Scapa® **Test and Performance Platform**

●Stress Testing ●Soak Testing   ●Benchmarking   ●Performance Optimization ●Migration Testing   ●Diagnostic Testing
●Load Testing   ●Scalability Testing ●Reliability Testing ●Bottleneck Identification ●Performance Comparison ●Right Sizing Systems
●Capacity Test ●Performance Testing ●Performance Tuning ●Maximizing User Densities ●Server Consolidation Testing ●Service Availability

Scapa Expedite
Managing Capacity, Continuity and Scalability

A Guide to Testing Activities

CONTENTS

Methodology Overview

Scapa Expedite is a proven methodology for system performance testing, designed to resolve issues of capacity, scalability, continuity and reliability. Evolved over a number of years, by Scapa Technologies' consultants and partner organizations whose experience of performance testing is unrivalled, it minimizes the risk that the capacity of a system is insufficient and will adversely affect the performance, scalability, capacity and continuity of your mission-critical IT systems and their delivery methods. These systems should be assured in order to avoid a significant impact on their return on investment (ROI) as well as end-user morale, service delivery and ultimately, the bottom line.

Performance testing, however, is often viewed as a complex and expensive process, in many cases completely divorced from the rest of IT implementation and management. The Scapa Expedite Methodology has been developed to re-engage performance testing with the rest of the IT function, and to offer significantly higher benefits at lower cost.

The key is to do testing sensibly. By this we mean,

- There is no point in generating a lot of information that cannot be acted upon, or information that is too late to impact on decision making processes.

- The cost of the testing process needs to be in proportion with the value to the business. As a rule of thumb, between 5% and 10% of the total cost of a project should be allocated to performance testing.

The Scapa Expedite methodology is based on a sequence of standard test activities, each of which outlines specific objectives and can be applied singly or in combination at various points in a project where most benefit can be derived. The methodology is designed to help you understand your system's capabilities and, where problems or inadequacies are uncovered, to advise on how to improve performance in a timely and cost-effective manner.

Planning

Scapa Expedite defines **8** types of test activity which cover a broad range of requirements. Generally, it will be the Project Manager, the Architect or Technical Design Authority and the Scapa Expedite leader who will align the test activities with the overall project timelines and introduce these into the project plan, following agreement from other stakeholders within the business, where required.

The following are factors that also need to be considered during planning:

- Known architectural and application risk factors

- Business priorities and risks

- Timing of project review points

- Scoping of tests to allow them to deliver maximum relevant information at project review points

Not all of the test activities are required in every project. In general, we would recommend, for even the smallest changes to a system, test activities should be performed at least twice, once at the beginning and once at the end of the implementation.

In small-scale projects where hardware and software architecture and sizing are well-understood, or can easily be taken from best-practice guidelines, a single test (known as an **Operational Acceptance Capacity Check**) can be used to confirm that the system has the necessary capacity. However, this is a high-risk strategy because the information comes too late to allow changes to be made.

Test Activities

The eight test activities in Scapa Expedite are:

• Start with simple tests. Simple tests produce a lot of rudimentary information and basic metrics useful for fine tuning.
• Introduce complexities to your tests in stages, rather than all at once.
• Following simple tests with the gradual addition of complexity allows you to run tests repeatedly and achieve fast test cycles.
• Try to use real data over synthetic data whenever possible. Real data produces more accurate results.
• Keep script creation time to a minimum - a few hours if possible!
• Rather than test one tier of your Remedy system at a time, aim to include all tiers of the Remedy system in the test.
• Strive to get results as early as the system can generate them. Early results mean more time spent with the system and more time means increased confidence in using the system.
• Understand what the results mean and how these may affect the system.
• Limit the time spent on an each test cycle. With careful planning a single testing cycle can be completed in a day.

Each test activity contains one or more actual tests which, following the BCS ISEB terminology, are either:

• Performance Tests, at fixed workloads anticipated to be within the capacity of the system

or

• Stress Tests, to drive the system beyond its capacity

Baseline Service Performance Test

Objective: To find out how the system behaves at a very low level of activity. If performance problems are seen at this level and at this stage, there is time to rectify issues with the application or the infrastructure.

Test Type: Performance test with extremely low load.

Project stage: Typically performed as early as possible on the first available environment, and may be repeated at various times during the project.

Scope: A few diagnostic transactions (e.g. log-in, logout, and perhaps a read-only transaction) for the key business application(s).

Scale: 1-10 users, small-scale environment.

Outcome: End-to-end service level metrics, and operational metrics for the diagnostic transactions. If the end-to-end service metrics are out of ballpark, the operational metrics can be used to analyze root cause. Some linear extrapolations can be made on the operational metrics, for example memory and CPU usage, to give early warning of capacity problems.

Sizing

Objective: Predict volume of hardware required to support defined levels of capacity.

Test Type: Stress test on one or more hardware configurations.

Project Stage: Best performed during design stage, when hardware configurations and associated budgets still have some flexibility.

Scope: Diagnostic transactions covering core system function (log-in, logout, basic read and write transactions) for key business application(s).

Scale: Single servers or small clusters/farms

Outcome: For proposed hardware configurations, the tests map the edge of the "capacity envelope" for the diagnostic transactions and derive associated operational metric signatures. Sizing tests provide clues to which elements of a workload are likely to give capacity problems. When used with reference to the anticipated production workload, these tests can also be used to predict hardware requirements. Some headroom should be applied to deal with approximations in the comparison between the diagnostic transactions and the anticipated production workload.

Note: The "capacity envelope" defines the area where system performance is acceptable, and outside of which performance becomes unacceptable.

Mid-Implementation Capacity Check

Objective: To validate sizing and identify any bottlenecks, whilst they can still be fixed.

Test Type: Performance test.

Project Stage: Should be performed roughly half-way through an implementation.

Scope: Diagnostic transactions covering good system function (login/logout, read, write, search) for key business application(s) and any applications that are viewed as potentially problematic. Additional elements such as login authentication and external system connectivity should be present in the architecture, if possible.

Scale: A system at around 30% or more of the size of the projected production environment at a load that the sizing process predicted would be within its capacity. Databases should contain realistic volumes of data.

Outcome: For the tested environment, a map is produced of the edge of the capacity envelope for the diagnostic transactions and the associated operational metric signatures. Elements of a workload likely to give capacity problems are pinpointed and, by reference to the anticipated production workload, can be used to determine areas for urgent fixes in the architecture. Some headroom should be applied to deal with approximations in the comparison between the diagnostic transactions and the anticipated production workload.

Production Calibration Test

Objective: For systems that are being migrated from an existing production system, this check confirms that the workload being used to test the new system can adequately predict the capacity of the production system.

Test Type: System metrics (and where possible transaction response times) are measured during a performance test on the new environment and compared with equivalent data collected from the production environment, either under production load or under the same performance test.

Project Stage: Typically associated with a capacity or sizing activity.

Scope: Same as associated test activity (capacity or sizing) scale.

Outcome: A set of expressions relating the performance of the production system and the test system under the same load. For example,
 "For a read-only workload at 25% of test system capacity the response time of the test system is 25% slower than the production system, and the CPU level is twice as high".

Typically between 5 and 10 such expressions are adequate for predicting production system capacity and performance, based on measurements taken in a test environment.

Capacity-Related Stability Check

Objective: To identify any ongoing "creep" in the application or environment that may lead to instability in production. Issues may be subsequently resolved by changes to the application or environment, or mitigation processes may be designed for use in operations.

Test Type: A long-running performance test, typically preceded by a stress test to identify an appropriate load.

Project Stage: From roughly half-way through an implementation.

Scope: Good system function (login/logout, read, write, search) for the key business application(s) and any applications that are viewed as potentially problematic.

Scale: Usually performed on a small part of the ultimate environment, although larger-scale tests may be required to determine stability of back-ends.

Outcome: Confirmation that the system is able to handle long periods of realistic user activity without suffering problems such as, memory leakage, handle leakage, or loss of stability. In some cases, where systems do degrade over time, it is necessary to suggest reboot schedules.

Session Stability Test

Objective: To identify any sporadic events which cause occasional unplanned termination of application or infrastructure sessions. This is typically required in remote desktop or virtual desktop environments, where a service level agreement (SLA) may be in place.

Test Type: A series of long-running performance tests.

Project Stage: Close to production, infrastructure should be in place (the correct routers, access gateways, authentication infrastructure, etc.). The test should be performed from a representative set of user locations (e.g. from multiple remote sites). The state of the applications is less critical.

Scope: Basic system function (e.g. log-in, logout, and perhaps a read-only transaction) for key business application(s).

Scale: Usually performed with a few hundred user sessions to ensure statistical viability, at one or more levels of load, from idle sessions through to realistic usage levels.

Outcome: Confirmation that the users will be able to log onto the system for extended periods without their sessions being dropped.

Scapa
Technologies

Operational Acceptance Capacity Check

Objective: To determine whether or not the system has enough capacity to go live.

Test Type: Performance test.

Project Stage: Immediately prior to the system going live.

Scope: A broad set of system functions for key business application(s) plus some functions for other applications as determined by risk analysis (based on business analysis and previous testing).

Scale: Anticipated production capacity, plus a defined level of headroom. On the production environment before it goes live, or the largest "domain" within a production environment.

Outcome: Final confirmation that the capacity of the system is adequate to produce the required level of performance and availability.

Operational Acceptance Continuity Check

Objective: Often employed as part of a compliance and/or a disaster recovery strategy to determine the resilience of the system when one or more system components are inoperative.

Test Type: Performance test during which elements of the system are deliberately rendered inoperative. Operational, end-user response time and availability are also measured.

Project Stage: Typically performed after an operational capacity test.

Scope: Core system function (log-in, log-out, basic read and write transactions) for the key business application(s).

Scale: Anticipated production capacity, plus a defined level of headroom. On the production environment before it goes live, or the largest "domain" within a production environment.

Outcome: Confirmation, or otherwise, that the system is capable of maintaining adequate capacity and availability during certain forms of failure. In cases where this is not the case, additional analysis would be performed to measure the risk to the business and suggest possible fixes (with associated costs).

Consulting Services

Scapa Expedite Methodology

The Scapa Expedite Methodology has been developed and is used by our highly skilled consultants and partner organizations working worldwide, either on-site or remotely to deliver a variety of performance testing activities across a wide-range of applications and environments, from off the shelf to custom applications.

Scapa consultants and partners specialize in performance testing of BMC Software® Remedy-based solutions, such as the ITSM® suite of applications, systems based on the Mendix™ application platform, web applications and desktop virtualization deployments (e.g. Citrix® XenApp ™/XenDesktop®, VMware® View™ and Microsoft® Remote Desktop Services® and Terminal Server®).

Our services include:

- End-to-end Testing (load, stress, capacity, performance, scalability, reliability, soak)
- Scalability Profiling
- Bottleneck Identification
- Benchmarking
- Capacity Planning
- Performance Comparison
- Performance Tuning
- Performance Optimization
- Maximizing User Densities
- Migration Testing
- Server Consolidation Testing
- Service Availability
- Diagnostic Testing
- Desktop Virtualization
- Right Sizing Systems
- Proof-of-concept

About Scapa Test and Performance Platform

Scapa Test and Performance Platform (TPP) features real data utilization, quick test creation, and dynamic multi-load manipulation among its capabilities, making it ideal to emphasize the influence of user data and for rapid test creation and real time analysis which are key concepts of an effective testing approach.

For a full list of features and to learn more about the Scapa Test and Performance Platform visit
http://www.scapatech.com