

SAFEPEAK TECHNOLOGIES LTD



**Accelerating Performance  
of Custom and 3<sup>rd</sup>-Party  
SQL Server Applications  
with  
SafePeak Automated Dynamic Caching**

**A SafePeak Whitepaper**

February 2014

Copyright. SafePeak Technologies 2014



## Contents

Objective.....	3
Introduction.....	3
SafePeak® - Immediate Acceleration of SQL Server Operational Apps .....	4
Automated Dynamic Caching for SQL Server OLTP Apps.....	4
How SafePeak Works - Query Process Flow.....	5
Scenario I – Query Result Returned By SafePeak .....	5
Scenario II – Query Sent to the Database; Result Stored In SafePeak Memory .....	6
Scenario III – Eviction of Results Set In Cache; Update to Database .....	7
SafePeak High Availability.....	8
SafePeak Dashboard – Deep Real-Time Performance Analysis.....	9
Telco Case Study: Upgrading Infrastructure Scalability and Utilization.....	10
Summary .....	11



## Objective

This paper describes the SafePeak solution, which delivers **plug-and-play performance acceleration and scalability for SQL Server OLTP / Operational applications**. This paper will explain how SafePeak can instantly enable any type of organization to resolve information access bottlenecks and latency, without any change to the existing applications or databases. In addition, the paper will clarify how the SafePeak solution enables organizations to resolve application and database performance and scalability challenges within hours, while preserving all data integrity and consistency.

*View an online demo: <http://www.safepeak.com/Product/Safepeak-Demo>*

## Introduction

The requirement for faster data access and application response times, along with the need for high availability of critical business information, has grown exponentially. In many industries, fast response time and continuous availability are an important competitive differentiator and a prerequisite for doing business. These industries include but are not limited to ecommerce, publishing, broadcasting, rich media, finance, medical, travel, online games and more.

Reduced costs of data storage and the advancement of technology for capturing and storing information are leading to the accumulation of massive amounts of data in organizational databases. The demand for increased database performance, increasing volumes of corporate information and requirements for operational efficiency provide a growing challenge for IT professionals, DBAs, architects and development managers. The database and application infrastructure needs to meet or exceed performance expectations while scaling to meet growing business needs.

While the requirement for faster data access is ever present, so are the causes of slowdowns and downtime for applications and databases, and these significantly affect businesses. Outages reduce revenues, reduce employees' productivity, and lower the quality of customer service, all of which can damage the corporate brand.

This paper reviews SafePeak's unique software capabilities to dramatically accelerate SQL Server-based applications by using innovative dynamic caching, sophisticated auto learning and automatic adaptation algorithms.



## SafePeak® - Immediate Acceleration of SQL Server Operational Apps

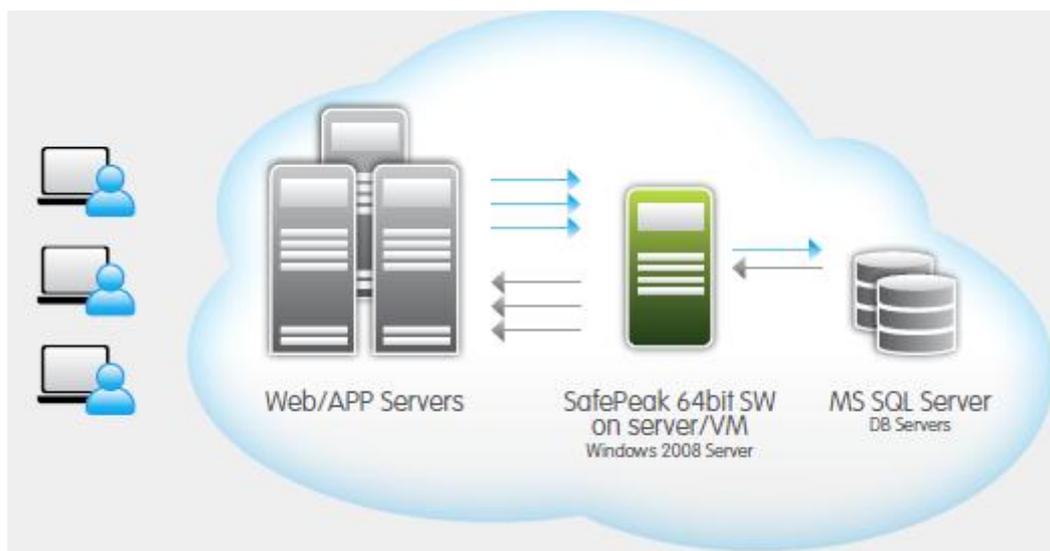
SafePeak provides immediate resolution of several SQL Server issues:

- Performance
- Query access speed
- Data bottlenecks
- CPU/IO load and latency

SafePeak dramatically accelerates data access to microsecond speeds, providing immediate benefits without any changes to application code or the database. SafePeak enables companies to maximize operational efficiency of existing infrastructure and scale 5x to 10x more data and work load, avoiding costly and complicated hardware and software upgrades. SafePeak ensures high availability for mission-critical custom and 3<sup>rd</sup>-party applications (like SharePoint, or CRM Dynamics), and safeguards against traffic spikes.

### Automated Dynamic Caching for SQL Server OLTP Apps

SafePeak delivers value through the implementation of an innovative solution for dynamic database caching, using sophisticated auto learning and automatic adaptation algorithms.



SafePeak's solution is deployed on Windows 2008 64bit virtual machine, a server (or two servers for an HA cluster), and acts as a seamless proxy between the applications and SQL Servers. Patent-pending auto-learning algorithms analyze connected SQL Server instances and study the application traffic, create patterns of SQL queries and stored procedure calls, to understand their nature and dependencies and to build caching rules for the SafePeak Core engine.

SafePeak offloads repetitive "Read" queries and stored procedures calls (queries and procedures that "select" data), while storing the result-sets in SafePeak's cache memory.

Requests for the same queries are not sent to the database servers and the result is retrieved from the RAM based cache at instant speed (microseconds). SafePeak's sophisticated algorithms manage the changes of data completely automatically, evicting in real-time the relevant items from cache with arrival of "write" commands (DML, DCL, DDL) ensuring 100% data integrity and 100% ACID compliance.

Database CPU and IO load is reduced dramatically, freeing important resources, leading to additional performance acceleration for both reads and writes.

The benefit from SafePeak's solution is immediate: Within few hours of installation, auto-learning and configuration fine-tuning your application gains significant performance acceleration and your database sees a scalability boost. Since SafePeak is transparent to the application or web server and database, the installation process requires minimal effort. No code changes in the application or database are required.

### How SafePeak Works - Query Process Flow

When a SQL query is issued from an application, SafePeak intercepts the query and determines whether to direct it to the SQL Server database for processing or to return a response from the SafePeak Cache Manager.

SafePeak examines whether the query is a repetitive read query whose results have been stored in SafePeak's RAM memory or if it is new and needs to be sent to the SQL Server database for retrieval. SafePeak also checks to see if it is a write request (update, insert, delete or other DDL/DCL types) that will cause a change in the target database (and possibly the cached results) and needs to be forward to the SQL Server database for execution. The major scenarios and flows are described in detail below:

#### Scenario I – Query Result Returned By SafePeak

The first action that SafePeak takes is determining if the transaction contains a repetitive query whose result set is stored in the SafePeak Cache Manager in RAM memory. If the query is found to exist, the result set is retrieved from the Cache Manager (C1) and returned to the querying application (C2). SafePeak takes no further action and the query never needs to reach the target database. In this scenario, the query cache holds the exact results that are sent to the querying application in a low level binary result set.

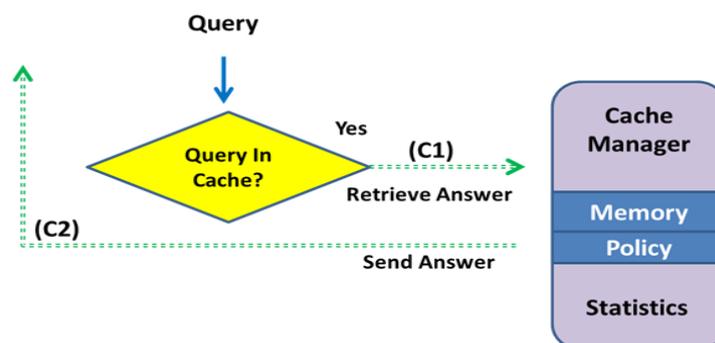


Figure 1 - Query Result Returned By SafePeak

The ability to rapidly retrieve identical result sets significantly improves response time, reduces overall network and database traffic and boosts system scalability, especially at times of peak usage and demand spikes.

### Scenario II – Query Sent to the Database; Result Stored In SafePeak Memory

In scenario two, SafePeak checks and determines that the query and result set are not stored in the Cache Manager. SafePeak continues to process the query and determines whether the request is a read query or a write request. In the scenario illustrated in figure 2, we have determined that the request is a read query (Q1).

SafePeak takes several steps at this point. The first step is to process the query on the target database and return the result set to the querying application to ensure the most rapid response possible (Q2).

Once the information has been sent to the querying application, SafePeak determines if the query is a repetitive cacheable query. If it is, then SafePeak saves its result set in binary code inside the RAM memory of the Cache Manager (Q3) to be accessed upon the next instance of the identical query.

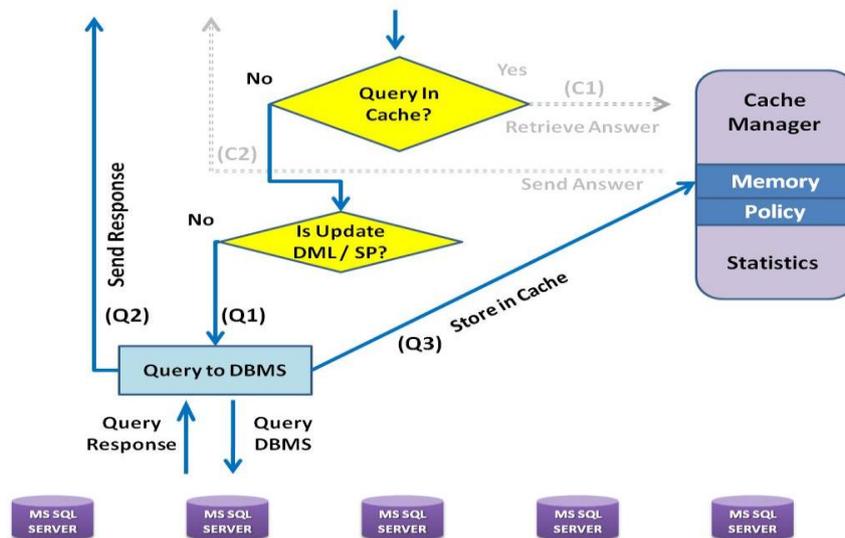


Figure 2 - Query Result Returned by Database; Result Stored In Cache Manager RAM

### Scenario III – Eviction of Results Set In Cache; Update to Database

In scenario three, SafePeak determines that the incoming query is an update, insert, alter or any other request that may cause a change in the database. In this case, SafePeak dissects the request and decides which tables in the database may be impacted by its execution. It then looks at the query results stored in the Cache Manager and evicts all results that have any connection to the impacted database tables (U1).

Once the Cache Manager has been cleaned to ensure data credibility and accuracy, the request is sent to the SQL Server database and executed (U2). The result set of the executed response is then sent back to the querying application (U3). By handling the update requests and eviction of cached result sets in this fashion, SafePeak is able to ensure the highest levels of data integrity and consistency. While the transaction is in progress SafePeak's Cache Manager is locked, preventing new queries to be inserted to the cache with relation to the objects affected by this update request.

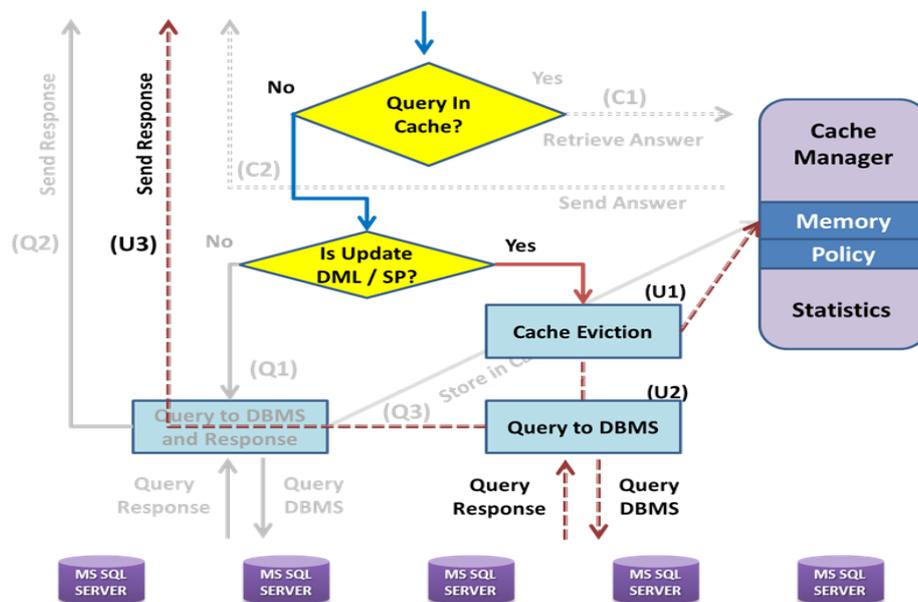


Figure Three – Eviction of Results Set In Cache; Update to Database

## SafePeak High Availability

SafePeak runs in mission critical, enterprise production environments providing continuous information flow. It is crucial to ensure the highest level of reliability and data availability with protection against both hardware and software failures and malfunctions.

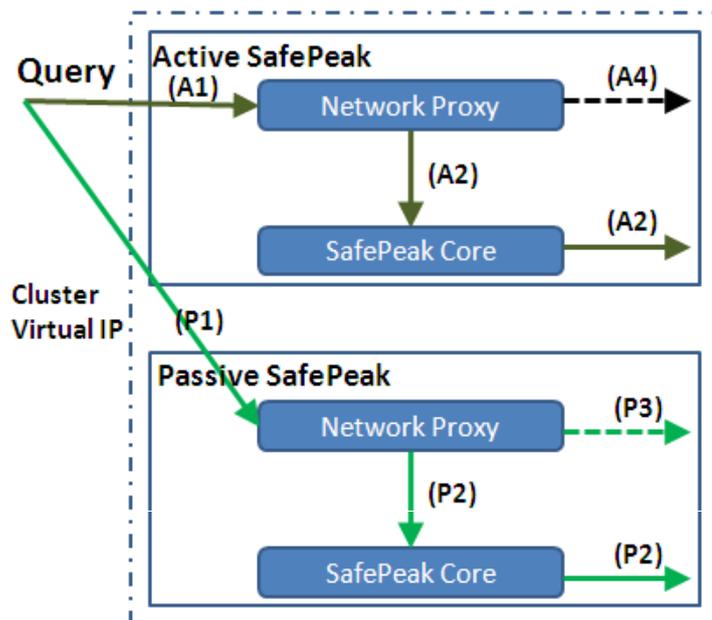
SafePeak provides a high availability solution, by utilizing a clustered deployment. Protection against hardware malfunctions requires a SafePeak Cluster installation, easily deployed and managed on two dedicated commodity stand-alone servers (no need for special servers, load balancers or shared storage). Protection against software failures is automatically embedded in both single and cluster installations.

SafePeak solution monitors and adjusts to system failures or malfunctions and provides automatic failover protection without needing IT intervention. It ensures the rapid recovery of critical functions for continuous database processing.

SafePeak's high availability solution supports various clustering or load balancing options, by configuring an active/passive cluster between the front-end client application and the back-end SQL Server database as shown in Figure 4 below.

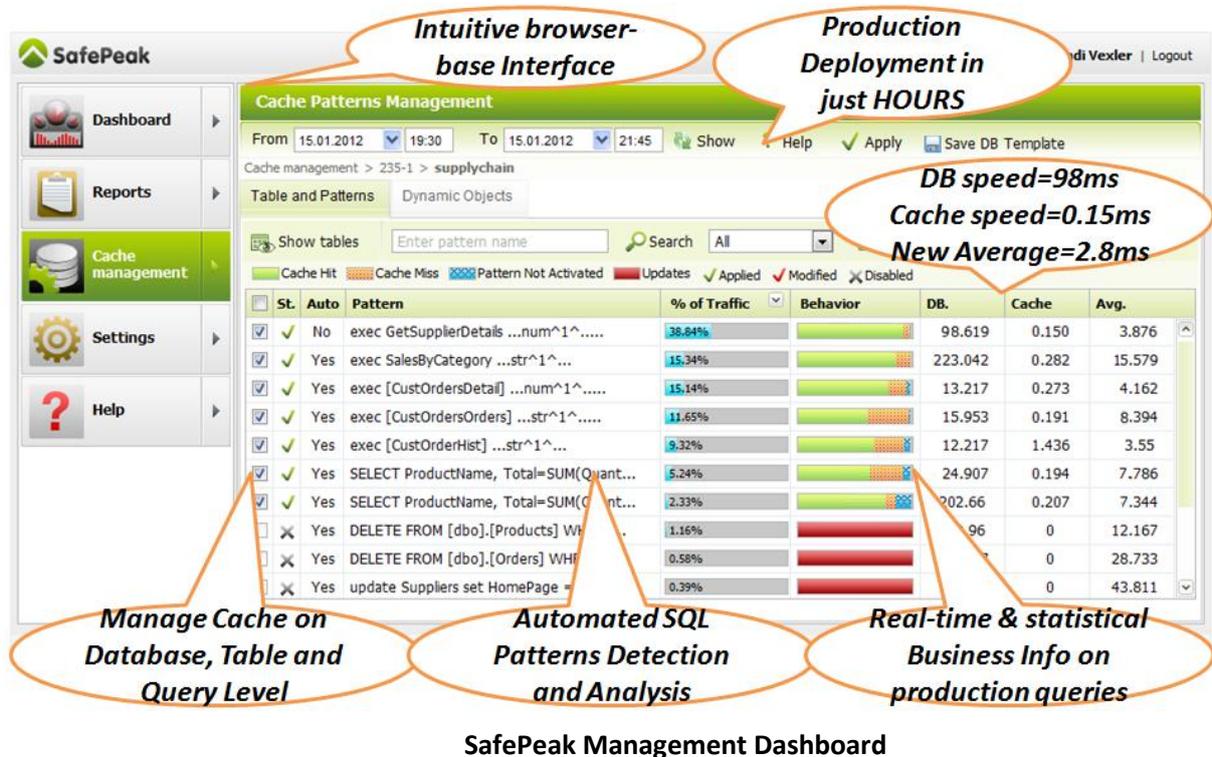
In normal operation, the incoming query is routed via the Network Proxy (A1) where the request is forwarded to the SafePeak Core for processing (A2). In the rare event of a SafePeak software failure, processing will be redirected to the database (A4).

If a hardware error or operating system error occurs due to a malfunction, processing is automatically shifted into failover mode (P1) ensuring that the loss of the server will not impact continuous processing. The queries are routed to the backup SafePeak instance. Processing at this point is handled by the SafePeak Core (P2) as long as SafePeak is fine. Otherwise, processing will continue directly from the Network Proxy to the SQL Server database (P3).



## SafePeak Dashboard – Deep Real-Time Performance Analysis

SafePeak comes with an intuitive browser-based user interface dashboard for configuration, management, tuning, reporting, real-time and historical analysis allowing both database and SafePeak performance analysis with a focus on finding bottlenecks. The auditing and analysis mechanisms are built into SafePeak and do not add additional performance overhead.



The dashboard is used for many aspects of the SafePeak operation, including the following:

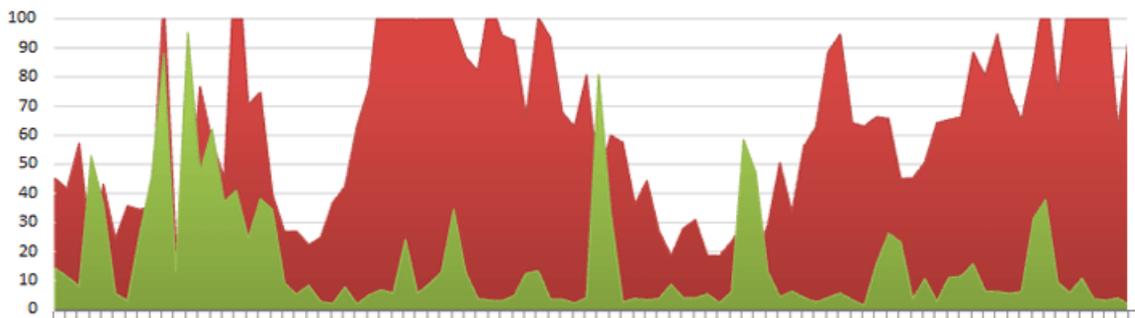
- SafePeak Management:
  - Configuration of multiple SQL Server instances managed by one SafePeak server
  - Enabling and disable caching for specific databases, tables (queries that access the tables), queries. Changing cache policy rules, such as default timeout, special eviction scheduling, etc.
  - Fine-tuning of dynamic stored procedures: objects dependencies, behavior etc.
  - Global cache start/stop, reload configurations
- Performance monitoring and analysis tools
  - Real time load information, performance and throughput statistics and reports
  - Analysis on different levels: instance, databases, tables and views, up to single query level: response time, percent of usage, dependent objects etc.

SafePeak's management dashboard provides DBAs, IT Managers, Architects and Developers with critical and valuable information. This information includes a drill down on database instances regarding usage, accessed tables and length of execution of queries. It also helps to uncover queries that are incorrect or those which can be fine tuned for better processing efficiency.

## Telco Case Study: Upgrading Infrastructure Scalability and Utilization

A trial conducted in a large Telco company (20,000 employees; over \$5B market cap) provides a clear benchmark regarding SafePeak's solution value and benefits:

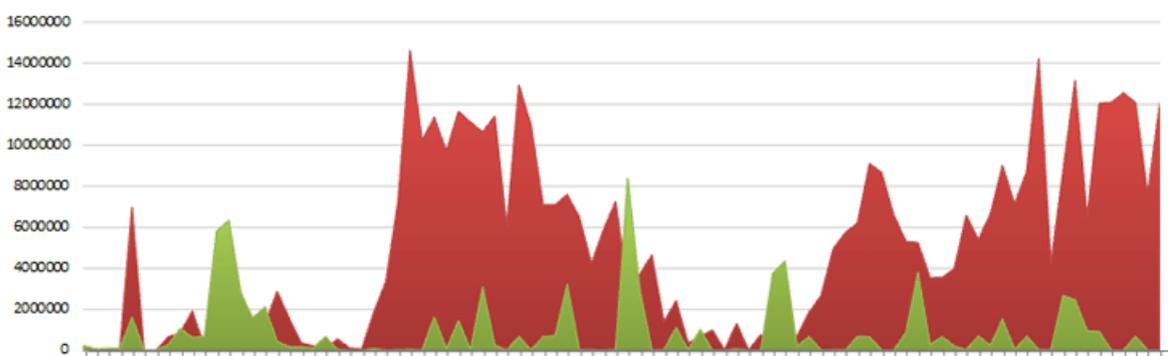
**Graph A: SQL Server: % Processor Time. SafePeak ON (green) and SafePeak OFF (red)**



Red represents the CPU load before SafePeak was integrated. Green represents the results after SafePeak was integrated and activated.

CPU consumption was reduced on average by 75%, enhancing the capability of the existing infrastructure to handle much higher workloads while also delivering superior performance results.

**Graph B: SQL Server: IO Data Bytes/sec. SafePeak ON (green) and SafePeak OFF (red)**



Red represents the CPU load before SafePeak was integrated. Green represents the recording after SafePeak was integrated and activated.

IO consumption was reduced on average by 85%, enhancing the overall infrastructure capability to handle much higher workloads while delivering superior performance results.

