

SAFEPEAK TECHNOLOGIES LTD



White Paper



Contents

Contents.....	2
Introduction	3
SafePeak	4
How SafePeak Works - Query Process Flow	5
Scenario One – Query Result Returned By SafePeak	5
Scenario Two – Query Result Returned by Database; Result Stored In SafePeak Memory	6
Scenario Three – Eviction of Results Set In Cache; Update to Database	7
Ensuring High Availability	8
SafePeak Management Dashboard	9
SafePeak Background Processes	10
Conclusion.....	11



Introduction

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

Significantly reduced costs of data storage and the advancement of technology for capturing and storing information, has led to the accumulation of massive amounts of data in organizational databases. The demand for increased database performance coupled with the increasing volumes of corporate information provides a growing challenge for IT professionals. The database and application infrastructure needs to meet or exceed performance expectations while scaling to meet growing business needs.

And while the requirement for faster data access is growing, so are the slowdowns and downtime of applications and databases that significantly affect business. These outages can reduce revenues, employee productivity and lower the levels and quality of customer service and support, often leading to the damaging of the corporate brand.

This paper describes SafePeak, a new patent-pending software paradigm from DCF Technologies Ltd., which enables customers to achieve and exceed their business goals by providing Plug and Play software for immediate resolution of information access bottlenecks and latency, without any changes needed to existing applications or databases. With SafePeak, organizations can resolve their database bottlenecks within hours and days, with significant measured performance improvements in as little as a week.



SafePeak

SafePeak provides immediate resolution of SQL Server query access and retrieval bottlenecks and latency. It dramatically accelerates data access and retrieval providing immediate results, with no need to change existing databases or applications. It enables companies to maximize their investment in existing infrastructure while avoiding costly hardware and software upgrades. Increasing database response time from SQL database queries and enhancing system scalability by orders of magnitude, SafePeak ensures high availability for mission-critical applications and safeguards against unpredictable traffic spikes and surges of usage.

SafePeak enhances scalability by offloading the handling of "Select" query traffic from operational databases, significantly increasing database capacity for handling transactional queries and updates. This serves to increase overall data and application throughput, while enhancing transaction efficiency. SafePeak leverages existing hardware and software to improve application and database performance with enhanced scalability. It is especially appropriate for Internet applications which experience unpredictable and varying peaks of usage. In addition, SafePeak helps to ensure real-time performance and high availability for traditional businesses running OLTP applications.

Large web sites and internal corporate systems often need to handle thousands of queries per second. To keep up with this immense workload and improve scalability, SafePeak utilizes a patented technique in which cache results of repeated identical queries are stored in binary result sets in read access memory (RAM) for rapid retrieval.

SafePeak acts as a proxy cache between application servers and database servers. Patent-pending auto-learning caching algorithms store SQL statements results in SafePeak RAM, greatly reducing database-server load, eliminating traffic bottlenecks and increasing scalability and transaction throughput.

A cache coherence protocol ensures 100% data integrity for both read and write database transactions. It enables scaling-out of database workload by simply adding additional SafePeak instances, providing a quick and cost-effective solution to customers' growing database requirements.

How SafePeak Works - Query Process Flow

When a SQL query is issued from a web-based or enterprise application, SafePeak intercepts the query and determines whether to direct the query to the SQL Server database for processing or to return the query 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; whether the query is new and needs to be sent to the SQL Server database for retrieval; or whether the query is a write request (update, insert, delete or other DDL/DCL types) that causes 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 One – Query Result Returned By SafePeak

The first action that SafePeak takes is to determine whether 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).

No further action is taken by SafePeak 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. When the query comes in, SafePeak does a very fast check to see whether the query is identical and then sends the stored results as in figure 1 below.

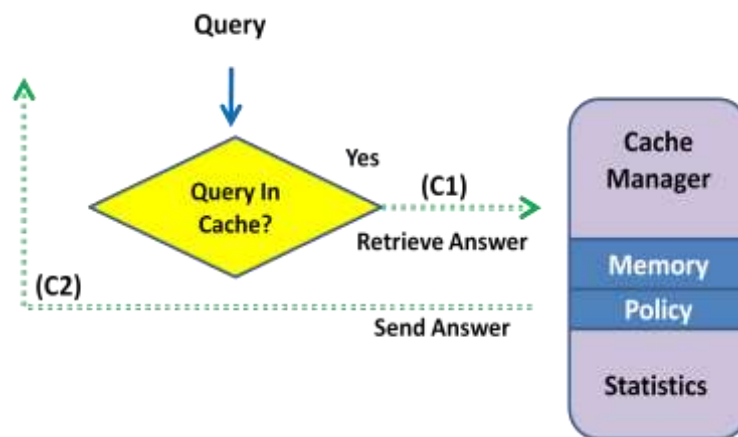


Figure 1 - Query Result Returned By SafePeak

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

Scenario Two – Query Result Returned by 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 below in figure 2, we have determined that the request is a read query (Q1).

SafePeak takes several steps at this point. The first step taken by SafePeak 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.

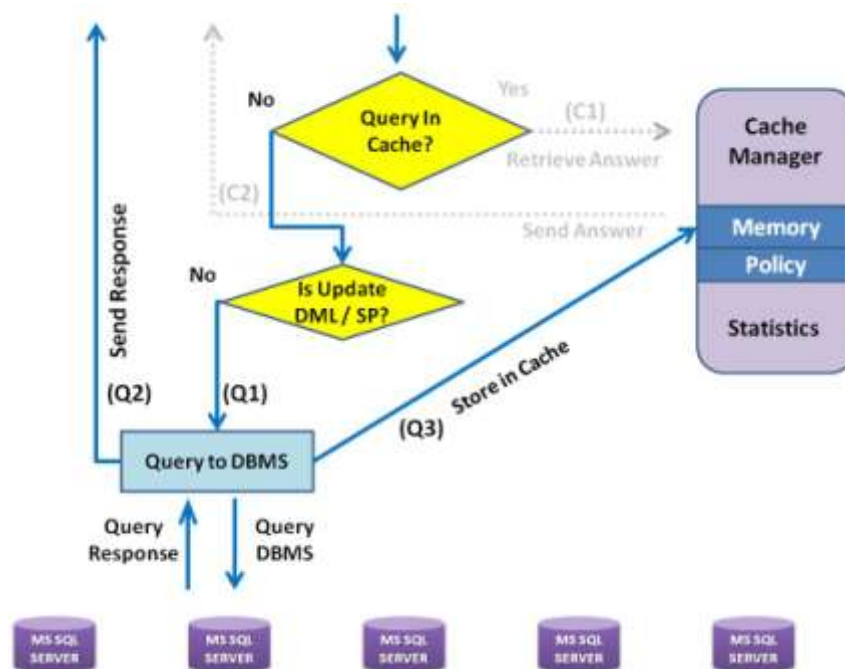


Figure 2 - Query Result Returned by Database; Result Stored In Cache Manager RAM

Scenario Three – 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 affected 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 affected 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.

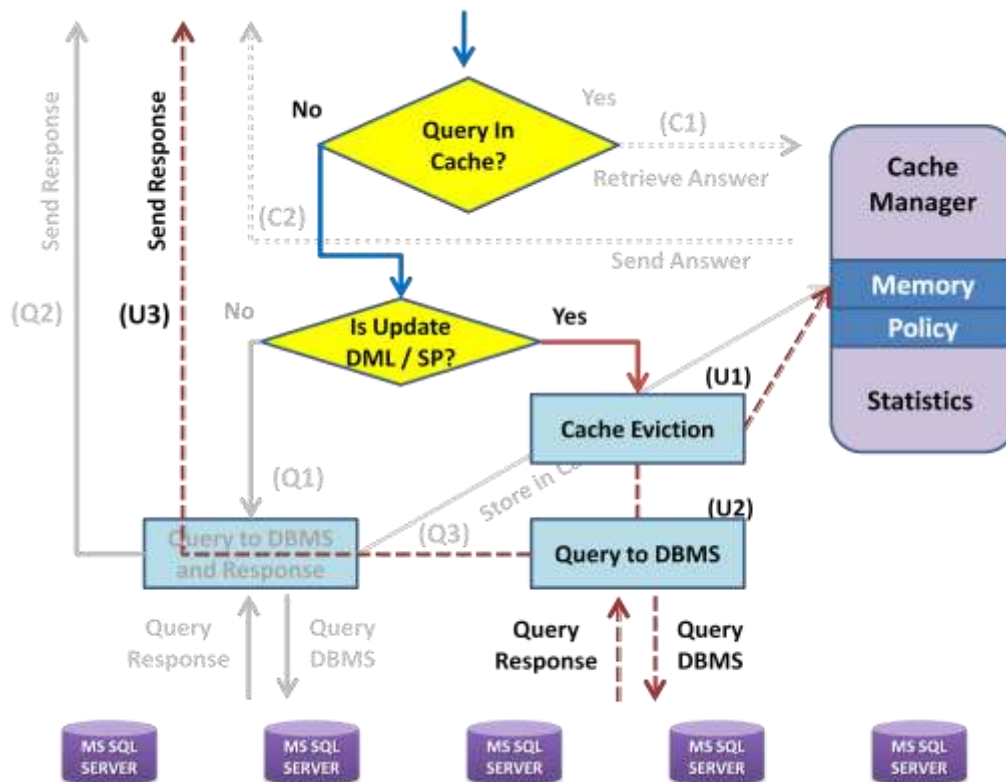


Figure Three – Eviction of Results Set In Cache; Update to Database

Ensuring High Availability

As SafePeak runs as a server in production environments, it is critical that SafePeak protect against both software failures as well as hardware malfunctions to ensure continuous information flow. To ensure data availability and to prevent data loss from outages, SafePeak provides a high availability cluster solution based on a reliable, high performance load balancing and failover proxy server in which each component is deployed on virtual machines. SafePeak's solution monitors and adjusts to system failures or malfunctions and provides automatic failover protection without the need for IT intervention. It ensures the rapid recovery of critical functions for continuous database processing.

SafePeak High Availability at Work

SafePeak's high availability solution is based on a load balancing network proxy deployed in an active/passive cluster configuration between the front-end client application and the back-end SQL Server database as shown in Figure 4 below. This high availability solution has been optimized and tested to run resource intensive applications and databases with minimal performance overhead.

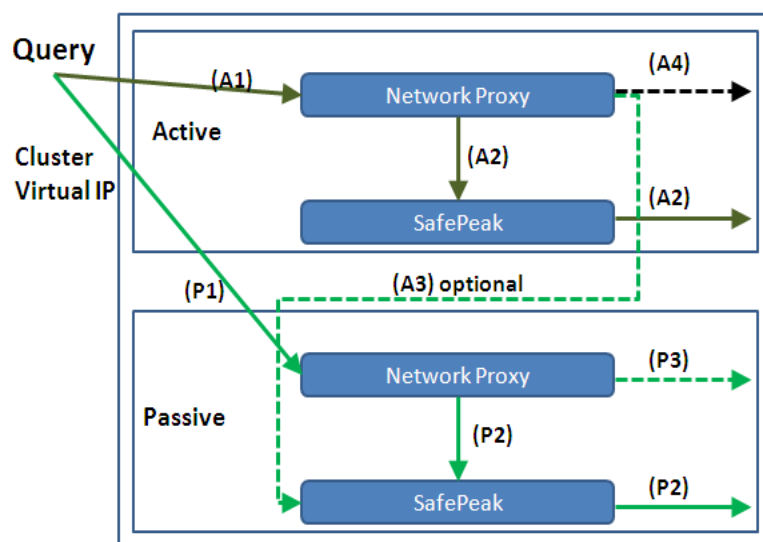


Figure 4 – SafePeak High Availability Architecture

In normal operation, the incoming query is routed via the Network Proxy (A1) where the request is forward to SafePeak for initial processing (A2). In the rare event of a SafePeak software failure, processing can take one of two routes. It can either go directly to the database (A4) or it can be automatically redirected directly to the backup SafePeak instance (A3). If a hardware error occurs due to a malfunctioning server, processing is automatically shifted into failover mode (P1) ensuring that the loss of the server will not impact continuous processing. The query is then routed to the backup SafePeak instance. Processing at this point is handled by SafePeak (P2) as long as SafePeak is installed. Otherwise, processing will continue directly from the Network Proxy to the SQL Server database (P3).

SafePeak Management Dashboard

SafePeak provides a web-based GUI management dashboard for configuration, management, operation, reporting and statistical analysis.



SafePeak Management Dashboard

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

- Configuration and addition of cache policy rules
- Management of rules such as default timeout, eviction scheduling, etc.
- Management of query and object dependencies for eviction and decision making
- Setting changes and management for cache start/stop, reload configurations, etc.
- Configuration of new instances of SafePeak
- Performance and throughput statistics and reporting

The management dashboard provides DBAs with 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.

SafePeak Background Processes

SafePeak utilizes a set of background processes which interface with each other and with the SafePeak repository as illustrated in Figure 5.

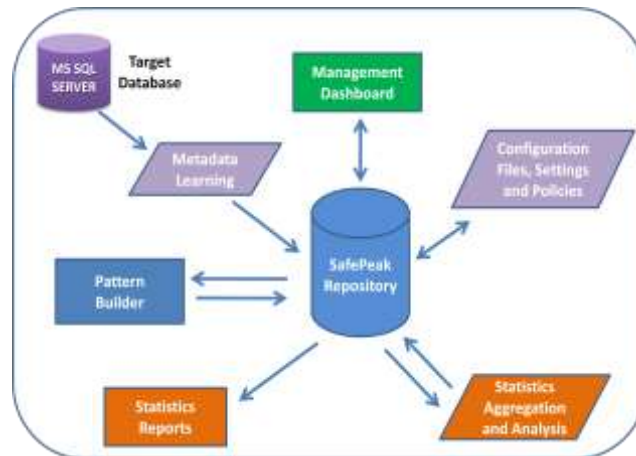


Figure 5 - SafePeak Background Processes

The background processes are used at varied times during SafePeak operations depending on the stage of operation. The major background processes are as follows:

Pattern Builder	SafePeak's Pattern Builder parses incoming queries to identify repetitive identical queries for and determine which result sets should be added to RAM in the cache manager.
Metadata Learning	A mechanism for determining the structure and attributes of objects in the target SQL Server database including tables and views, columns, dependencies and constraints, indexes and keys.
Configuration Files, Settings and Policies	Configuration files are automatically created and configured during the installation of SafePeak. These files contain the settings and policies that determine how queries are handled, cache override settings, policies for eviction, and more. The configuration files are updatable from the management dashboard.
Statistics Aggregation, Analysis and Reporting	Aggregates cache entries and evictions and read, write and timing statistics, for performance fine tuning and reporting.



Conclusion

SafePeak provides a revolutionary approach to database performance challenges. It enables customers to achieve and exceed their business goals by providing Plug & Play software for immediate resolution of information access bottlenecks and latency, without any changes to existing databases or applications.

By offloading repetitive queries from the database, SafePeak significantly improves data retrieval performance while decreasing database resource consumption. This also serves to improve the response time of other database operations, increasing overall data and application throughput and scalability.

And to ensure continuous information availability and to prevent data loss or outages, SafePeak provides a high availability cluster solution based on a reliable, high performance load balancing proxy server. SafePeak's solution monitors and adjusts to system failures or malfunctions and provides automatic failover protection without the need for IT intervention. It ensures the rapid recovery of critical functions for continuous database processing.

To find out more about SafePeak, please visit www.safepeak.com.