



VMware Migration and Consolidation Without the VMware Capacity Planner

Louis F. Springer

November 2007

Contents

	Contents	ii
	Revision History	iii
1	Overview and Introduction	1
2	Requirements for Empirical Testing	3
3	Approach Steps for Capacity and Suitability	3
3.1	Workload Characterization	3
3.2	Initial Infrastructure Architecture Implementation	4
3.3	Conceptual Capacity Model	4
3.4	Pilot Migration Planning	4
3.5	Pre-Pilot Load Testing	4
3.6	Pilot Migration and Elaboration	4
3.7	Staging Workload Compression	5
3.8	Pilot Workload Instrumentation Cycle	5
3.9	Rollback Capability	5
3.10	Elaboration Completion	5
3.11	Production Migration and Transition	5
4	Post Transition Facilities	6
5	Other Migration Planning and Execution Steps	6
6	Migration Plan Approaches and Requirements	7
7	Conclusion	7

Printed 12 February 08

Revision History

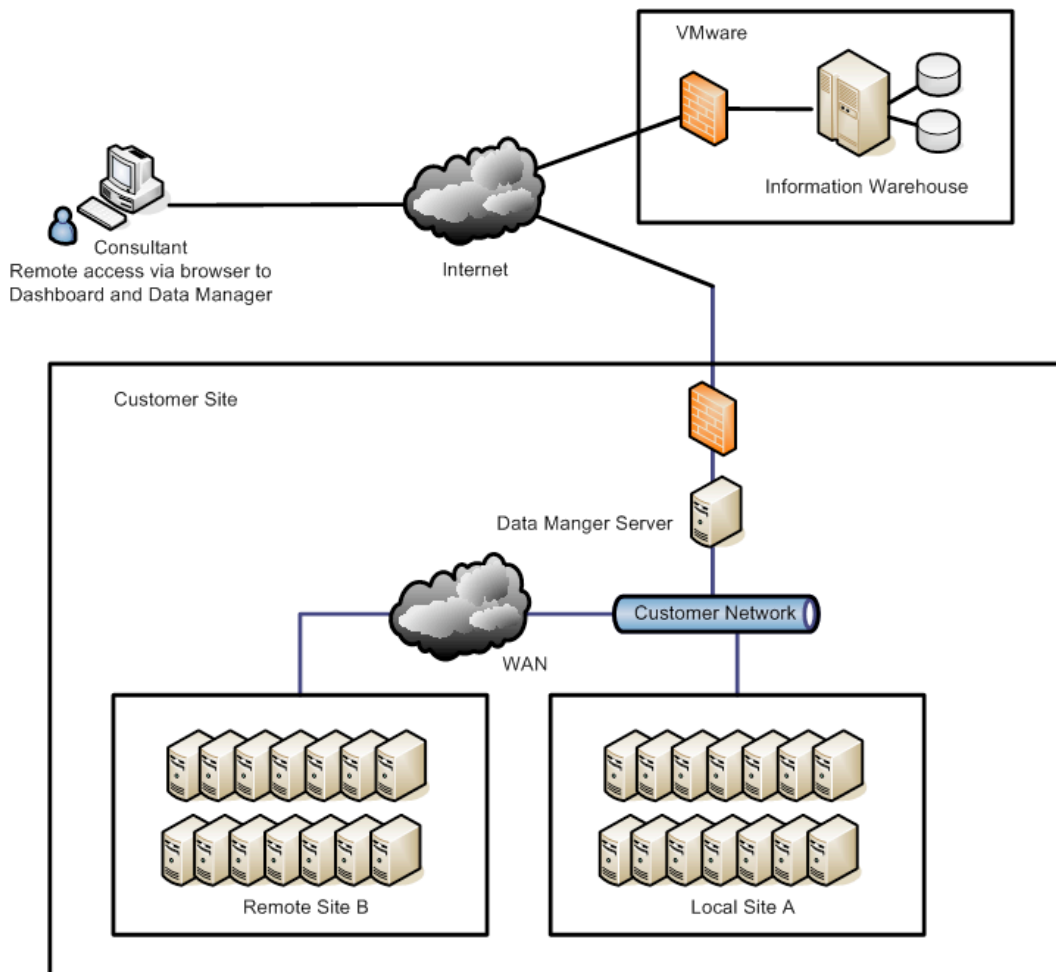
Revision	Date	Author	Description
1.0	11/9/2007	Louis F. Springer	Initial.
1.1	2/12/2008	Louis F. Springer	Minor rewording. Added conclusion.

1 Overview and Introduction

The typical method employed for migrating existing workloads to ESX from physical machines involves instrumentation of the workload in the current environment, coupled with technical and business analysis to determine:

- Workload suitability for virtualization
- ESX resource requirements to support the workload

The technical resource sizing method most often used leverages the VMware Capacity Planner. VMware capacity planner employs an offsite analysis of the workload in a VMware hosted service. This service utilizes collected VMware Capacity Planner metrics for over 20,000 machines, coupled with experience with factors affecting virtualization and ESX resource utilization, to provide a “black box” solution to consolidation target composition. The solutions provided are generally accurate enough to substantially mitigate risks associated with oversubscription of ESX resources. Typically, organizations are able to migrate with minimal additional empirical workload testing from the physical server to the ESX virtualized environment.



Organizations that cannot use the VMware Capacity Planner facilities for workload analysis, due to policy or technical constraints that disallow the use of the VMware offsite analysis facility, must employ more elaborate and extended empirical testing of workloads to ensure that virtualized workloads:

- Do not oversubscribe ESX server resources,
- Do not incur unacceptable performance degradation.

Currently, there are no adequate workload modeling facilities, outside the VMware Capacity Planner, that are able to substantially mitigate risks associated with improper sizing and qualification of workloads. There are a number of very mature modeling tools for accurately sizing consolidation solutions when migrating from physical-to-physical servers, such as TeamQuest Modeler. These modeling tools provide extremely accurate predictions for typical consolidate scenarios for several reasons:

- Many manufacturers, such as Sun Microsystems, provide these vendors with the requisite, normalized performance benchmarks to perform accurate comparisons.
- Many years of experience with modeling complex workloads have refined the processes and algorithms to a point that the predictive capability is acceptable in almost all cases.

Unfortunately, the virtualization marketplace, vendor relationships and third-party technical capability is insufficiently developed for third-party capacity planning tools to accurately model physical-to-virtual consolidation scenarios. In the future, this type of modeling capability will almost certainly become available as:

- Competitive market pressures force all virtualization vendors to supply normalized benchmarking metrics to third-party modeling vendors such as VMware and,
- Industry experience in the complexities of modeling physical to virtual consolidation improves.

Traditional approaches to modeling workloads for consolidation without benchmarking data will fail when extrapolating solutions for VMware. These methods generally revolve around estimating capacity requirements based on observed CPU utilization of the workload and the relative processor speeds in the current and target configurations.

The two most critical components for accurate models revolve around:

- Complex interactions between the hypervisor scheduler and the OS scheduler of the virtual machine, particularly for multi-threaded workloads and
- Difficulty in predicting the amount of CPU resources required for providing idealized virtual hardware, particularly as related to network and disk IO.

Because of these factors, organizations electing not to leverage the VMware Capacity Planner analysis facilities must perform more elaborate empirical testing and instrumentation of workloads to ensure migration and consolidation risks are adequately mitigated. The need for accurate, extended instrumentation is heightened for target architectures with high compression ratios and low excess capacity designs.

2 Requirements for Empirical Testing

Adequate testing ensures that the production workload behaves as predicted after migration to the ESX server. This generally can be accomplished in one of two ways:

- Instrumentation of test workloads designed to mimic production behavior,
- Instrumentation of actual production workloads in a manner that mitigates the risks associated with migration and virtualization.

The development and testing of qualified test workloads to mimic production workloads is costly, error prone and difficult. Because of this, as well as the time constraints and the typical number of migrations and workloads to be tested, the recommended approach tends toward later alternative: instrumentation of actual, virtualized production workloads in a manner that mitigates the risks associated with the initial deployment of the migrated and virtualized production workload :

- ESX server capacity and resource oversubscription,
- Unintended effects of the workload on other virtual machines and resources on the ESX server, and
- Unacceptable virtual machine performance independent of resource availability.

Another approach would involve using VMware resource containment strategies, such as configuring assumptive limits and minimums to all resources when introducing the workload, such as pinning the workload to particular CPU utilization constraints or resource limits on network and disk I/O, but this technique can tend toward implementations with underutilized excess capacity, as well as extremely fragile and complex tuning parameters. This approach also does not address the needs:

- To build an initial baseline of unconstrained workload that is necessary for the development of the effective, necessary and sufficient resource constraints that minimize the negative impacts of constraining resources while
- Ensuring the workload maintains the desired performance under the resource constraints that are applied, and
- The desire to develop reasonable shared, homogenous resource constraints based on experience with workloads of a similar type and character.

3 Approach Steps for Capacity and Suitability

In this situation, the following steps are recommend to build and maintain the desired robust virtualized computing environment offering the maximum efficiency within typical time and cost constraints.

3.1 Workload Characterization

Develop a complete workload characterization model covering the inventory to discover and document the prevalent patterns and characteristics for a so-called '80%' target solution with a combination of technical

metrics, application types, machine types, I/O types, service types and service levels. This analysis will yield patterns and opportunities that can be leveraged to minimize risks and maximize the value of the consolidation and virtualization.

Interviews and data collection will result in an inventory of servers and workloads, characterized by the risks associated with virtualization and consolidation, sufficient for reasoned prioritization of the machines for virtualization.

3.2 Initial Infrastructure Architecture Implementation

Parallel with this effort, the conceptual hardware architecture should be implemented, along with preliminary testing of the facility. These facilities must include any tools required for instrumentation of workloads for initial testing, over and above typical production instrumentation, such as TeamQuest probes and modeling tools.

3.3 Conceptual Capacity Model

Based on the analysis, develop a conceptual design of simplified target configurations, primarily focused on resource constraints, minimums and maximums, for elaboration in early migration iterations. This analysis should use the best available information, including any available, applicable benchmarks, to derive a conceptual capacity model for test, verification and refinement throughout the migration project. The analysis will also derive the requirements for post transition facilities to be included in the migration process design.

3.4 Pilot Migration Planning

Develop an initial migration or small set of migrations, specifically designed to elaborate the configuration, proving the quality and capability of the configuration and the capacity model, as well as proving any other risks or unknowns in the architecture, such as backup and recovery facilities, management and monitoring facilities and so forth.

3.5 Pre-Pilot Load Testing

Perform initial non-production load testing of the virtual machines and applications, using any reasonable cost-effective means to generate real-world loads. Based on these tests, the predictive capacity model should be updated to reflect the results.

3.6 Pilot Migration and Elaboration

After initial non-production testing, elaboration during the pilot should migrate actual production virtual machines and workloads to ESX servers with abundant excess of resources and little, if any, concurrent load.

The initial pilot migrations should result in a further refinement of the predictive model for the type of workload. These initial tests must have heightened sensitivity service levels to ensure production service levels are not compromised, These migrations must be accompanied by fully tested rollback-to-physical contingency plans that fit within migration windows and outage constraints.

3.7 Staging Workload Compression

After initial unconstrained use of virtual machines in ESX servers with substantial excess capacity, the workloads should be migrated to higher-workload density production servers. This process should be accomplished in as many stages as necessary to validate the predictive capability of the resource capacity models, and should involve reasonable permutations of workload mixes to build experience, confidence and operational capability around managing concurrent workloads with the represented characteristics.

3.8 Pilot Workload Instrumentation Cycle

The elaboration monitoring of production workloads should extend through an entire business cycle of workload variations, typically a month for most businesses.

SLA monitoring for each application workload should be established and conducted throughout all phases of the process.

3.9 Rollback Capability

Every successive, compressing migration, or virtual machine to configuration change, must reserve sufficient capacity throughout the migration period to allow for rollback to the prior state, in the event of any unacceptable service impacts, until the new configuration is deemed acceptable and the change accepted.

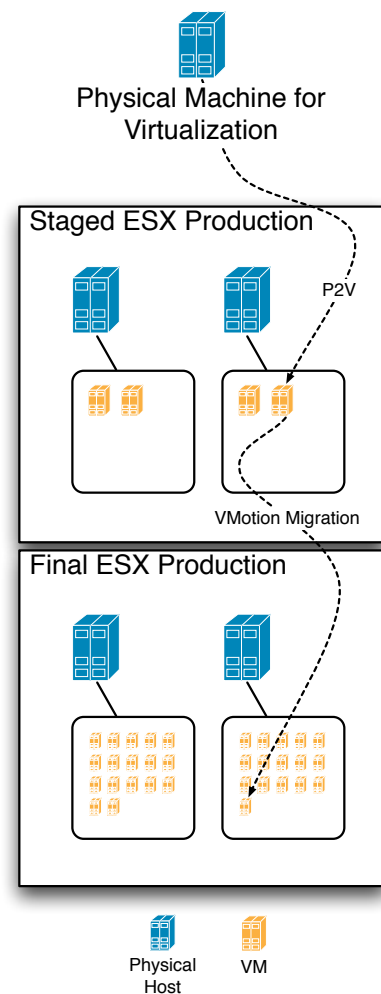
3.10 Elaboration Completion

After all elaboration migrations are complete, analysis must generate an informed and documented recommendation to proceed with transition and further migrations, or recommend further elaboration to mitigate any unacceptable risks or unknowns.

3.11 Production Migration and Transition

Once elaboration migrations are complete, transition migrations of additional migrations can begin. Succeeding migration iterations may proceed at a greater pace, or with fewer transition compressions, as the predictive model is refined and workloads behave closely enough to the model to acceptably mitigate and manage migration and consolidation risks.

Every migration must start with at least one initial migration to an ESX server with abundant excess capacity to perform initial, baseline virtualized workload measurements and perform final qualification of the workload for transition to high-density ESX servers.



4 Post Transition Facilities

Regardless of the approach to migration planning and initial architectural testing, IT operations will require ongoing management of facilities that efficiently manage the VMware service, and ensure timely resolution of incidents when they occur, as well as supporting operations requirements to ensure new workloads are suitable for the environment.

After transition and final production turnover, the operation must always maintain production ESX servers with sufficient excess capacity to isolate workloads that inexplicably operate inconsistently with the predictive capacity model for the workload. This provides a “cushion” for quickly isolating non-conforming or inconsistent workloads from other workloads. This also provides an isolated environment for rapid analysis, evaluation and incident management activities as well as for evaluating urgent configuration changes for the problem workload. These facilities must never be scavenged for functional testing or other purposes than supporting urgent production Incident Management processes.

After transition and final production turnover, the production operation must also maintain sufficient low-density facilities for qualifying pre-production and full production workloads and production workload changes for qualification and migration to high-density servers. These facilities must also be available for extended analysis activities by Problem Management processes and other processes that may require the facilities for elaborating systemic changes to the shared facilities. These facilities must not be used for functional testing of applications, training or any other purpose, and are reserved for “fit for service” analysis and testing of VMware service and configuration changes.

5 Other Migration Planning and Execution Steps

“**Approach Steps for Capacity and Suitability**” is primarily focused on aspects of design, testing and planning focused on mitigating the inability to leverage best practice tooling for this effort, the VMware Capacity Planner analysis facilities, due to the unacceptability of exporting capacity and utilization data outside the enterprise facilities to VMware.

There are other components of migration planning and execution proposed regardless of this aspect of the program.

The optimal migration approach also includes the development, refinement and execution of documented migration procedures meeting the following requirements and characteristics:

- The development of documented migration plans in concert with owning development teams, test teams and the application business owner,
- Multiple checks and verifications on each virtual machine for hardware dependencies and conflicts,
- Functional and integration tests of migrated virtual machines,
- Developing and executing Change Requests and Release Management plans for migrations, including interim steps, such as testing, as required,
- Running pre-migration backups of machines,

- Scheduling of all impacting migration activities with the affected business units,
- User acceptance testing on all migration plans,
- Monitoring of performance at all stages of migration execution for variance with accepted or expected measures,
- Using VMotion for optimal final workload placement,
- Extended workload monitoring post migration throughout the applicable business cycle for at least 15-20 days,
- Fully tested rollback contingency plans that fit within migration windows and outage constraints for each phase of the migration process.
- Decommissioning of the source server after formal acceptance of the migration and
- Updates to any operations documentation reflecting the service changes.

6 Migration Plan Approaches and Requirements

The initial stage of the ideal migration plan extends serially from cold shutdown of a physical server, though copy of data to initial deployment of the virtual server and the requisite testing as outlined. Not only must migration plan initial step fit in the available service outage window, the back out contingency plan must also fit within this window, in the event it must be exercised.

Services that do not have sufficiently long service outage windows to accommodate a serial approach will require various techniques to accommodate compressed outage windows, such as:

- Hot as opposed to cold duplication of the physical machine disks.
- Concurrent testing of a migrated instance with the production instance with subsequent final migration.
- Application level backup and restore to a point-in-time, such as though log file application to databases.
- The use of web and application tier characteristics to minimize migration outage impacts, such as though the use of load balancers to distribute all or part of a workload during transition to and from pre-migration and post-migration service instances as need.

The ideal plan will derive the smallest set of acceptable, tested migration patterns and plans for the facilities and services to be migrated. This approach will yield the optimal set of tested and validated migration strategies and plans for the source footprint, based on the technical and business characteristics for the services.

7 Conclusion

Although use of the VMware Capacity Planner yields good consolidation solutions reliably and cost-effectively, it is not the only approach. Organizations that are unable to use this facility can use other, empirical methods to prove the workability of the consolidation solution.