SAP HANA

SAP HANA® Platform – Technical Overview

Driving Innovations in IT and in Business with In-Memory Computing Technology

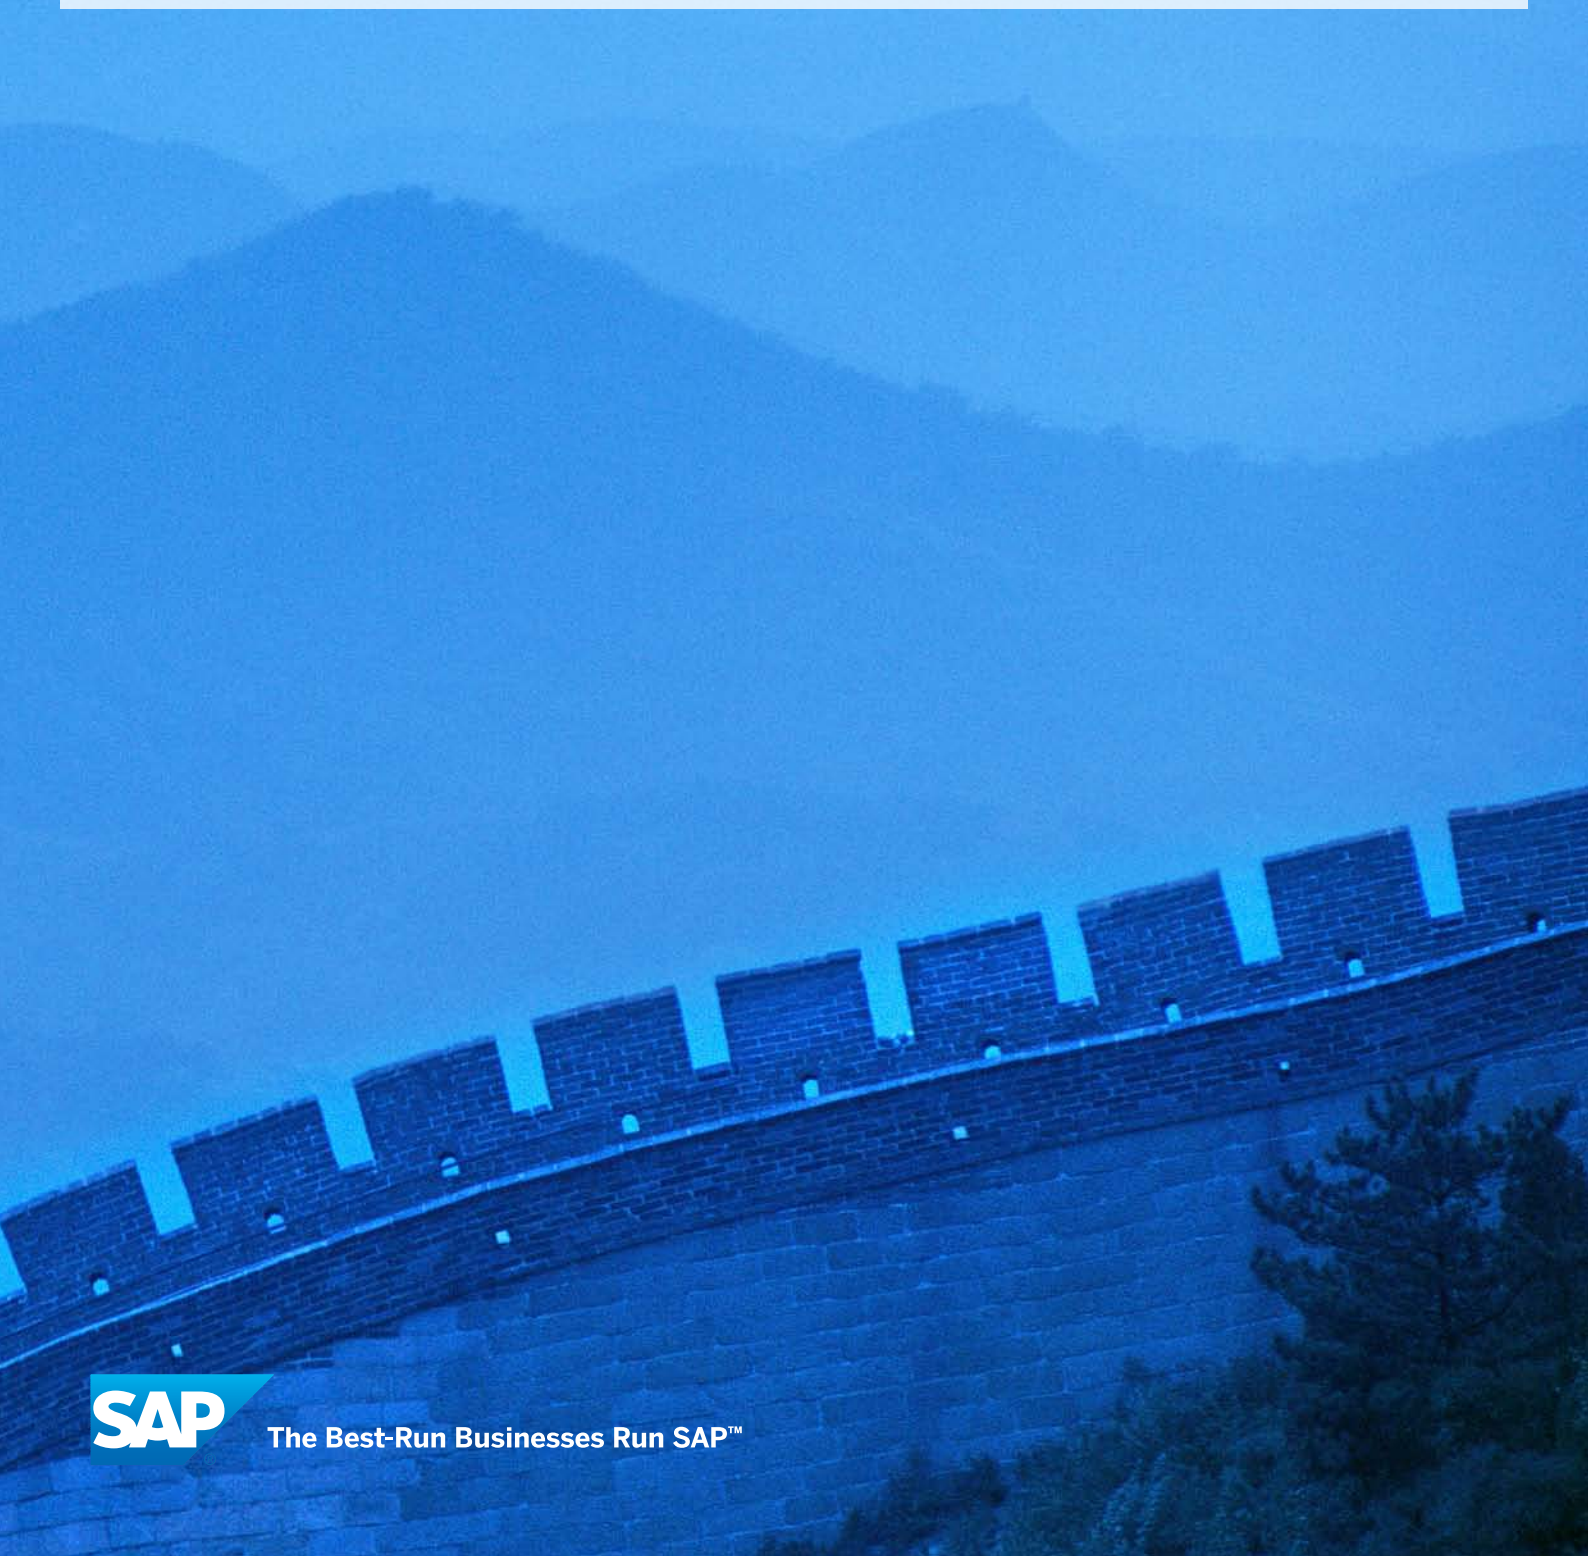

SAP

The Best-Run Businesses Run SAP™

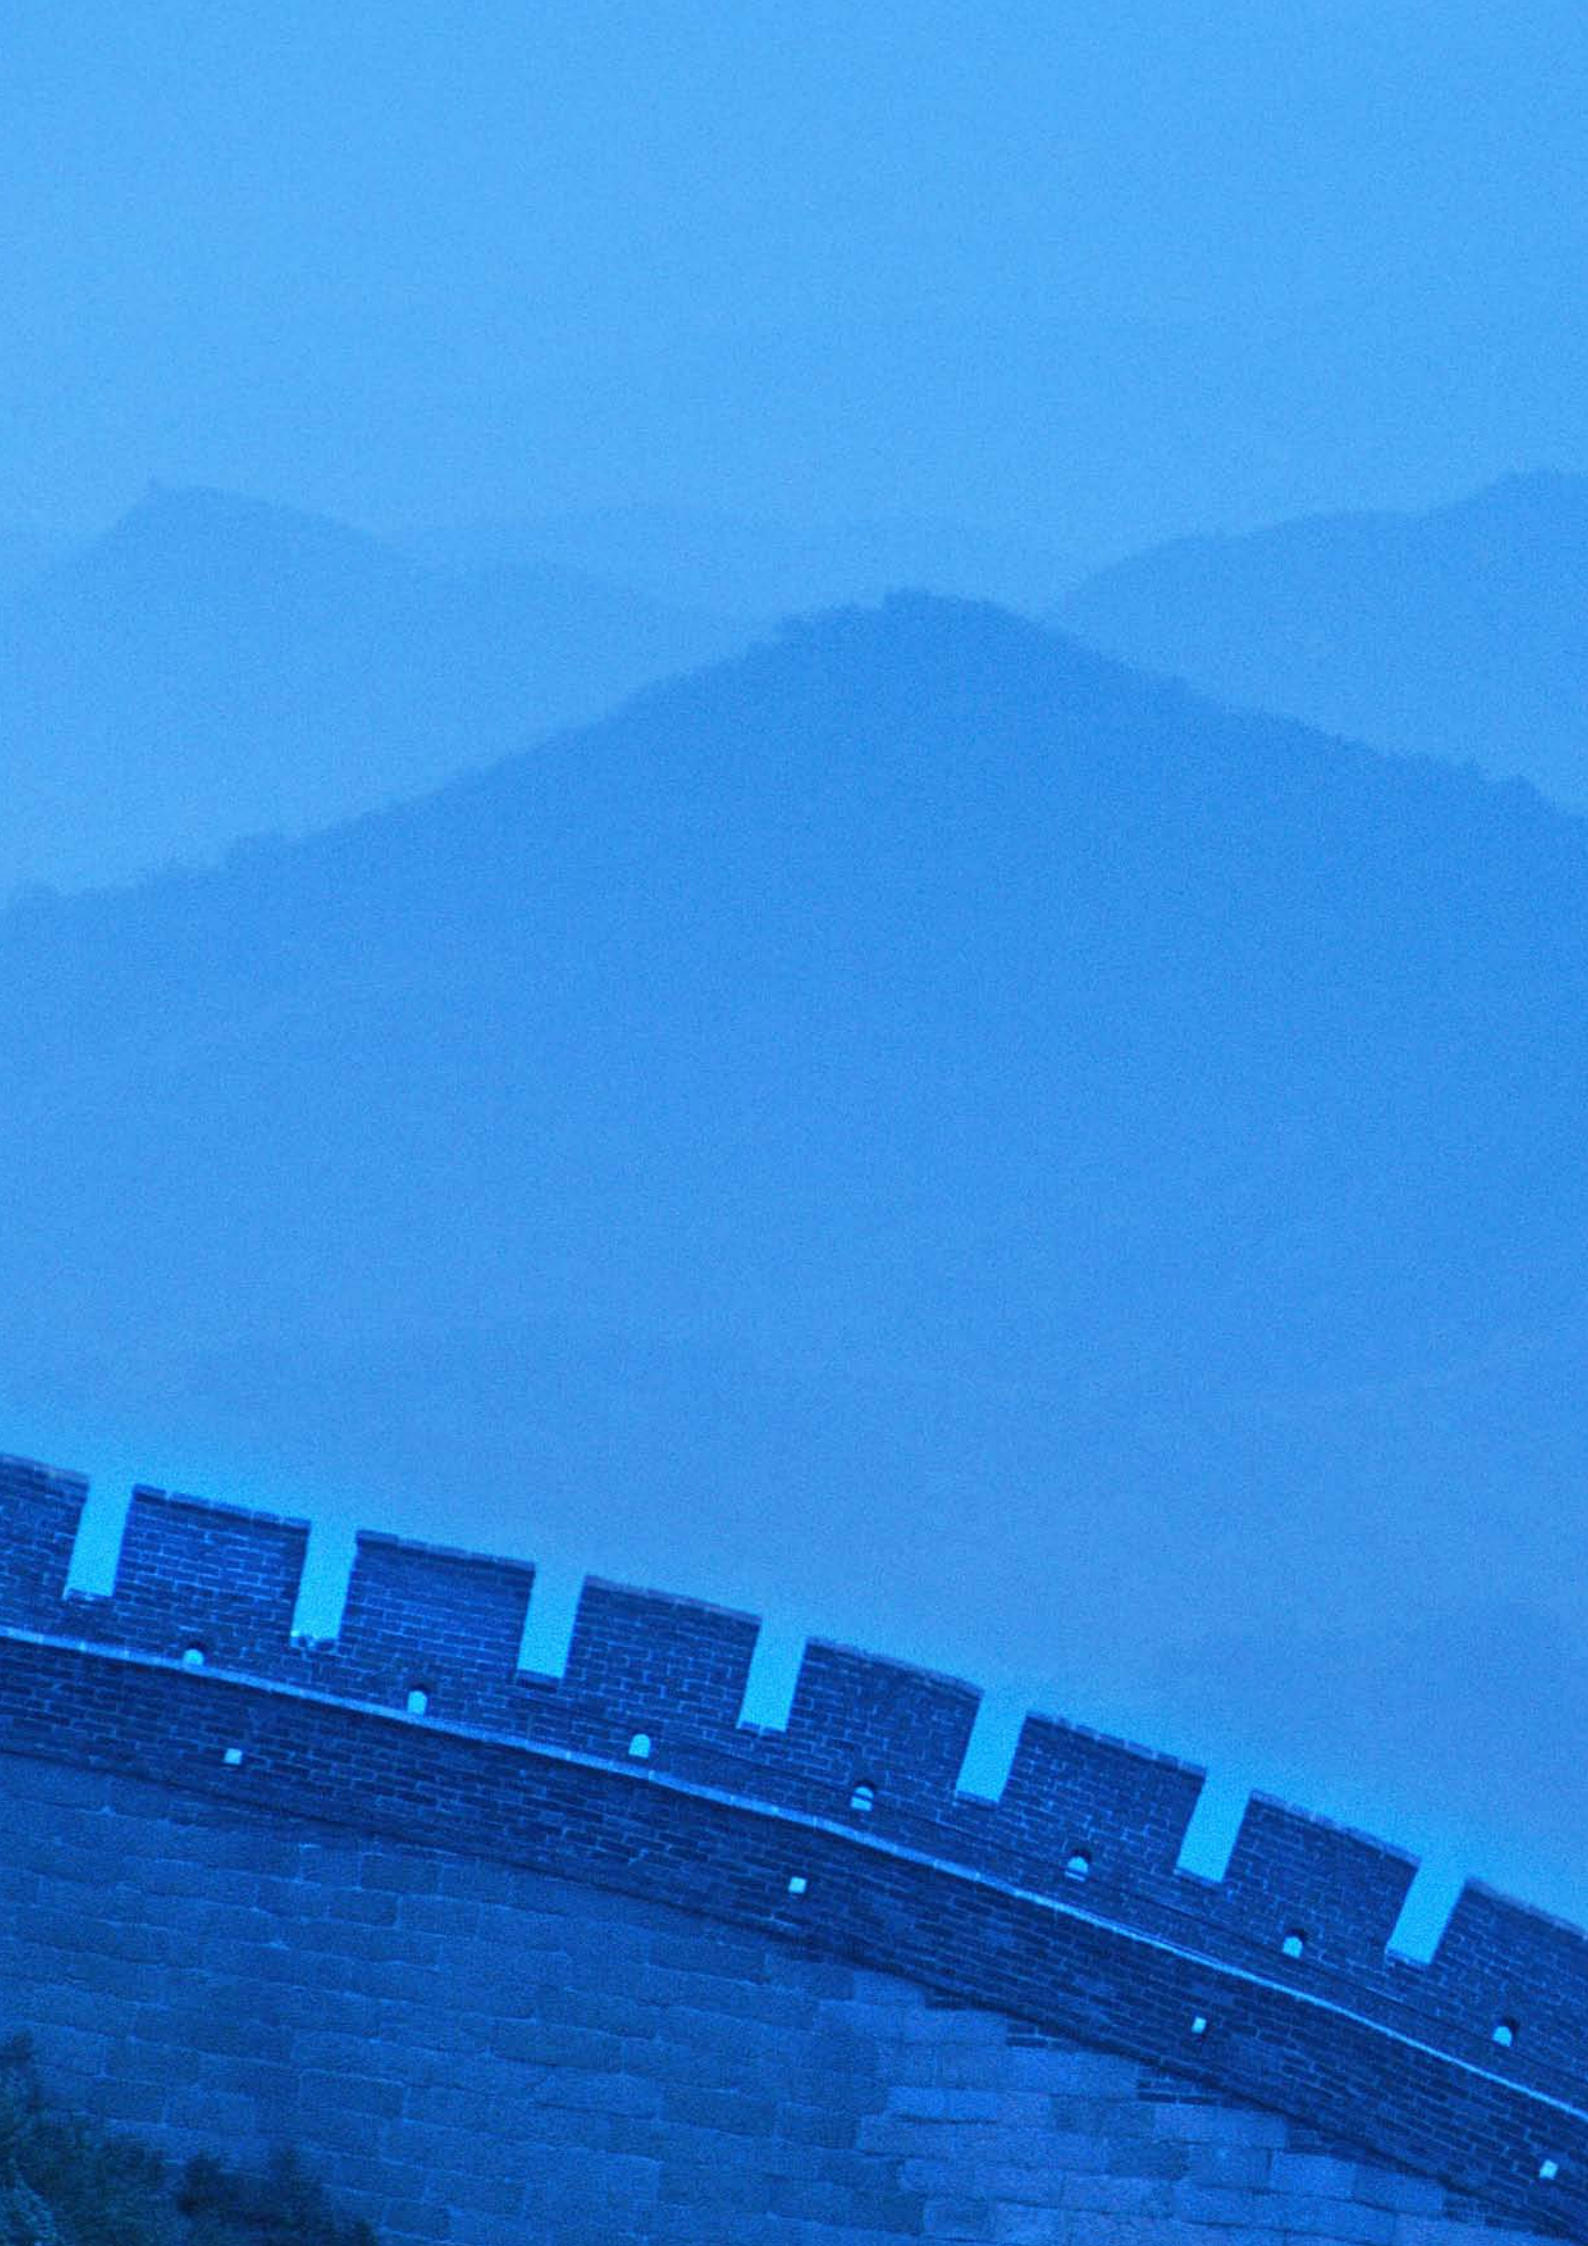

Table of Contents

About the Authors
Erich Schneider is a senior director in the Premier Customer Network (PCN) and has been implementing SAP® solutions for over 20 years. He was chief architect in the internal SAP program to enable data and business intelligence (BI) architecture using SAP in-memory computing technology across all SAP lines of business worldwide.

Raghav Jandhyala is a solutions expert for the SAP HANA® platform in the SAP Customer Solution Adoption (CSA) group, focusing on early customer adoption of innovations from SAP for strategic industries. Prior to joining CSA, Raghav was a solution manager for the SAP Banking group, where he supported go-to-market activities for SAP HANA in the banking industry.

Introduction

The announcement of the SAP HANA® platform has created a lot of buzz in the IT and business world. As new business demands challenge the status quo, the scale is larger, expectations are greater, and the stakes are higher. New-breed IT systems must be able to evaluate, analyze, predict, and recommend – and do so in real time. An in-memory approach is the only way to tackle a real-time-data future that includes new data types such as social media monitoring and Web-automated sensors and meter readings.

Today's business users need to react much more quickly to changing customer and market environments. They demand dynamic access to raw data in real time. SAP HANA empowers users with flexible, on-the-fly data modeling functionality by providing nonmaterialized views directly on detailed information. SAP HANA liberates users from the wait time for data model changes and database administration tasks, as well as from the latency required to load the redundant data storage required by traditional databases. The elimination of aggregates and relational table indices and the associated maintenance can greatly reduce the total cost of ownership.

Some use the term "in-memory" in the context of optimizing the I/O access with database management, centering on accessing data from the hard disk by pre-storing frequently accessed data in main memory. The term is also used for a traditional relational database running on in-memory technology. Some solutions offer columnar storage on traditional hard-disk technology, while other platforms offer the option of storing data on solid state disks (SSD). Although these disks have no moving parts and access data much more rapidly than hard disks, they are still slower than in-memory access.

Only SAP HANA takes full advantage of all-new hardware technologies by combining columnar data storage, massively parallel processing (MPP), and in-memory computing by using optimized software design.

Many SAP customers have already successfully deployed SAP HANA to drive innovations in IT and in business. In one scenario, SAP HANA supports data marts for ultrafast analytics with extreme data volumes, where users require more detailed information with access unrestricted by aggregations or caching of some of the data sourced from systems running SAP® and non-SAP software.

In a second scenario, SAP HANA replaces traditional relational databases underneath SAP applications. The SAP NetWeaver® Business Warehouse (SAP NetWeaver BW) component, a proven enterprise data warehouse solution, is the first application where SAP customers can migrate their existing database to an SAP HANA database.

In a third scenario, using a database connection (DB Connect), the SAP HANA database can be connected as a secondary database to the SAP ERP application, for example, and provide accelerated data processing for existing SAP applications. SAP CO-PA Accelerator software is the first solution delivered to SAP customers using this functionality.

This document provides a general understanding of SAP HANA as of support package 3 (SP03). It presents insights to the impact of SAP HANA on IT portfolio management and enterprise architecture; metadata and master data support; people, processes, and structures in business and IT; and data lifecycle management. It complements the SAP white paper *SAP HANA® for Next-Generation Business Applications and Real-Time Analytics* and the SAP thought leadership paper *SAP® In-Memory Computing Technology: Changing the Way Business Intelligence Is Managed*. For details on how to access these papers as well as other information about SAP HANA and in-memory computing, see the "Additional Information" section at the end of this document.

SAP HANA Platform

SAP HANA is a modern platform for real-time analytics and applications (see Figure 1). It enables organizations to analyze business operations based on a large volume and variety of detailed data in real time, literally at the speed of thought, from a human perspective. Initial deployments of SAP applications on SAP HANA have shown that business users can act on sub-seconds system response times, which opens the door to application possibilities that may not yet have been imagined.

The platform can be deployed as an appliance or delivered via a cloud. SAP in-memory computing is the core technology underlying the SAP HANA platform.

Figure 1: SAP HANA Platform Overview

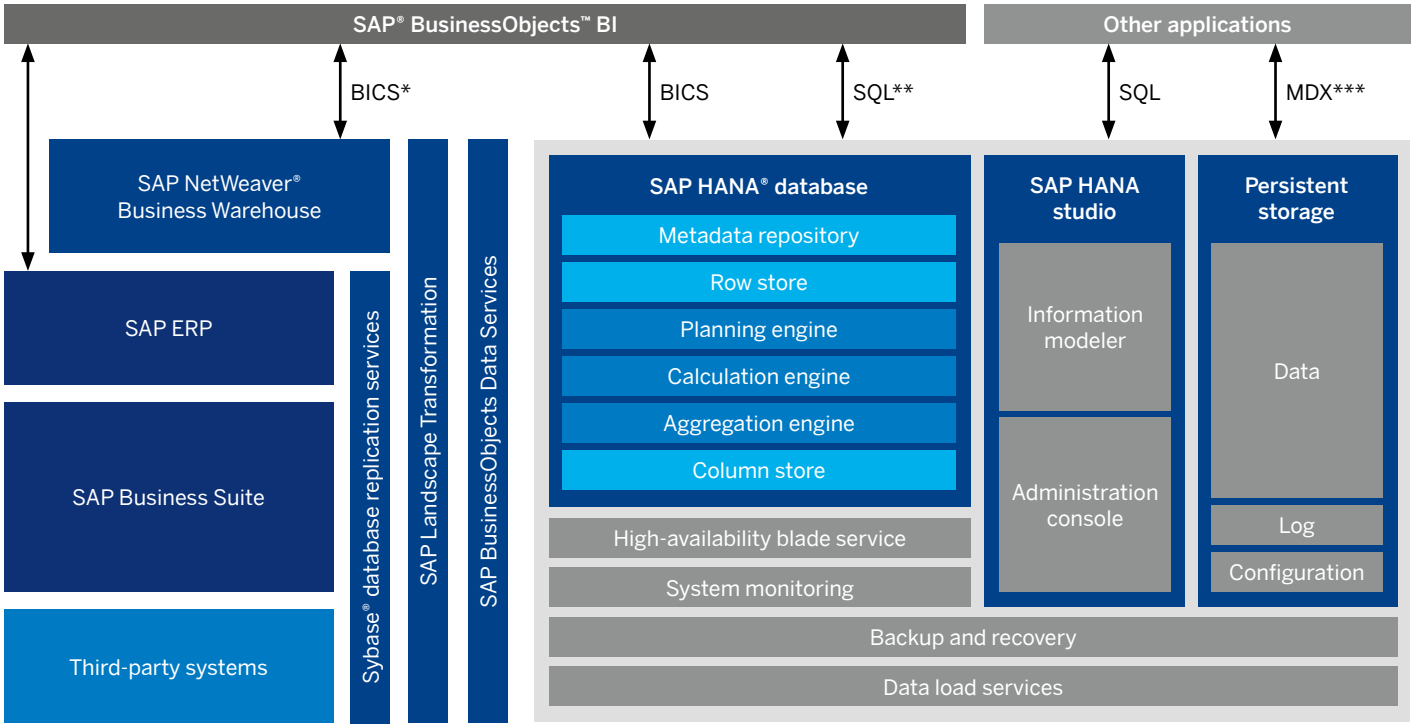

*BICS = business intelligence consumer services, **SQL = structured query language, ***MDX = multidimensional expression

SAP HANA Appliance Software

SAP HANA appliance software is a flexible, multipurpose, data-source-agnostic, in-memory appliance that combines SAP software components optimized on hardware provided and delivered by leading SAP technology partners, such as Cisco, Dell, IBM, HP, and Fujitsu, using the Intel Xeon processor. The appliance includes a number of integrated SAP software components, including the SAP HANA database, data and lifecycle management applications, support for multiple interfaces based on industry standards, and the SAP HANA studio, an easy-to-use data modeling and administration tool.

SAP Landscape Transformation software, with real-time replication services, and SAP BusinessObjects™ Data Services software integrate with SAP HANA but are not preinstalled. For the latest updates on SAP HANA appliance software, refer to SAP Note 1603671. Structured query language (SQL), Java Database Connectivity (JDBC), Open Database Connectivity (ODBC), and multidimensional expression (MDX) are industry-standard interfaces supported with SAP HANA.

Following the SAP appliance delivery model, the hardware vendor provides factory preinstallation for the hardware, operating system, and SAP software, and may add specific best practices for software configuration of SAP HANA to optimize performance. The hardware vendor finalizes installation with an on-site setup and configuration of the SAP HANA components, including deployment in a customer's data center, connectivity to networks, SAP Solution Manager application management solution, support for Secure Socket Layer (SSL), and connectivity for SAP support with an SAP software program known as SAP router.[1]

HARDWARE TECHNOLOGIES FOR SAP HANA®

SAP HANA® is fully adaptable to the dramatic advances of hardware storage technology, on premise and in the cloud. Multicore CPUs and 64-bit systems offer a new reality in scalability, where, for example, a blade server with 4 blades, each with 4 CPUs, and each CPU with 8 cores, provides a total of 128 cores for software processing. New technology available provides a multi-socket, 8-CPU system with 10 cores in 1 server, translating into 1 terabyte of memory.[2]

The customer starts by establishing replication and connectivity to data source systems and business intelligence (BI) clients. This would include the deployment of additional replication components on source systems and client components from the SAP BusinessObjects BI suite.

Although the term appliance suggests a black box ready to plug into a power outlet, in reality there is still individual work and cooperation required on a highly technical level. Superior project preparation and good coordination, detailed planning with project management resources, and quality training for the technical resources are essential prerequisites for a smooth deployment. Appliance technology reduces the implementation effort significantly, but does not reduce the effort to zero.

Combining OLAP and OTLP into one database, SAP HANA creates a unified view on data from transaction, analysis, decision, and planning systems. Real-time analytics and transactional applications, including planning, can run in mixed operations.

SAP HANA Database

Traditional database management systems are designed for optimizing performance on hardware with constrained main memory. Disk I/O was the main bottleneck. The focus was on optimizing disk access, for example, by minimizing the number of disk pages to be read into main memory during processing. The SAP HANA database is designed from the ground up around the idea that memory is available in abundance, considering that roughly 18 billion gigabytes or 18 exabytes are the theoretical limits of memory capacity for 64-bit systems, and that I/O access to the hard disk is not a constraint. Instead of optimizing I/O hard disk access, SAP HANA optimizes memory access between the CPU cache and main memory. SAP HANA is a massively parallel (distributed) data management system that runs fully in main memory, allowing for row- and column-based storage options, and supporting built-in multitenancy.

SAP HANA serves as a foundation to develop future in-memory analytic and transactional applications. The SAP HANA database can potentially provide performance improvements for existing SAP applications where, for example, SAP applications that use Open SQL can run on SAP HANA without changes.

New applications developed natively on SAP HANA and powered by SAP HANA can improve performance of business process and analytical scenarios. Application development techniques optimized for parallel in-memory processing can take advantage of new enterprise data management and application development logic to fully exploit advances in hardware technologies.

The table summarizes the benefits offered by specific features of the SAP HANA database.

The SAP HANA database manages data in a multicore architecture for data distribution across all cores to maximize RAM locality using scale-out (horizontally) and scale-up (vertically) functionality.[3]

In the scale-out scenario, the SAP HANA database scales beyond a single server by allowing multiple servers in one cluster. Large tables can be distributed across multiple servers using round-robin, hash, or range partitioning, either alone or in combination. SAP HANA has the functionality to execute queries and maintain distributed transaction safety across multiple servers.

Specific server configurations for SAP HANA deployments are the responsibility of certified SAP technology partners. These partners can balance better performance per CPU at lower cost, enabling customers to take advantage of larger memory-address spaces, lower data-center operating costs, and simpler management.

One of the major contentions and the reason for slow performance in traditional DBMS is locking data when data updates are being performed. SAP HANA avoids this issue and enables high levels of parallelization using insert-only data records. Instead of creating new records in a database table, deltas are inserted as net-new entries in existing records stored in columns.

Using columnar data stores, SAP HANA can achieve major compression rates unheard of in traditional databases. On one example, the analysis of SAP customers' systems showed that only 10% of attributes in a single financial database table was

SAP HANA® Database Feature	Benefit
Multicore CPU Large memory footprint	Greater computation power Faster than disk
Row and column store	Faster aggregation with column store
Compression	Highly dependent on actual data used
Partitioning	Analysis of large data sets Complex computations
No aggregate tables Nonmaterialized views	Flexible modeling No data duplication
Insert only on delta	Fast data loads

used in an SQL statement, shrinking the actual size of data volume to be accessed from 35 GB in traditional relational database management system (RDBMS) storage to 800 MB in a column-store design, just over 2% of the volume in the traditional storage. As this example shows, much higher compression rates will be accomplished with high-sparsity data than with dense data.[4]

SAP HANA DATABASE ADMINISTRATION

The administration console of the SAP HANA studio provides an all-in-one support environment for system monitoring, backup and recovery, and user provisioning.

In case of scenarios like data center failures due to fire, power outages, earthquake, and so on, or hardware failures, such as the failure of a node, SAP HANA supports the hot-standby concept using synchronous mirroring with a redundant data center concept, including redundant SAP HANA databases (see Figure 2). This is in addition to the cold-standby concept using a standby system within one SAP HANA landscape, where the failover is triggered automatically (see Figure 3).

SAP HANA is an ACID-compliant database supporting atomicity, consistency, isolation, and durability (ACID) of transactions. In addition to recovery for online analytical processing (OLAP), SAP HANA also provides transactional recovery for online transactional processing (OLTP) through the administrative console in the SAP HANA studio.

Currently supported processes are:
- Recovery to last data backup
- Recovery to last and older (previous) data backup
- Recovery to last state before crash
- Point-in-time recovery

User provisioning is supported with authentication, role-based security, and analysis authorization using analytic privileges, which enable security for analytical objects based on a set of attribute values.

The administration console in the SAP HANA studio enables the version control mechanism for models of SAP HANA and SAP BusinessObjects Data Services. Supported is the export and import function of XML files, which can be used to back up and restore versions of models. SAP HANA can be run in a single production landscape, especially if the initial use-case scenario is not business critical and the data load performance for the initial load is acceptable for reloading the data. However, we recommend aligning the SAP Landscape Transformation and SAP BusinessObjects Data Services environment with the existing source system landscapes. In an enterprise-grade business support mode, SAP HANA needs to run in standard SAP development, quality assurance and staging, and production environments.

Figure 2: Example Configuration for Disaster-Tolerant Redundant Hardware

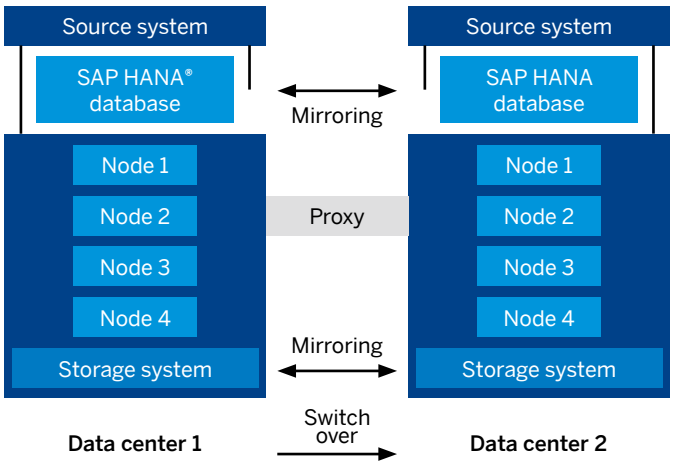

Figure 3: Distributed System: Multiple Nodes Using Shared Storage for Workload Distribution

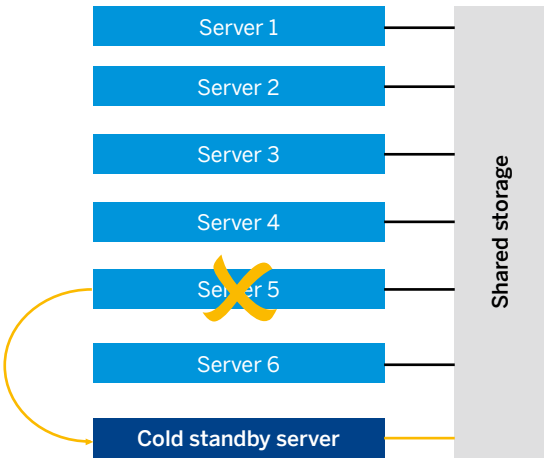

Sizing for the SAP HANA Database

Sizing for SAP HANA is continuously optimized based on customer implementations and can vary. The sizing consists of individual calculations for memory sizing, each for static data and for data objects created during runtime (data load and query execution); for disk sizing to allow for persistency of data logs and data models; and for CPU sizing. The number of users determines the size of the CPU, considering query complexity and user behavior, such as click rates. SAP provides Quick Sizer tool support for sizing SAP HANA.

In general, the size of the data-source footprint and its compression rate determine the memory sizing. SAP Note 1609322 in the SAP Notes tool describes in detail how to determine the data-source footprint. The memory sizing is based on an average compression rate of 5 for the source tables without indices and aggregations for the static as well as for the dynamic memory. The compression rate is highly dependent on the sparsity of the source data and is also relative to the metadata volume, so the compression rate based on initial data loads will change with the increase of data volume loaded into SAP HANA.

It is important to consider the access frequency of data for in-memory-computing sizing. In the case of SAP NetWeaver Business Warehouse, only "hot" data, such as data for the current and previous year, would be loaded in the memory of the SAP HANA database. "Cold" data (history of the last 10 years, for example) would be loaded into near-line storage solutions.

The initial sizing needs to be continuously validated, as sudden increases in data volume can cause swapping from the memory to the persistent data storage, which, of course, causes dramatic performance loss. Therefore, tight integration between the business and IT is required to monitor the data growths for the selected business process to be supported by SAP HANA. Consult the SAP Services organization and your technology partner for the sizing of your SAP HANA implementation.

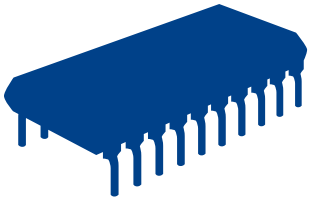

Only SAP HANA takes full advantage of all-new hardware technologies by combining columnar data storage, massively parallel processing (MPP), and in-memory computing by using optimized software design.

Data Load Architecture Scenarios for SAP HANA

SAP HANA offers both real-time-replication and periodic-load options to move data from source systems to the SAP HANA database. Replication-based data provisioning in SAP Landscape Transformation provides near-real-time synchronization of data sets between source systems and SAP HANA. The data is pushed from sources to SAP HANA as and when it is available. On the other hand, extraction-based data provisioning, like SAP BusinessObjects Data Services, loads snapshots of data periodically as a batch and is triggered from the target system. Depending on the deployment landscape of the customer environment, the following data load architectures can be used.

NEAR-REAL-TIME REPLICATION WITH SAP LANDSCAPE TRANSFORMATION

The replicator in SAP Landscape Transformation software provides near-real-time and scheduled data replication from SAP and non-SAP sources to SAP HANA. It is based on proven system landscape optimization technology used for many years for near-zero-downtime upgrade and migration projects.

Trigger-based data replication using SAP Landscape Transformation (see Figure 4) is based on capturing database changes at a high level of abstraction in the source enterprise resource planning (ERP) system. It benefits from being database and operating system agnostic and can parallelize database changes on multiple tables or by segmenting large table changes. The key building blocks are database triggers and table-based delta logging.

SAP Landscape Transformation provides high availability and security by utilizing reliable replication features such as auto-reconnect functions and buffering of database changes in source systems to aid in case of power outages. It uses standard SAP optimization techniques for providing high-performance replication of data. It also provides flexibility to merge data from different source systems along with reliable central monitoring of the entire replication process.

SAP Landscape Transformation software can be installed on an existing SAP source system or on an additional lightweight SAP system alongside the source system.

Figure 4: SAP Landscape Replication: Architecture and Key Building Blocks

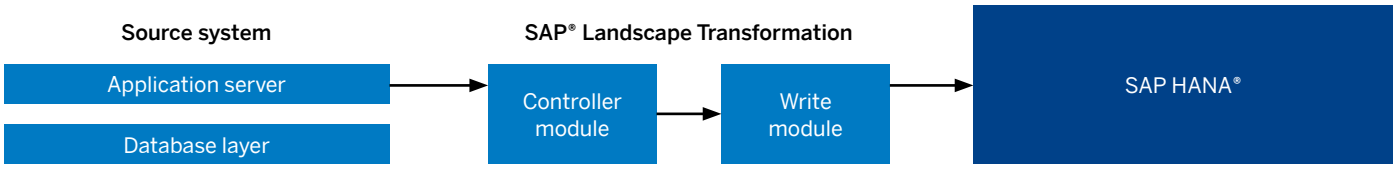

REAL-TIME REPLICATION WITH DIRECT WRITE

SAP HANA supports real-time replication with direct write using a database shared library (DBSL) connection. With DBSL, the SAP HANA database can be connected as a secondary database, for example, to a system running SAP ERP, and provide accelerated data processing for existing SAP applications. Applications can use DBSL on the application server layer to simultaneously write to traditional databases and the SAP HANA database.

EXTRACTION/PERIODIC LOAD

Data loading based on extraction, transformation, and load (ETL) uses SAP BusinessObjects Data Services to load relevant business data from any source system (SAP and non-SAP) to the SAP HANA database (see Figure 5). SAP BusinessObjects Data Services is a proven ETL tool and supports broad connectivity to databases, applications, legacy file formats, and unstructured data. It provides a modeling environment to model data flows from one or more source systems along with transformations and data cleansing. The data load jobs are scheduled using the job server to load to the SAP HANA database.

RECOMMENDED SCENARIOS TO LOAD DATA INTO THE SAP HANA® DATABASE

- Use SAP® Landscape Transformation software for real-time direct table-level integration for SAP and non-SAP data sources.
- Use SAP BusinessObjects™ Data Services software for all periodic-load scenarios from SAP and non-SAP data sources, including loading data from the SAP NetWeaver® Business Warehouse component.

Figure 5: Extraction, Transformation, and Load (ETL) Using SAP BusinessObjects Data Services

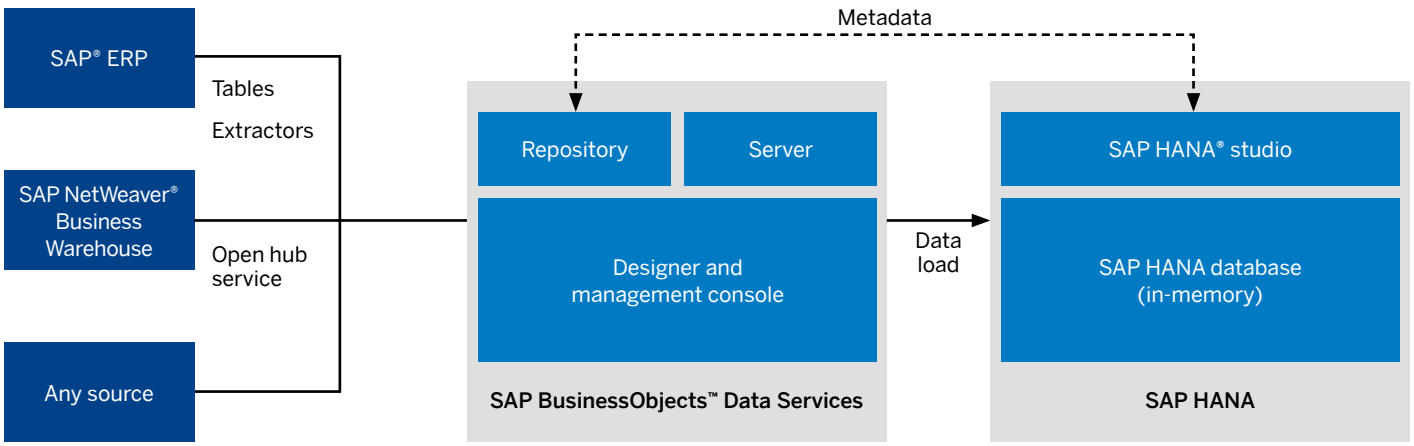

Deployment Scenarios for SAP HANA

The SAP HANA software appliance can be plugged into your landscape to accelerate both operational reporting and data warehousing without disrupting the existing IT landscape. As a side-by-side approach, SAP HANA addresses today's challenges of increasing data volumes, redundant storage, calculation speed, information latency, and additional data aggregation in data marts.

SAP HANA supports three types of data marts:
- Operational data marts for real-time operational reporting
- Agile line-of-business data marts where users can add their own data and data models
- Architected data marts on the corporate enterprise data warehouse layer supporting data from multiple data sources

The sections below describe best practices of deploying SAP HANA for operational, agile, and architected data marts to enable new functionality for your current landscape.

OPERATIONAL DATA MART

Operational data from source systems can be replicated to SAP HANA for flexible operational data marts and query performance (see Figure 6). With either real-time replication or periodic load, data from operational systems is loaded into the SAP HANA database. After the database structures and data are loaded into SAP HANA, they can be modeled with on-the fly aggregation for operational reporting needs, without any data duplication.

AGILE DATA MART

With SAP BusinessObjects Data Services, selected analytical models can be copied from data warehouses (SAP or non-SAP) along with ad hoc or line-of-business data sources to the SAP HANA database for agile data marts (see Figure 7). Any complex transformations from multiple systems can be handled in SAP BusinessObjects Data Services and loaded to SAP HANA for highly flexible modeling, extensions of analytical models, and reporting.

ARCHITECTED DATA MART

SAP HANA supports architected data marts by combining architected data sources from data warehousing systems along with operational sources (see Figure 8).

Figure 6: SAP HANA for Operational Data Marts

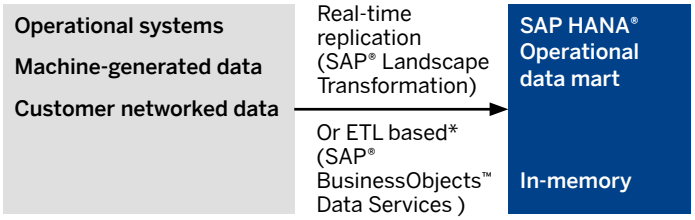

*ETL = extraction, transformation, and load

Figure 7: SAP HANA for Agile Data Marts

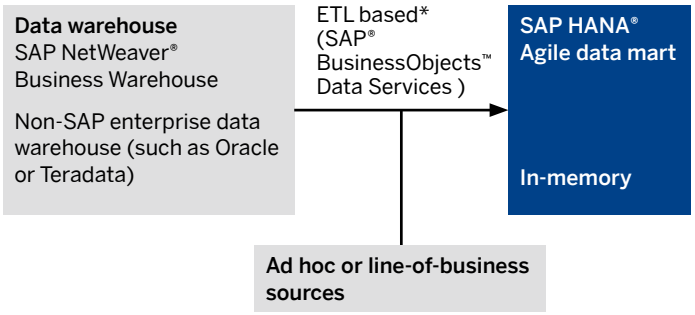

*ETL = extraction, transformation, and load

Figure 8: SAP HANA for Architected Data Marts

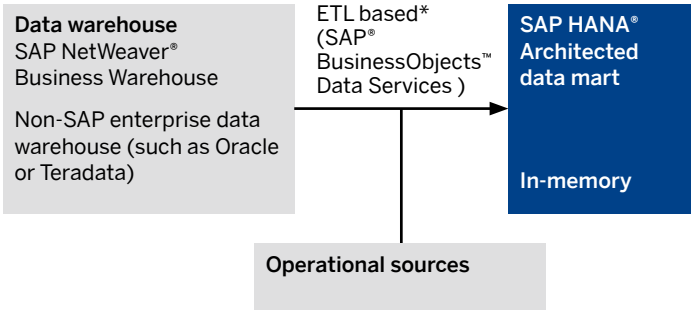

*ETL = extraction, transformation, and load

SAP HANA DATABASE FOR SAP NetWeaver BUSINESS WAREHOUSE

SAP NetWeaver BW can be accelerated by using SAP HANA as its underlying database. In this scenario, SAP NetWeaver BW functions, such as those of SAP NetWeaver BW Accelerator software and new SAP HANA–optimized information providers (objects within a database that act as data providers within a query definition), are performed within SAP HANA to benefit from its in-memory and calculation engine functionality. This scenario also includes the SAP BusinessObjects Planning and Consolidation application, version for SAP NetWeaver. Available as of release 7.30 of SAP NetWeaver BW with support package 5, this feature eliminates the need of a separate instance of SAP NetWeaver BW Accelerator after migration from a traditional RDBMS to the SAP HANA database.

SAP HANA DATABASE FOR SAP BUSINESS SUITE ACCELERATORS

SAP HANA further supports a new breed of applications called "suite accelerators," which accelerate existing SAP Business Suite applications through a side-by-side approach where SAP HANA is deployed with SAP Business Suite applications. These accelerators leverage SAP HANA as a secondary database, keeping the existing application UI. Database tables that are needed for the accelerator are replicated into the SAP HANA database. Data access of the application is executed against the replicated tables in the SAP HANA database instead of the primary database below the application server. SAP CO-PA Accelerator software is one of the first solutions delivered to SAP customers using this functionality.

SAP HANA DATABASE FOR IN-MEMORY BUSINESS APPLICATIONS

SAP has started a new wave of business applications that are designed, developed, and optimized to take full advantage of SAP in-memory computing technology. New applications that are powered by SAP HANA include SAP Smart Meter Analytics software, the SAP Dynamic Cash Management application, and the SAP Sales and Operations Planning application.

SAP HANA DATABASE FOR CORE SAP BUSINESS SUITE APPLICATIONS

Following the successful launch of SAP NetWeaver BW on SAP HANA, SAP will offer SAP ERP 6.0 and core SAP Business Suite 7.0 applications on SAP HANA as well. SAP ERP will be the second SAP product for which database migration from a traditional RDBMS to the SAP HANA database will be possible.

With this development, SAP demonstrates its commitment to timeless software by defining architectures that assume continuous change and by building solutions that can consume innovation nondisruptively, even as fundamental an innovation as SAP HANA. With SAP Business Suite applications running on the SAP HANA platform, SAP protects customers' existing investments while delivering additional value.

SAP HANA APPLICATION CLOUD

The SAP HANA application cloud is an on-demand application environment based on SAP in-memory computing. Applications delivered with the SAP HANA application cloud allow customers to immediately leverage the benefits of SAP HANA through easy-to-use and scalable cloud-based offerings. The first regional data center, located in the United States, supports SAP BusinessObjects OnDemand offerings as well as Recalls Plus. This app, available to consumers in the United States from the Apple App Store, enables searching recall history by brand or product category and sharing recalls with others. Recalls Plus is the first mobile, cloud-native app designed by SAP for mass consumer adoption.

Data Analysis in SAP HANA

SAP HANA supports open interfaces, such as for Microsoft Excel. Nevertheless the "face" of SAP HANA for all analytical operations is the SAP BusinessObjects BI suite. SAP BusinessObjects BI suite is optimized to access data on SAP HANA and provides instant access and split-second response time to queries on massive volumes of data. SAP BusinessObjects BI suite accesses the nonmaterialized data views created in the information modeler of the SAP HANA studio.

Figure 9: Example of a Simplified SAP HANA Landscape for Analytics

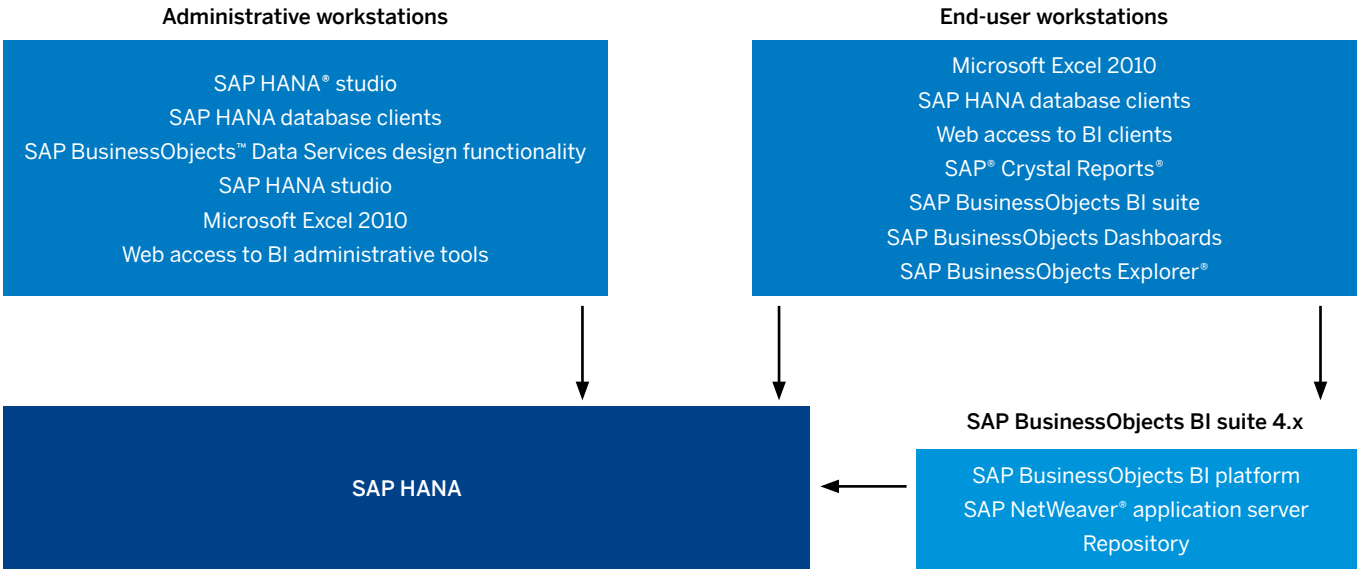

The table below describes the SAP BusinessObjects reporting tools, as shown in Figure 9, and key benefits with SAP HANA.

SAP HANA is designed to support the search and explore paradigm with extreme query performance. For example, by running SAP BusinessObjects Explorer® software on an iPad or iPhone (software available for download from the App Store), consumers can run fully enabled GPS analytics to find a store near them with the product they are searching for. To further support geolocation functionality, SAP has partnered with Google Maps to provide users of SAP applications with unrestricted Web integration.

Purpose	SAP® BusinessObjects™ offerings	Use Case	Benefits with SAP HANA®
Formatted reporting	SAP® Crystal Reports® software	Access and transform corporate data into formatted enterprise reports for greater insight	Native support of SAP HANA® software relational semantic layers accessing information models in SAP HANA
Dashboards and visualization	SAP BusinessObjects™ Dashboards software	Visualize key performance indicators for management decision making	SAP BusinessObjects BW Universe Builder software for passing queries from the reporting tools to SAP HANA for execution
Interactive or ad hoc analysis	SAP BusinessObjects Web Intelligence® software SAP BusinessObjects Web Intelligence Interactive Viewing software	Use self-service access across heterogeneous data sources for ad hoc reporting and interactive analysis	Faster on-the-fly query refresh that potentially eliminates the need for scheduling in SAP BusinessObjects Web Intelligence
Online analytical processing (OLAP) analysis	SAP BusinessObjects Analysis, edition for OLAP	Uncover trends from historical data and make possible better forecasts	Multidimensional analysis using the analytical models created in SAP HANA
Searching and data exploration	SAP BusinessObjects Explorer® software	Use self-services to get answers to business questions on large volumes of data	SAP BusinessObjects Explorer using the in-memory data of SAP HANA for on-the-fly data exploration instead of disk-based indexes on SAP BusinessObjects BI server

Information Modeling in SAP HANA

Business and IT users can either create on-the-fly data views or build reusable nonmaterialized data views via an intuitive user interface. This UI utilizes SQL script and stored procedures to perform business logic on the data models, complementing the functionality of the SAP BusinessObjects BW Universe Builder software.

The information models created in SAP HANA can also be consumed by SAP BusinessObjects clients using the SAP BusinessObjects BW Universe Builder (semantic layer) built on top of SAP HANA views. SAP BusinessObjects BW Universe Builder software delegates query requests to SAP HANA where the query is processed to benefit from in-memory and parallelization features of SAP HANA. Additionally, the semantic layers can extend and provide data processing and complex calculations built on the underlying models in SAP HANA.

Information models in SAP HANA are a combination of attributes, dimensions, and measures. SAP HANA provides three types of modeling views:

- Attribute views are built on dimensions or subject areas used for business analysis.
- Analytical views are multidimensional views or OLAP cubes that enable the analysis of values from single fact tables related to the dimensions in the attribute views.
- Calculation views are used to create custom data sets to address complex business requirements using database tables in addition to attribute and analytical views.

AGILE DATA MODELING WITH THE INFORMATION COMPOSER

SAP HANA has an information composer tool that is targeted to business audiences looking for agile data modeling. This tool allows business analysts who are not familiar with data modeling to create new analyses that blend private data with centrally published data by uploading data into SAP HANA using a spreadsheet or by cutting and pasting from a clipboard. Because many of the more technical steps involved in typical data modeling processes are automated through intuitive, wizard-driven interactions, business users can rapidly perform new analyses to answer business questions in real time without the need to engage IT support.

Figure 10: Information Modeler in SAP HANA® Studio for Creating Real-World Business Models

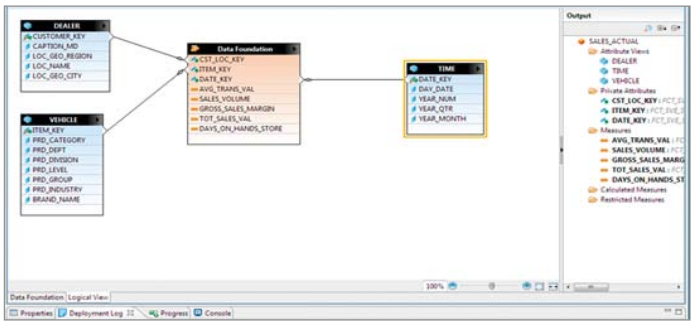

DYNAMIC DATA MODELING WITH THE INFORMATION MODELER

While highly experienced SQL analysts would access data directly in the SAP HANA database, we currently recommend using the information modeler (see Figure 10) in the SAP HANA studio as the only way to create nonmaterialized views. This can help ensure viability during release upgrades.

In traditional databases, users experience bottlenecks when new or modified business requirements result in changes to the existing data model. This subsequently required deleting and reloading data into materialized views. SAP HANA provides dynamic data modeling on the lowest granular level loaded into SAP HANA. In SAP HANA, raw data is exclusively available in memory for analytical purposes. Data is not preloaded in cache, in physical aggregate tables or index tables in relational databases, or in any other redundant data storage for the sake of performance management.

The nonmaterialized views enable flexible data modeling on the fly, liberating end users from the traditional data-modeling bottleneck. These views enable the business to react more quickly to changing customer and market environments. For example, change log information can be viewed in SAP HANA in a calculation view (nonmaterialized view) instead of a materialized view, such as a separate physical table.

Outlook

Combining online analytical processing and online transactional processing into a single database, SAP HANA creates a unified view on enterprise data from transaction, analysis, decision, and planning systems. Real-time analytics applications and real-time transactional applications, including planning, can run in mixed operations.

SAP HANA can drive system landscape simplification by eliminating the need for purpose-built databases such as operational data stores, and thus reduce overall system proliferation.

SAP HANA will become a staple in the best-run companies in the world, on premise as well as supporting SAP applications in the cloud for large enterprises and small businesses. With the SAP HANA platform, SAP enables co-innovation for SAP customers and partners with the SAP Services and SAP Custom Development organizations, as well as with the SAP ecosystem and partners.

FIND OUT MORE

For any company interested in IT innovation, the time is now to take advantage of an uncompromising, real-time, in-memory database. To learn more, contact your SAP representative or visit us on the Web at www.sap.com/hana.

ADDITIONAL INFORMATION

Global SAP HANA®
www.experiencesaphana.com

SAP® Help Portal
help.sap.com/hana/

SAP HANA 1. 0 Overall Install Guide
https://service.sap.com/~sapidb/011000358700000604562011

SAP HANA Database – Administration Guide
http://help.sap.com/hana/hana1_imdb_studio_admin_en.pdf

SAP HANA 1.0 Modeling Guide
http://help.sap.com/hana/hana1_model_en.pdf

SAP HANA Security Guide
help.sap.com/hana/hana1_sec_en.pdf

SAP Note for SAP HANA
https://service.sap.com/sap/support/notes/1514967

SAP Note for SAP HANA operational concept
https://service.sap.com/sap/support/notes/1599888

M. Bernard, "SAP High-Performance Analytic Appliance 1.0 [SAP HANA] – A First Look at the System Architecture" (February 23, 2011, Webinar presentation at http://www.sdn.sap.com/irj/scn/events?rid=/library/uuid/70d16119-ad21-2e10-de8b-eaaedf86b9cd)

Analyzing Business as It Happens: SAP In-Memory Appliance Software (SAP HANA®) runs on the Intel® Xeon® processor to generate superior, real-time business intelligence (Intel|SAP, April 2011, available at http://www.intel.com/en_US/Assets/PDF/whitepaper/mc_sap_wp.pdf)

SAP HANA® for Next-Generation Business Applications and Real-Time Analytics (SAP AG, February 2012, available at http://download.sap.com/download.epd?context=E67AFF9FD4CFC6693FC443EB965A7B4FE599BA88ED6C9F622486D0E5BEB350EB7DAD751393D16A2D916A3C9EA834D1CB2A8138467760590E)

E. Schneider and R. Jandhyala, *SAP® In-Memory Computing Technology: Changing the Way Business Intelligence Is Managed* (SAP AG, March 2011, available at http://fm.sap.com/data/UPLOAD/files/SAP_In-Memory_Computing_Technology_.pdf)

H. Plattner, *A Common Database Approach for OLTP and OLAP Using an In-Memory Column Database* (SIGMOD'09, June 29–July 2, 2009, available at http://www.sigmod09.org/images/sigmod1ktp-plattner.pdf)

H. Plattner and A. Zeier, *In-Memory Data Management – An Inflection Point for Enterprise Applications* (Springer-Verlag, April 2011)

Customers' experiences with SAP HANA at
www.sap.com/hana/reviews/index.epx

DISCLAIMER

This presentation outlines our general product direction and should not be used as a basis for making a purchase decision. This presentation is not subject to your license agreement or any other agreement with SAP. SAP has no obligation to pursue any course of business outlined in this presentation or to develop or release any functionality mentioned. This presentation and SAP's strategy and possible future developments are subject to change and may be changed by SAP at any time for any reason without notice. This document is provided without a warranty of any kind, either express or implied, including but not limited to, the implied warranties of merchantability, fitness for a particular purpose, or noninfringement. SAP assumes no responsibility for errors or omissions in this document, except if such damages were caused by SAP intentionally or through gross negligence.

FOOTNOTES

1. SAP router is an SAP program that acts as an intermediate station (proxy) in a network connection between SAP systems or between SAP systems and external networks.
2. Source: *Analyzing Business as it Happens: SAP In-Memory Appliance Software (SAP HANA®) runs on Intel® Xeon® processor to generate superior, real-time business intelligence* (Intel|SAP, April 2011).
3. The SAP certification program for technology vendors validates functionality by hardware product; individual hardware functionality can differ.
4. H. Plattner, *A Common Database Approach for OLTP and OLAP Using an In-Memory Column Database* (SIGMOD'09, June 29–July 2, 2009).

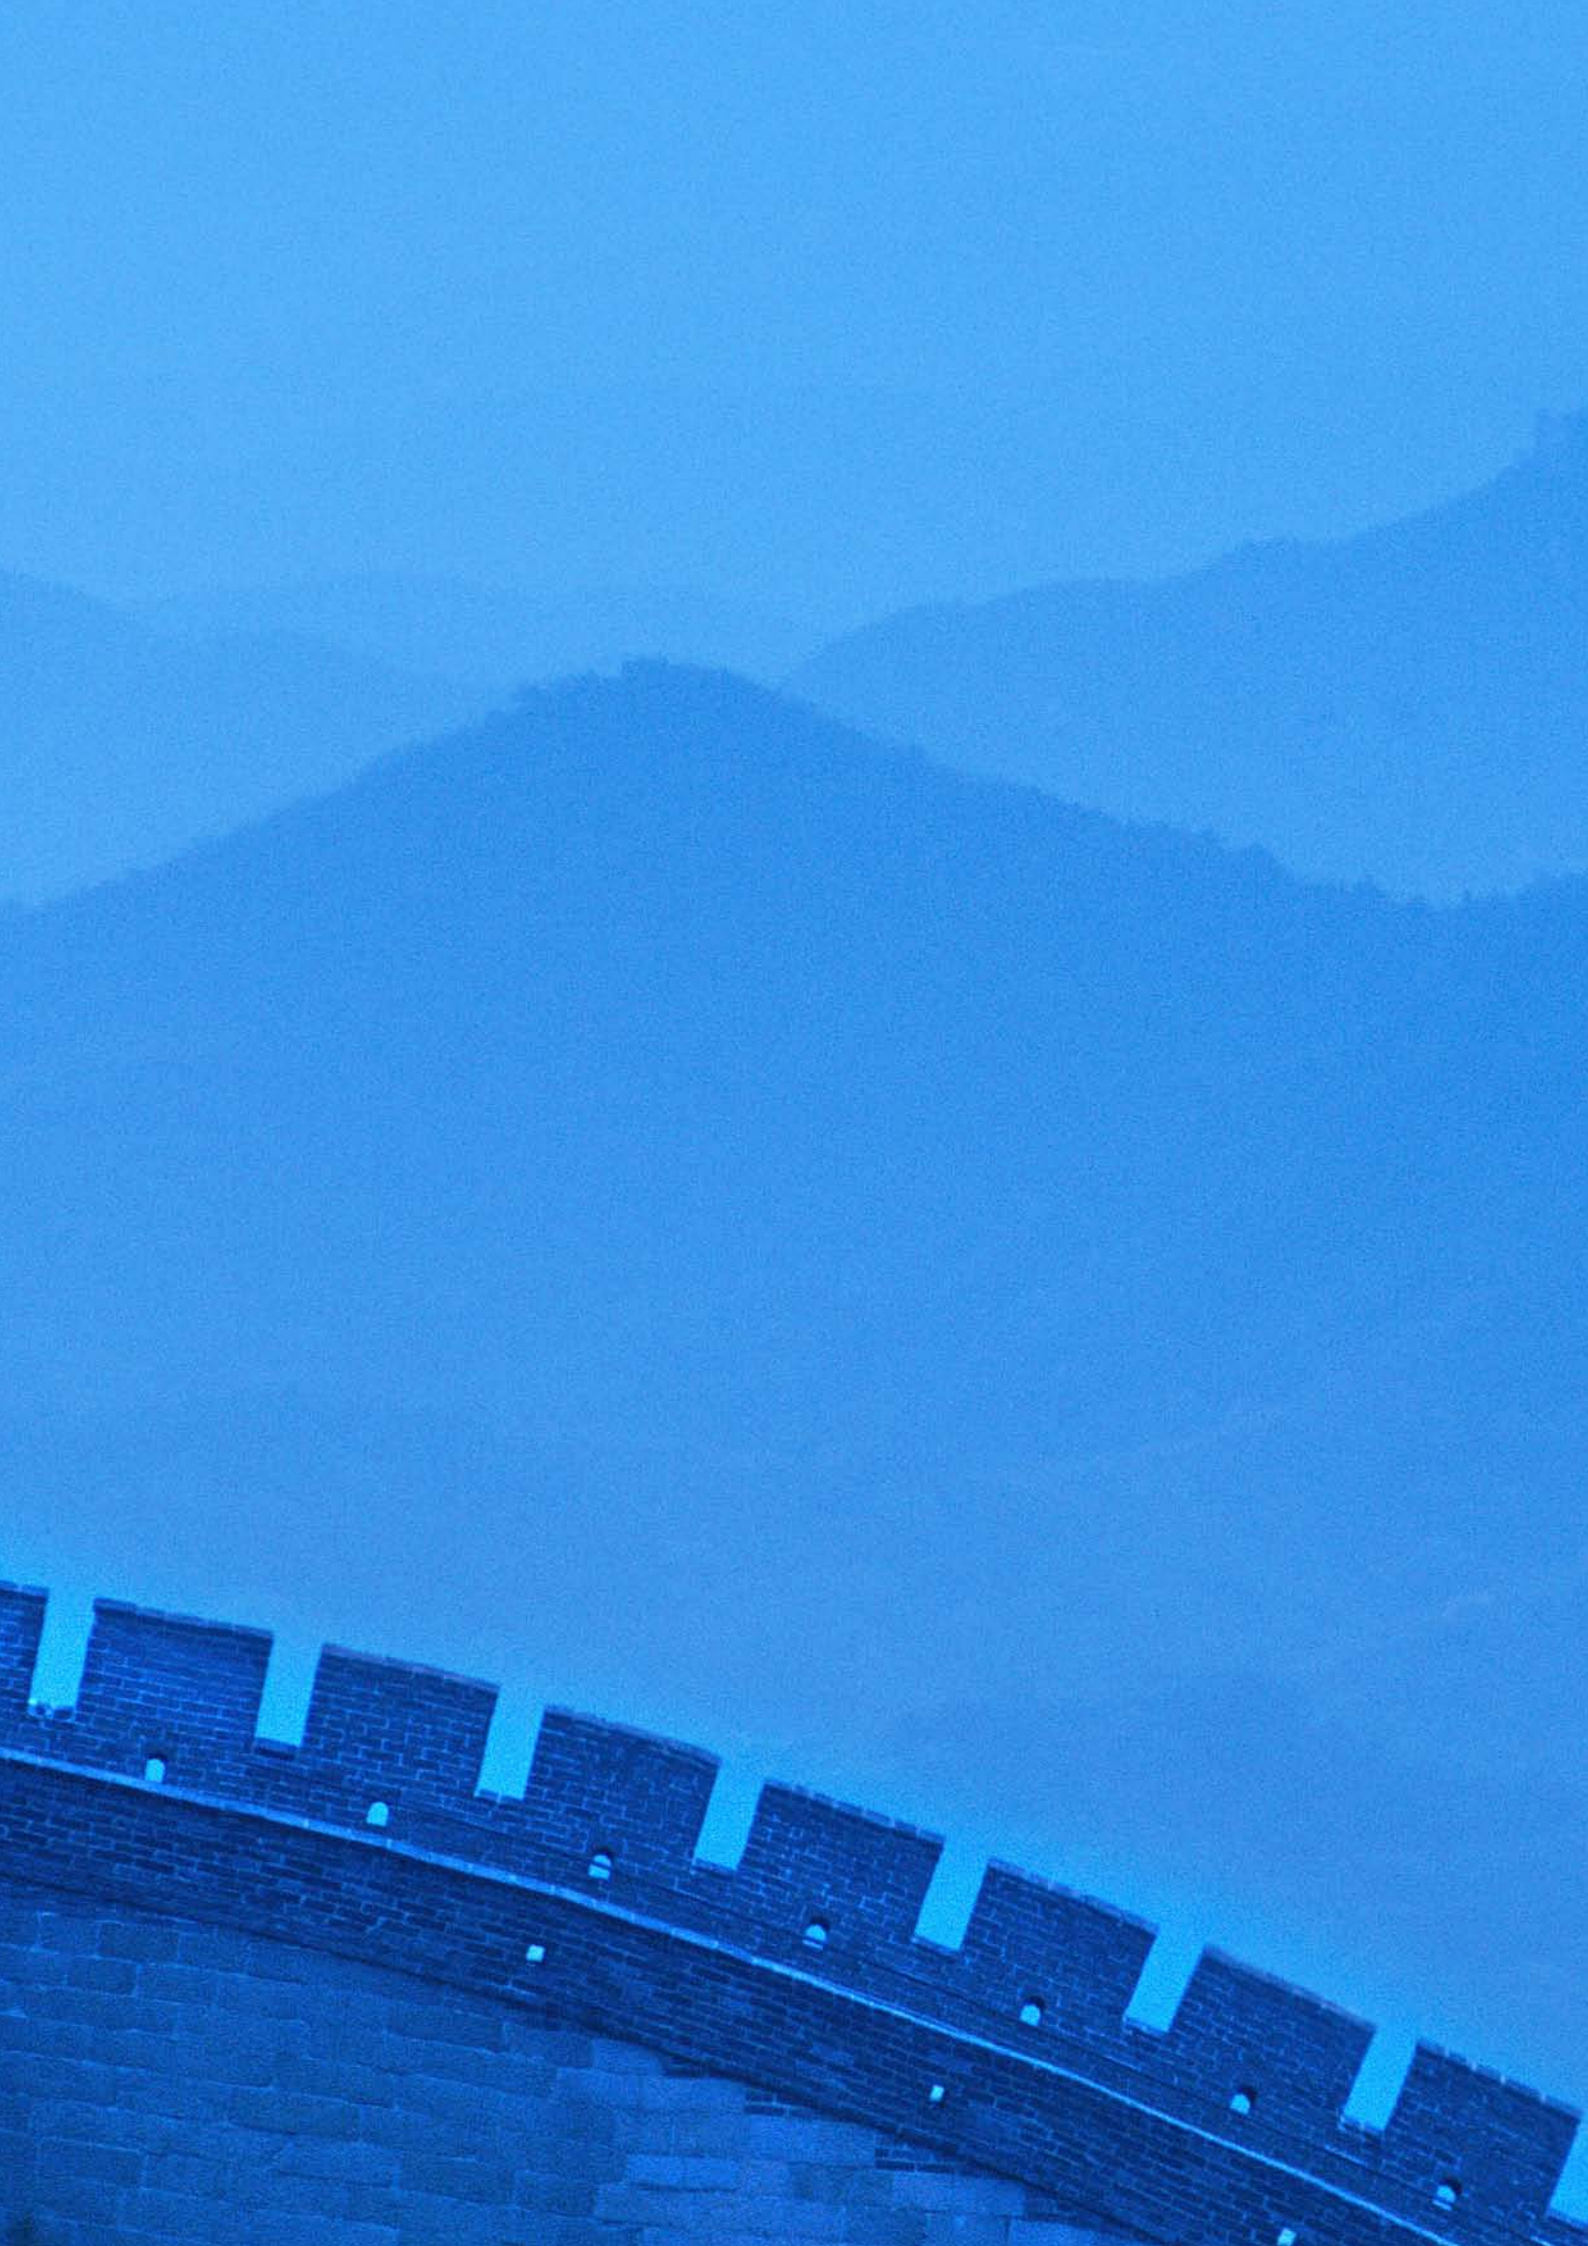

www.sap.com/contactsap

SAP

The Best-Run Businesses Run SAP™