September, 2012

IBM IT Optimization - Tools and Technologies White Paper

R1.0

Abstract: This paper describes the advantages of methods-based tooling to support IT Optimization, and tools and technologies IBM uses for Server Consolidation, Virtualization, Operational Optimization, Cloud Enablement, and Migration.

Introduction and Overview

Migrating toward a cloud infrastructure and consolidating IT capabilities can yield significant long term benefits in cost reduction due to savings in space, power, cooling, maintenance support, and reduced investment in new capabilities. Such efforts are inhibited when extensive and detailed information about the current environment is lacking. Servers with heterogeneous configurations, distribution, ownership, and management greatly magnify the challenge and increase risk.

Any major enterprise transformation effort depends on having complete, accurate, and pertinent information available in a timely manner to ensure a successful transformation. Knowledge about the inventory, configurations, utilization, workloads, interdependencies, networks, and mission functionality of the technical components are critical to developing a repeatable and scalable approach utilizing an integrated method and tool set that ensures: 1) continuing mission operations, 2) a well justified and anticipated return on investment, 3) consistent, rapid execution of the migration, 4) a cost effective and agile target environment, and 5) efficient use of target resources by migrated capabilities. IBM's integrated approach circumvents challenges that may arise when relying on manual inventory gathering, which has proven to be insufficient in accurately prioritizing, estimating migration costs, assessing risks, or uncovering the details to successfully execute a migration initiative.

Manual inventories are typically 70-90% complete/accurate, yet the need exists to identify approximately 100% of the assets to perform a successful migration. Dependencies between applications and equipment are more complex than most IT staff realizes.

IBM has developed a methodology and a supporting suite of tools that enable moving beyond straightforward consolidation into IT Optimization – that is, maximizing the efficiency and effectiveness of a data center. This is accomplished through an in-depth analysis of the complete system stack across the complete set of IT components (such as applications, services, appliances, servers, and storage) provided by the data center. The analysis is largely conducted using these methods and assisted by automation.

Using automated data collection and an analytics engine reduces the cycle time and labor cost and fosters the development of the right design, approach, and business case to support the consolidation and migration objectives. Automation also facilitates the repeated collection of data, which allows continued refreshment of available data.

IBM's automated discovery tooling supports collecting this data on a large scale. IBM's analytics capabilities enable effective use of this configuration, utilization, and dependency data to create a detailed implementation roadmap for future consolidation and optimization that will deliver maximum return on investment and effective migration execution. Then, automated migration enables successful staged/phased transition to the target solution using the analytics based roadmap, with a back-out plan in place to mitigate risk.

Advantages of Method-based Tooling

The huge size and the pace of change required for optimizing IT environments means that it cannot be done manually.

Collecting reliable and complete data and using this data for associated analysis encompasses all aspects of the source environment to be migrated. Automation reduces the cycle time and labor cost for gathering complete and comprehensive data, even if some data can only be collected manually. This data can then be confidently used to fully develop the right design, approach, and business case to support the consolidation and migration objectives.

Automation tools can detect a wide range of IT components (servers, applications, appliances, network devices, etc.) to verify IT inventory, and identify how servers and applications are connected to other applications (affinity). They can evaluate how interdependencies impact potential optimization opportunities, and identify technical constraints that affect the order in which servers and applications are optimized. Tools can support the appropriate methods and skills to evaluate potential business, technical, security, and financial dependencies to determine and associate each with optimal target platforms (a "Fit-for-Purpose" roadmap). Ultimately, tools can expedite the process used to determine which of the applications and infrastructure (within the data centers) should be optimized (consolidated, virtualized, etc.), and identify opportunities to reduce IT workload and infrastructure.

Business cases needed to justify IT budgets can be supported with automated tooling as well. Using industry standard costs combined with the actual number of devices detected in the IT environment, we can calculate the potential high-level reduction/savings from the IT optimization effort and produce a useful Total Cost of Ownership (TCO). Hard ROIs typically include power, space, and cooling savings, and data center reduction. Soft ROIs may include increased mission agility, producing a known baseline, and help mission owners focus on mission rather than IT.

The remainder of this paper focused on tooling as used in various phases of a typical project.

Inventory Data Collection

During this phase, IBM identifies IT processing assets by understanding the current IT environment through discovery of IT assets and their relationships (supported by automated tools such as Galapagos and TADDM) and leveraging key artifacts such as the overall architecture, inventory, risk assessments, and Standard Operating Procedures (SOPs). The data collection process includes discovery of the IT environment (normally iterative), and can include identifying targeted servers, establishing access, and performing discovery.

TADDM: Tivoli Application Dependency Discovery Manager (TADDM) delivers network based, automated discovery and configuration tracking capabilities to build application maps that provide up-to-date visibility into application complexity. This includes affinities, i.e. application/workload interactions and dependencies, which are critical to discover and manage application dependencies.

TADDM enables visualization of several key types of server affinities and application dependency modeling:

- Server Affinity refers to one or more devices communicating with one another. (e.g. Server 1 makes a connection to itself, Server 2, or unknown device).
- Application (Process) Affinity refers to one or more process communicating with one another

(e.g. 'iis.exe makes a connection to db2sys.exe).

- Application Dependency refers to one or more process instances communicating with one another (e.g. site "www.ibm.com/team" makes a call to the "Busapp1a" cell that then calls the "CONTACTS" Database).
- Relationships between business applications, identified manually, are also considered.

SCMD: When we cannot reach a server over the existing network, we use Script Collected Manual Data (SCMD) where scripts are executed at the individual server and the output files are returned for processing and analysis.

Galapagos: Galapagos is an automated model-driven tool used to discover end-to-end applications and data relationships in multi-tiered distributed systems. In a planning context, or on an ongoing proactive use basis, it can be used to discover which applications are impacted by other applications. For planning application deployment transformation, it is useful for partitioning middleware to avoid overloading resources and for consolidating middleware for infrastructure transformation. Galapagos has been employed to aid business objective focused IT optimization projects for customers such as DoD and DARPA.

The data collection tools have the following features and benefits:

- Comprehensive: self contained appliance that scans the environment for comprehensive server, application and configuration information. Note: This can be implemented as software running self-contained on existing hardware, an external device that is temporarily used, or a permanent inclusion of software left in place for the organization's future benefit.
- Non-Invasive: agent-less (but properly authorized) discovery and data collection tools run with a data collection schedule & monitoring to minimize impact to your infrastructure.
- Security: The secure interface for entering credentials allows system administrators retain control over sensitive information.
- Flexible-Alternatives: Data collection mechanism ensures that required data is collected even when credentialed access is not possible.
- Accelerated data collection: Consolidated and automated data collection appliance simplifies data collection for system inventories, accelerating IBM's wave plan development for IT Optimization support, including migration, and consolidation.
- Streamlined processes: Faster and efficient packaging of collected data for uploading to IBM's automated Analytics Engine.

Discovery and application dependency mapping tools help identify the assets to move together to meet availability requirements. Addressing the need to identify critical issues to reduce risk prior to the move: for example, reducing or eliminating single points of failure.

Analytics

Maximizing efficiency and effectiveness is accomplished through an in-depth analysis of the complete system stack across the complete set of applications and services provided by each data center. The analysis is largely

conducted using these methods, assisted by automation using tools such as SCOPE, Fit for Virtualization/ Cloud, and WTA Analysis.

During the analytics phase, IBM maps assets (such as servers) to business applications, identifies device performance and relationships, researches discovered devices, develops a mapping of the devices to the mission applications supported by the devices, and generates configuration and performance reports and mission application affinity maps.

Analytics Engine: automates application analysis to provide a rapid prioritization of the workloads for migration. IBM's proprietary analytic tools are used to create the dependency maps of the device, network and application infrastructure. Discovery data will be periodically migrated from TADDM, SCMD, and Galapagos to the Analytic Workbench environment to support this analysis initiative. IBM Mission Analysts map the discovered configuration items to mission applications.

The automated tooling used for analytics is focused on system derived data collection for the configurations, utilization, and affinity for the assets discovered. The intelligent use of analytics supports server optimization, consolidation, and virtualization, and can be used independently of data center consolidation/migration, during it, or after data center transitions are complete.

IBM analytics include the use of: server, application, and dependency maps to help identify all assets that need to be moved. Information gleaned from these maps aids in moving dependent applications together, bringing them online together, and migrating data online almost instantaneously to meet application availability requirements. The result is mitigation of the inherent risk in the data center relocation process.

An overall Fit-for-Purpose Roadmap is typically produced, containing:

- Target architecture and configuration
- Source state mapping to target state
- Scheduling constraints (e.g., in-flight project, infrastructure readiness, security, software licensing, etc.)

The benefits of such analytics include:

- Integrated Target Platform Sizing Analyzer that accurately sizes the capacity required while eliminating error prone data handling tasks.
- Single central repository for collected data that is the "golden copy" to all analysis activities.
- Integrated power and cooling analytical function that calculates energy savings using advanced power management techniques.
- Single, comprehensive view of engagement progress and status, reducing time and effort for weekly project status reporting.
- Integrated financial and business case functions that enable architectural decisions to be made by taking both technical and financial factors into consideration.
- Integrated visualization of application dependencies to enable optimal wave plan design.

Planning: Recommendations for Architecture & Implementation

This phase identifies and plans the target state of the business application components database and application server middleware using the data collected in the prior discovery efforts. This methodology uses fit-for-purpose

analytics against the data collected to develop a target state recommendation. The IBM team will analyze all the collected data (applications, databases, environments, servers, etc) as part of our fit-for-purpose approach and construct a Fit-For-Purpose roadmap for a follow-on phase of IT Optimization that provides business justification inputs for consolidation and optimization initiatives.

During the planning phase, information is also collected that only exists in human systems and knowledge bases (ex. Down-time window, application criticality, application constraints, etc.). An Application Profile is pre-populated with application data from data collection to minimize Application Team impact in validation and collecting Non-Functional Requirements (NFRs). Source systems could be added and/or removed from project scope during this process. Newly identified systems are collected as discovered through workshops.

The system-collected database of server affinity and application dependency data gathered from IBM's automated data collection tools makes possible the analytics necessary to determine optimal "migration units" of components to transform, focusing on the interdependencies between applications and each other, and between applications and other components such as data, appliances, middleware, etc. Interdependencies between waves are considered during application assessment during detailed design to determine migration impact management, which may change the grouping (combination/recombination) of servers and apps in each wave to facilitate more optimal interoperability during the transition, or may determine the scheduling of waves in relation to one another.

Projects are sized and prioritized leveraging analytics data by technology diversity in source platform, the number of servers – source and target, migration method to be used, client time and resource dependencies. Risks are mitigated by implementing a phased methodology, starting with complete and detailed analysis of migration scope, and using a standard, proven work breakdown structure supported by the automation method planned for each workload. Generally, prioritization of migration units are based on workload characteristics, TCO analysis, proposed workload placement, and operating system (OS) and middleware modernization.

Transformation Design (Detailed)

During this phase, IBM produces an actionable plan which contains a detailed wave plan with application affinity, a migration workbook, test plan coordination and execution, and cutover and decommission plans. IBM would build a detailed Wave Plan based on Application Affinity and define target architecture configurations.

Execution of Implementation/Test (Install, Provision, Migration, Test)

During this phase, IBM performs installation/provisioning, wave migration, and wave testing. We move the workload into the ultimate environment according to the Wave Plan and the Technical Solution Design. Reliably and completely collecting and using data for associated analysis encompasses all aspects of the source environment to be migrated. This is normally accomplished via small IBM migration teams working with automation tools and the application owner to minimize disruption and expense.

Tools are used to provision virtual servers in target environments with OS

and middleware, using TPM for Images, Platespin, VMware, Mirage, ICON, and CohesiveFT.

Migration of middleware configurations, file systems, and databases are accomplished using tools such as Darwin and UCM.

During test and cutover, remediation and testing is performed on the migrated workloads using tools such as Shift Left.

The benefits of the migration tooling include the following:

- Simplified process: Specialized routines that simplify the discovery and collection of middleware, database, and application configurations.
- Controlled application migrations: Dynamic interface for controlling and directing application migrations.
- Informed communications: Accurate migration status and progress communications.
- Automation: Automated server, middleware, and database provisioning speed cycle time and ensure consistent results.
- Optimization: Application configuration can be optimized for target platform.
- Streamlined testing: Fast and efficient testing of migrated applications.
- Proven processes: Fully tested and validated procedures for provisioning, configuring, and migrating a variety of middleware, database, and other software products.

Such automation reduces the cycle time and labor cost for gathering complete and comprehensive data, and for developing the right design, approach, and business case to support the consolidation and migration objectives.

Summary

IBM's Clients are increasingly realizing the value of IT Optimization for a variety of reasons. The efforts typically consider value from consolidation, migration, operational efficiency, and enabling solutions for the cloud. Most commonly, our clients are driven by the need for cost reduction (e.g. power, space, cooling, and support), increased capacity, or increased flexibilities. They are occasionally driven by broader efforts to follow standards, compliance regulations (e.g. following government regulations and directives), and/or simplification.

IBM's cross-brand specialists, supported by automation tools and methods, provide significant value to our clients by: 1) significantly improving the schedule by which the value can be realized, 2) reducing risks of not realizing the values expected from these activities, and 3) overcoming inhibitors to success. Value can be realized more quickly through repeating processes used successfully in the past, automation, and leveraging skilled and experienced personnel. Inhibitors to success can be overcome when those inhibitors can be made visible and objective. For example, by establishing baselines of IT inventories owned and managed by mission stakeholders, or by revealing the level of maturity in adopting standards that would enable successful migration of assets to a cloud.

.

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
Printed in the United States of America, 05/2008