management for only the most expensive and/or business critical systems. Those systems which were ignored were generally over-provisioned and their capacity needs addressed in a reactive manner. However, in a virtualized environment, the capacity requirements of these “neglected” systems add up. This was always the case, however because they now share the same set of resources as larger or more critical systems, their impact can no longer be ignored. Indeed it is here that the cost-saving promise of virtualization may be realized. All this means that capacity management must be performed for all virtualized systems. Complexity increases because of the number of systems (once physical, not virtual) which require capacity management.

Multi-layer Complexity

In the virtualized environment it is essential to provide capacity management from both the perspective of each virtual system as well as the perspective of the host systems or cluster on which the virtual machines operate. The virtualized environment introduces a new level of complexity. Within a standalone system, the operating system software, arbitrates demand for available resources. In a virtualized environment a second level of arbitration is introduced as resources are apportioned between VMs. In this second level of arbitration the host system’s hypervisor will attempt to isolate performance impact on one VM from variable resource demands of other VMs using resource allocation rules and priorities. The task of capacity management is made tractable by dividing the effort into host-wide capacity management and per-VM capacity management.

Capacity management at both host and VM levels is illustrated in Figure 1. For each VM, predictions of resource demands are made and VM provisioning adjusted as needed. This is the iterative process of capacity management. As VM consumption varies with usage and provisioning, predictions are made as to overall resource demand on each host. When necessary host resources are upgraded or VMs rebalanced. Both VM and host capacity management activities are iterative.

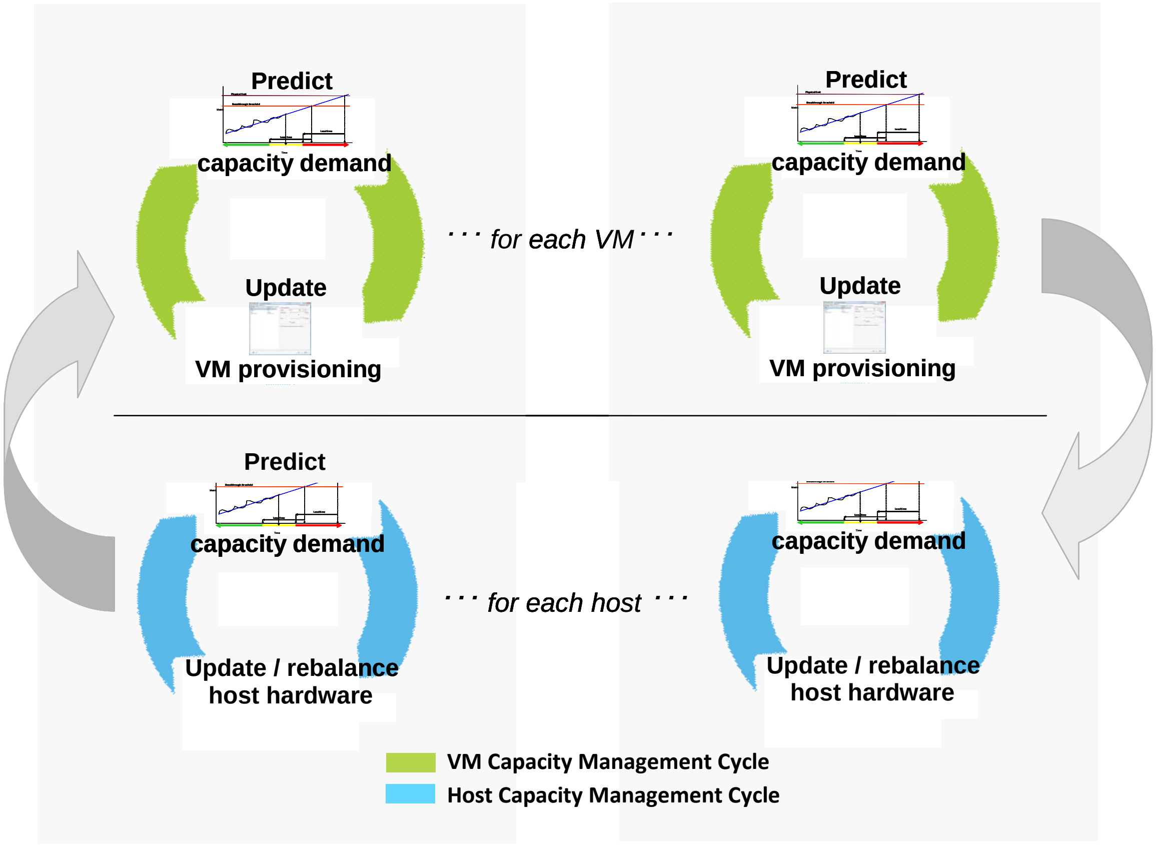

Figure 1 - Capacity Management Cycles in a Virtualized Server Environment

Although the virtualized environment presents added complexity, the essential methodology of capacity management remains the same. It is this fact which will enable provides the basis for a tractable solution to realistic capacity management in a virtualized server environment.

Capacity Management Methodology

Capacity *management* is the discipline of ensuring that the appropriate IT resources are available at the right time in the right amount. By definition, this activity is pro-active in nature. However, unexpected workloads suddenly appear; resource failures trigger workload failovers; operational errors happen. Because of unanticipated resource demands and the imperfect nature of any implementation of the capacity management discipline, capacity management will always embrace some level of reactive exercise. For this reason, we distinguish between *near-term* capacity management and *future-term* capacity management which we call *capacity planning*. Management of current resources is both reactive and proactive in nature with an emphasis largely on the former. Capacity planning is proactive in nature. If performed correctly, capacity planning largely eliminates the reactive component of capacity management. Figure 2 illustrates the PerfCap approach to both aspects of capacity management.

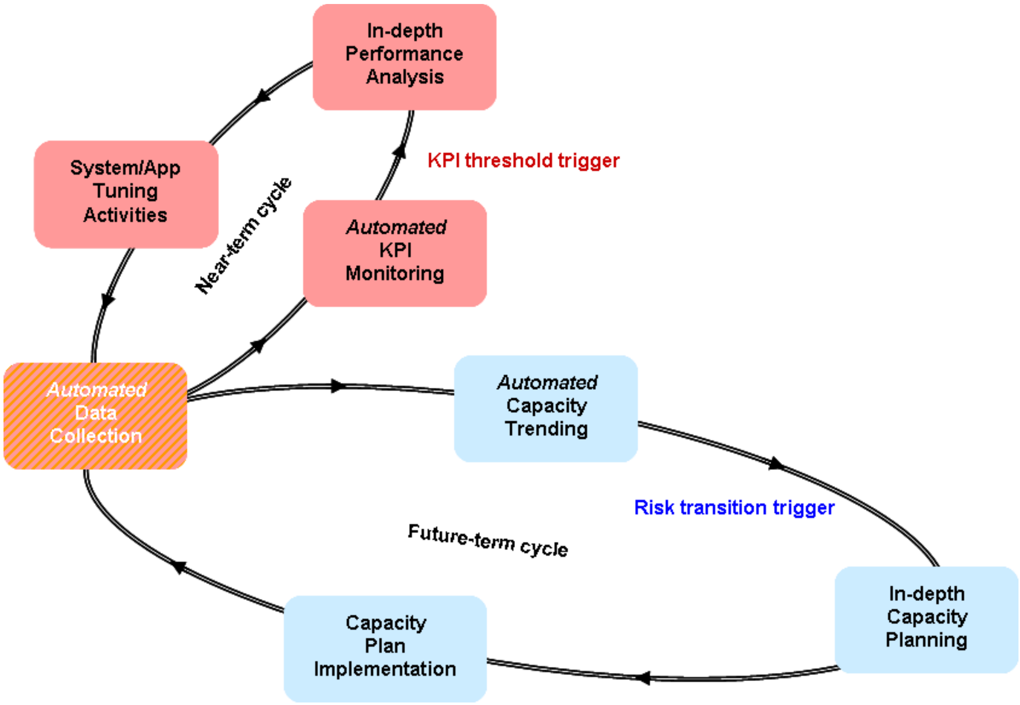
In-depth Performance Analysis
System/App Tuning Activities
KPI threshold trigger
Near-term cycle
Automated KPI Monitoring
Automated Data Collection
Automated Capacity Trending
Risk transition trigger
Future-term cycle
In-depth Capacity Planning
Capacity Plan Implementation

Figure 2 - Near-term *vs* Future-term Capacity Management

Near-term Capacity Management

Near-term capacity management is a well-established methodology. Key Performance Indicators (KPIs) such as CPU utilization, storage latencies, network utilization, and response times are continuously monitored and compared to a set of performance alert criteria. These criteria may involve specific thresholds, durations for which the threshold is breached, and times of the day or week during which the criteria apply. When the alert criteria are satisfied, a notification trigger is established. This notification initiates in-depth analysis and remediation activity.

Future-term Capacity Management – Capacity Planning

Future-term capacity management, or capacity planning, foresees capacity issues within a defined event horizon and ensures that the appropriate steps are taken to mitigate the problem.

Capacity planning for organic growth is based on past behavior of KPIs. The KPI history is used to predict future behavior as illustrated in Figure 3. The purpose of prediction is to determine if or when in the future the KPI will reach a level indicative of a performance/capacity problem. KPI risk analysis is a concise approach to summarizing the projected behavior of a KPI. When a KPI's risk status reaches a specified level, an in-depth capacity plan exercise is triggered.

KPI Risk Analysis

The key to effective future-term capacity management is KPI risk analysis. Figure 3 illustrates the essential concepts of KPI risk analysis.

Trends are based on a set period of past history. Thresholds are set based on the nature of the metric being monitored, and lead times are set based on time required to resolve capacity issues. For each monitored node, for each key performance indicator, PAWZ automatically determines the risk color status on a daily basis. If a metric is projected to reach the physical limit within the lead time, the risk status color is red. If the metric is projected to cross the breakthrough threshold within the lead time, the risk color status is amber. Otherwise the risks status color is green. When a key performance indicator transitions to amber or red, a notification is raised via a web report or an e-mail to subscribers for the monitored system. In this manner the capacity planner's attention is only raised when the automated risk analysis determines risks have risen to a defined level.

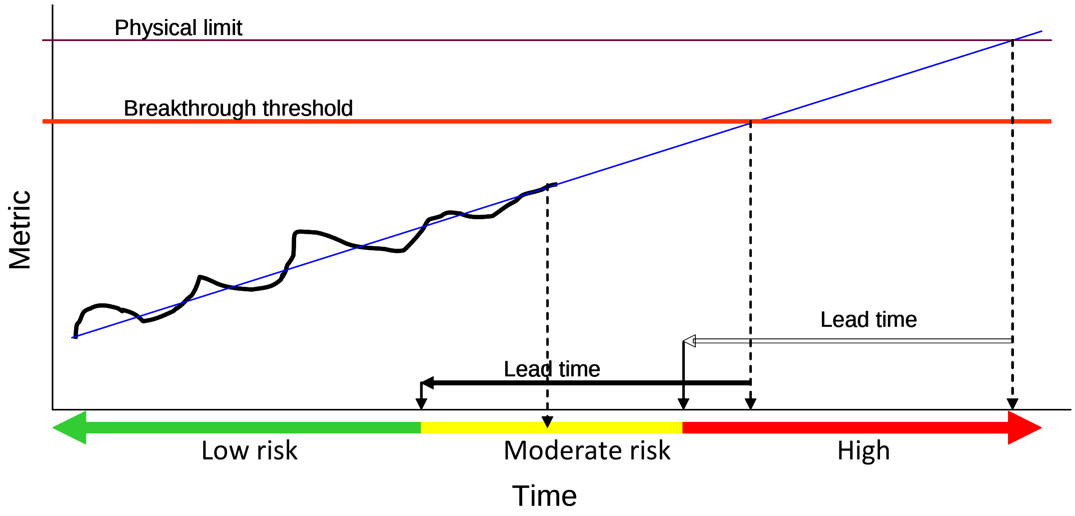

Figure 3 - Performance Risk Color Determination

As a KPI's risk color status is evaluated day to day, it is the transition to a higher level of risk color which triggers the capacity planning effort required to mitigate the projected performance problem.

Automation, the Answer to Managing Complexity

Whether one must manage an estate of 5 servers or 10,000 servers, this fundamental methodology of capacity management is the same. However, capacity management of large virtualized server estates requires extensive automation. It is impractical to require the capacity planner to review large numbers of reports. Even the review of health or risk status summaries becomes tedious.   Human skills are best activated only when analytical automation has identified the need for such skill!

In order to enable the capacity planner to manage large numbers of virtualized servers, the solution must automate:

1. performance/capacity data collection
2. data transfer
3. data archiving
4. data threshold analysis and reporting (near-term management)
5. performance projections and risk analysis (future-term management)
6. notification of emerging capacity risks (future-term management)

Additionally, these automated capacity management activities must be applied to both the virtualized systems and the physical hosts. This automation frees the capacity analyst to devote time and skill to the relatively small number of systems where capacity problems have occurred or are projected to appears within a given future time horizon.

Capacity Management Using PAWZ

PerfCap Corporation's PAWZ software provides facilities for near-time and future-term capacity management. The first four automated activities listed above are realized by PAWZ's robust data collection, transmission, and archiving, features. Automated data analysis and performance risk analysis and emerging risks notification is the key feature to capacity management of large virtualized estates.

Fundamental to these facilities are the PerfCap's low-footprint data collectors for both VMs and vSphere hosts. The metrics which are collected within each VM are extensive and beyond the scope of this article. Table 1 in the Appendix lists the vSphere host metrics which are collected using the PerfCap collector.

Data collected may be displayed as time series graphs either in real-time or historically. Additionally reports may be generated providing various rollups by day, month, or year. However, capacity management by reviewing and evaluating graphs and reports becomes extremely difficult when managing large server estates. As already stated, PAWZ automates collection, storage, and analysis of data. When performance or resource capacity issues are detected by PAWZ, notifications in reports, log files, and e-mails are generated for subscribers.

PerfCap's approach to near-time capacity management is as follows:

A. For each VM and each vSphere host,

1. Continuously monitor key resource usage metrics

2. Automatically analyze key resource metrics for extended periods of excessive utilization (e.g. CPU utilization has exceeded 95 % for at least 15 minutes) or indications of performance problems.

3. Automatically notify of such events by web-page report or by e-mail notification to subscribers. Notification may be real-time or end-of-day analysis

4. Automatically determine key resource metrics future risk status color

5. Automatically notify of risk status color changes.

B. In-depth analysis facilities to:

1. Provide web-based graphs of key performance indicators

2. Examine and compare detailed performance metrics

3. Observe VM performance metrics in real-time using the PAWZ web interface.

4. Provide graphing facilities to compare daily time-series performance data with historical values.

5. Provide aggregate graphs of vSphere cluster resource utilizations

PAWZ provides visibility of how CPU, memory, disk, and network resources are used by each hypervisor host as well as aggregate CPU, disk, and memory metrics for clusters of vSphere hosts. Threshold exceptions may be created on all appropriate performance metrics. Other metrics may be viewed using detailed graphs. The resource consumption pattern on a given day may be compared with historical resource consumption using the daily profile graph.

With the exception of VM-view CPU utilization, VM-view metrics are reliable indicators of resource consumption within the VM. In the case of CPU utilization reported within the VM, virtual timing effects render this metric unreliable. An effective replacement is CPU queue length.

PerfCap Recommendations for Monitored Metrics for vSphere VMs and Hosts

With the exception of CPU utilization, the procedure for monitoring risk color status for physical servers is the same for virtualized machines.

For vSphere hosts, PerfCap recommends that risk color status be created for the following metrics:

- Overall CPU utilization
- Peak 20-sec CPU utilization
- Per-processor CPU utilization
- Peak 20-sec per-processor CPU utilization
- Overall I/O and data rates
- Per-datastore response times
- Queue length for each LUN
- % physical memory used
- Data rates for each physical NIC

Risk color status definitions for VM should be identical to those defined for physical servers with the addition of CPU queue length.

Because of the large number of hosts and VMs to be monitored, trend graphs and color status reports become difficult to monitor in a timely manner. For this reason, PerfCap supports e-mail alerts for all color status changes be enabled. This frees the capacity analyst to devote time to analyzing and resolving real problems.

Conclusion

The complex challenges of capacity management of a large virtualized infrastructure require the use of automation. The PAWZ architecture provides automated data collection, transmission, data archiving, near-term analysis, and projection of performance risk with timely notification to the capacity planner. This approach supports PerfCap's near-term and future-term methodology for capacity management. The capacity planner's attention is only engaged when PAWZ automated processes determine that the level of risk has reached a level requiring timely investigation and intervention.