Technical white paper

Enterprise Solution Design for SAP HANA

Table of contents

Author's note: This paper is written based on the best information available at the time (February 2013) and based on the knowledge and experience of the author and the review committee. The document is an opinion piece and should not be used as the basis for the technical design of entire solutions. It can, however, be used as a guide to enterprise solution design approaches.

Introduction

SAP HANA is a powerful, in-memory computing platform that streamlines business suite applications, analytics, planning, predictive analysis, and sentiment analysis on a single platform, so businesses can operate in real time. The design approach for enterprise-level solutions involving SAP HANA, and the best practices surrounding them, isn't intrinsically different from the approach to any other enterprise-level solution for technology implementations. This paper is written to address those elements of good solution design and apply them to the SAP landscape, with particular focus on the SAP HANA element.

It's important to note that not all SAP HANA implementations need to be enterprise-level. Locally developed and managed solutions based on SAP HANA are entirely acceptable and encouraged. But if your organization is qualifying SAP HANA for use across the enterprise, a consolidated approach and enterprise-quality design will help ensure a cost-effective model that's implemented using best practices. And that's likely to result in a solution that delivers both the increased longevity and decreased TCO you seek. Implementing a cross-organization solution while not following the principles of good enterprise design will result in a solution that is costly to maintain and difficult to manage and grow.

The three principles of good enterprise design

Despite the fact that HANA is an appliance, its incorporation in an enterprise solution means the implementation must align to the three core principles of good enterprise design: stability, scalability, and agility.

Enterprise-level stability

Stability incorporates the concepts of system stability from a data center perspective. And it focuses on security, maintainability, and supportability from a technical support perspective. That means an enterprise-level solution should be maintainable within existing data center "support-and-run" structures, which help ensure that any growth and subsequent maintenance do not incur spiraling complexities and the elevated costs associated with them. Furthermore, your approach to security should be lean, rigorous, and automated; you want no access breaches and no audit breaches.

Enterprise-level scalability

Any enterprise solution should be designed initially to allow for future growth, both planned and unplanned. Is "planning for the unexpected" a contradiction in terms? No, it's a key factor in any enterprise IT solution design. Any design classified as "enterprise-level" should include the expectation of data volume growth in planned functional areas. It should anticipate entirely new functional areas as well.

This doesn't mean the infrastructure design should from the outset incorporate volume capacity for the entire sprawling organization. But be sure to consider the ability of your solution to scale with less disruption to live environments. Aim, too, for the most cost-effective model that can encompass alternative approaches like virtualized environments, hosted solutions, and solutions as a service.

Enterprise-level agility

For the administrator maintaining a large enterprise landscape, "flexibility" and "agility" can be frightening words. But they shouldn't be. These words don't need to mean the loss of control. They needn't suggest it will take uncontrolled mutation of the enterprise environment to enable the support of requirements outside the initial design scope. Rather, they should be positive terms, referring to the ability of the solution to easily grow, transform, and evolve in line with the modernization of the entire landscape. A flexible and agile solution will mean there is no stagnation during its evolution that might cause a failure to meet the requirements of business users.

Start with a self-assessment

If you're constructing an enterprise solution design, ask yourself the following questions. If you can answer them confidently, then you're on the right track.

Stability

- What is the availability service-level agreement (SLA) for your solution? Is your hardware provider or hosting partner committed to support you to that level? What is your disaster tolerance requirement?
- How does this solution integrate with your organization's support structures for other elements of the enterprise landscape?

- Have you clearly defined a maintainable, automated security concept that aligns to your enterprise security requirements? Have you documented it and have the relevant stakeholders signed off on it?

Scalability

- What is the design principle for adding functional areas to the initial design?
- Will the solution require intercompany billing? Are there other chargeback/showback requirements for business units migrating to the stack, or developing on it?
- What's the approach for protecting against unauthorized viewing of secure data from specific functional areas?
- What are the projected data volume and user growth? If these projections are exceeded, what is the solution scaling plan?

Agility

- What are the initial boundaries of scope for development on this element of the landscape?
- How does your modeling practice ensure that your solution can evolve for other purposes?
- For development outside these boundaries, what is the change request process and who are the key signoff parties for that change?

You should be able to answer most, if not all, of these questions. If you can't, then your solution isn't yet an enterprise solution.

Figure 1. Enterprise solution design for SAP HANA from HP
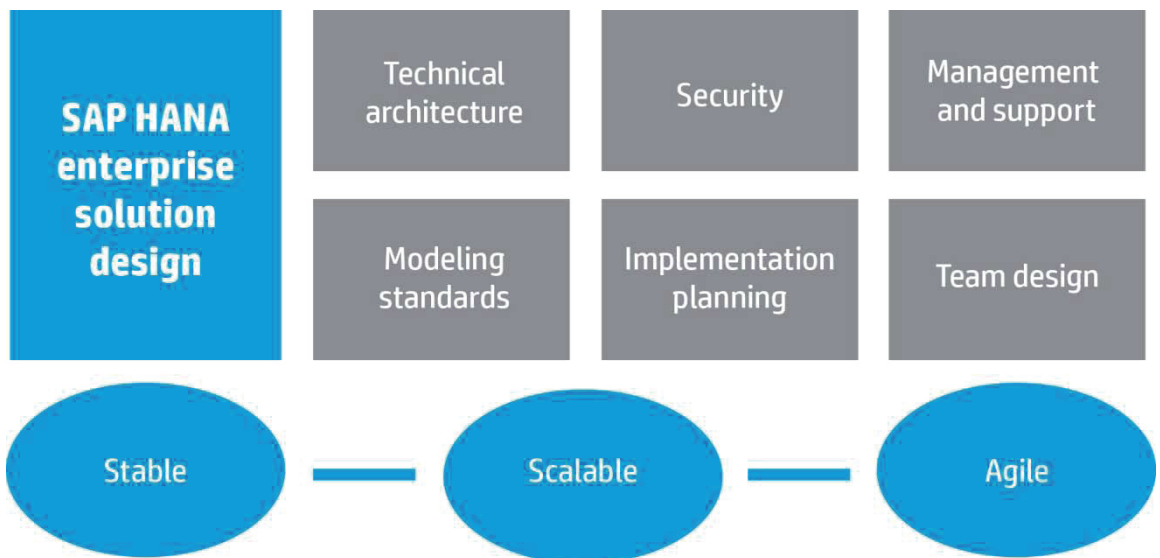
Key design considerations for your SAP HANA solution

Figure 1 shows the key considerations for an SAP HANA solution design. As with all designs for technology-based solutions, the starting point should be the business needs you want to address. The implementation then should focus on these needs, and should ultimately improve bottom-line financial results. Until you've identified and documented the business pain points, the need for a solution won't be clear—which means the justification for embarking on a search for technology to eliminate these pain points won't be clear, either. This is the case with SAP HANA as with any other implementation.

So the first step in solution design is to document the business pain points and their related technical requirements. Then weigh the ability of the proposed technology to address those pain points and requirements. These two steps will help ensure that your design aligns with your business needs.

The solution design should integrate the technical design—at a high level—with the implementation approach and deliverables. It should include the selection of technologies within the solution stack and the method of integration of those technologies. And it should contain references to the best practices or design methodologies to be used during implementation.

Any successful design will be the result of careful attention to details like architecture, security, management and support, modeling standards, implementation planning, and team design. Let's look more closely at what you'll need in each of these six areas of design focus.

A capable, scalable architecture

In your design, take into account the entire landscape, and not just the size of the HANA appliance required. Of course, the size of the appliance is important. But equally important are all other elements of the solution through which data will pass on its way to the end user. If you have a bottleneck at any other point in the data flow to the end point, your investment in the HANA element will have been wasted. In a broad sense, data residing has a source from where it must be extracted and a target, or front-end, to which it must be relayed. Take a good look at the entire path:

• Ensure that your network can provide the required throughput

• Make certain that the application servers can adequately handle the amount of data and the user concurrency

• Confirm that any portal or mobile servers are up to specifications, and that you've considered the impact of this solution on them

To help ensure that you uncover potential issues with these components, design stress tests that apply the extremes of data replication, loading, and user concurrency you expect. Include tests that will push the limits of application server and network throughput, as well as the ability of the HANA software to crunch the data. Testing strategies generally include unit, volume, regression, concurrency, and user acceptance. A test methodology should be defined during solution design.

Scaling your HANA environment for future growth is important in the design of an enterprise solution. This scaling extends beyond the HANA appliance itself and includes the application servers and portal servers used in the solution landscape.

You can carry out initial sizing in several ways. The most accurate is with the SAP QuickSizer tool. True, this tool takes time and effort to populate, but it will be effort well spent, as the tool will provide an accurate sizing.

Clearly document all availability requirements and SLAs with all parties involved in system setup, support, and operation. By so doing, you'll be better able to understand the expectations of the landscape, helping to ensure that the infrastructure solution design will be fit for its purpose and can meet the demands put on it. Especially important are documentation for the high-availability (HA) and disaster tolerance consideration of the landscape. Document these elements and discuss them with hardware vendors, data center managers, and applications consultants. The CIO should sign off on these requirements based on the needs of the business.

Consider using SAP HANA AppSystems from HP in your solution. These are based on either rack servers (the HP ProLiant DL series) for what are termed "scale-up" solutions, or blade servers (based on HP ProLiant BL series blades, IBRIX shared file system, and SAN storage) for "scale-out" options. HP has a history of designing some of the most resilient and stable server solutions available on the market.

Scale-up covers requirements up to 1 TB of RAM. While there is some ability within the system options to grow your RAM volume, this is limited and RAM cannot scale past 1 TB. Our "scale-out" solution offers linear scalability—currently up to and beyond 20 TB, in 2 TB increments. These HP scale-out solutions offer high availability and synchronous disaster tolerance up to 50 km. At this writing, we're in the process of developing an asynchronous disaster tolerance solution that extends beyond 50 km.

Appliances for nonproduction (and non–disaster recovery [DR]) instances can support multiple system installations. This means you can install both development and QA environments on one appliance. The memory allocated to each of the instances can be controlled in the configuration settings of the database via the HANA studio, which allows for economies of scale in the purchase and setup of the environment. Your DR environment can host development or test environments; however, additional storage for the file systems is required and the DR system will take complete priority during a DR situation.

You'll find different solutions from other vendors, of course. When evaluating different solution designs from an infrastructure point of view, consider the same criteria around sizing and availability.

A strong security and authorization framework

As with all enterprise solutions, SAP HANA needs a strong security solution, one that will meet your internal standards in place for information security. The solution should include the obvious firewall security of the data center, as well as the user authorization components managed in the application layer.

The implementation of the enterprise security solution needs to:

• Improve availability and protection of networks, systems, software, and database management system (DBMS)

• Protect data at rest, data in transit, and data in motion

• Maintain a secure audit trail across the organization

It's imperative that any new database or solution, such as SAP HANA, is integrated into your enterprise security framework to help ensure that the new solution does not become the weak link. This will be especially important where the solution is exposed in any way to data feeds from outside the corporate firewall.

From an application perspective, user authentication should make use of Kerberos or SAML for single sign-on. At the time of writing there is no officially approved approach for the automated addition of analytic privileges to users within

HANA, whether based on tables or hierarchies for specific objects such as profit/cost centers, organization units, or other criteria. Workarounds have been tested in which SQL programs run and assign privileges based on table entries.

It is always advisable to automate security authorizations in order to avoid human error when creating users and assigning access rights. Although at the time of writing there is not yet an approved method of table-based security, this is expected to change in the near future. Discuss with SAP or your systems integrator (SI) to determine if this is available at the time of your implementation.

For other SAP applications based on the HANA database, it is generally not necessary to implement HANA security for information consumers, because this is managed through the application (such as the SAP BW or Business Suite applications). Best-practice security approaches for these applications do not generally change with HANA and should be followed as with any other database.

Integrated management and support

The integration of SAP HANA management and support into existing enterprise agreements within your IT support and run structure is central to a positive and consistent user experience. Management and support is a standard component of all enterprise solution implementations. But its consideration is often left to the last minute, which can cause issues as you approach the go-live date.

For this reason, design the support and management approach for any new element of your enterprise landscape at the outset of the implementation program. This will help ensure that the support team has enough time to incorporate the new technology skills required to support the solution.

We can help with HANA management and support. In fact, support of the SAP HANA appliance up to and including the OS level is easily covered by HP support services. You can treat these services as a single point of contact for HANA support; they will redirect to SAP any issue seen as lying at the application layer, rather than the appliance layer. Both the software and the appliance support can be integrated with existing SAP and HP support structures.

We also offer management and run services for SAP HANA. These come in two key areas:

• ITO covers the management and running of the solution up to the OS layer. It is consistent with HP ITO services for entire SAP landscapes, and covers all elements of the solution stack up to the operating system. This gives you a consistent best-price landscape management option for large environments.

• From the applications management perspective, we offer the same services for SAP HANA as for all other elements of the SAP landscape. Applications management integrates with ITO and covers all elements of the landscape above the operating system. Thanks to economies of scale, we're confident our applications management for SAP offers you the best price solution for entire SAP solution landscape, including SAP HANA.

Whether you manage the system with internal IT teams, or contract it externally, plan early for the incorporation of HANA. Consider the costs associated with it and the skills required. And agree upon all structures and handoff between parties to deliver rapid solution deployment and reduce any risk from poor support after the go-live date.

A consistent modeling approach

Whether you are implementing a HANA native solution or an application running on the HANA database, we recommend following a consistent approach for modeling that is in line with the best practices for the given application. For SAP NetWeaver BW, this is Layered Scalable Architecture (LSA) or some variant of it that suits specific requirements. A modified version of LSA has been developed for SAP NetWeaver BW solutions running on HANA. For HANA modeling directly, refer to the best-practice documents.

To provide consistency, it's important to formally document the modeling approach within your organization. This will benefit the development team, because team members will be able to pick up and modify models from other members. Good documentation also lessens the risk of losing key knowledge when a developer leaves. Moreover, documentation helps ensure that best practices are followed, which will likely yield the best system performance.

Many best practices for development on HANA are published in SCN and on the Experience SAP HANA website, though they don't always cover all areas. Consider the potential requirements from the data flows you will build and determine the best practice of modeling for each one.

For example, if you have a real-time requirement for data, then SLT is recommended. But this allows very limited transformation at load time, so you'll need to build any data manipulation, if required, into a calculation view. This may impact the end performance of the reporting layer, as the manipulation occurs at runtime. If having a batch load every 30 minutes will suffice, then it may be more beneficial to load using data services and create the transformation in the extraction, transformation, and load (ETL) tool. This will remove the requirement for heavy Structured Query Language (SQL) at runtime and improve report performance. Considerations like this are crucial to a successful implementation.

Document the modeling standards, or list references to specific documents owned by the software company or systems integrator (SI) that developed them. This helps ensure that as you add new developers to the team, you'll have a point of reference for those individuals. Building based on a standardized approach results in data model consistency, and ultimately you'll have a solution that's easier to maintain.

Planned and phased implementation

Unstructured solution implementations are likely to fail. This is the reason that most SIs follow a prescribed methodology in project implementations for technology-based solutions. These methodologies act as accelerators for successful, structured project execution and provide a consistent and robust approach. They increase the likelihood of project completion within planned timelines and, most importantly, within budget. Follow an implementation methodology, whether it is your own or one designed and shared by an SI or SAP itself.

When implementing on new technology, the "Big Bang" approach to implementation is not generally the recommended approach. (This is nothing new, but it is good to reiterate.) So if you're building an SAP HANA solution that is not using the existing SAP applications such as BW or Business Suite for the application layer, we strongly advise you to take a phased, slow rollout approach. Implement a small number of solutions, or SAP Rapid Deployment Solutions (RDSs), if appropriate, in each phase. This will help ensure that the audience receiving the solution is limited for the initial go-live date, allowing all areas of the support team to skill up at a comfortable pace.

RDSs from SAP have internally built methodologies for their implementations. These are extremely refined to the specific build, and as such are an invaluable tool for the implementation of RDS projects. Of course, an RDS is unlikely to be the whole story for your solution, so plan beyond its boundaries if that's what you will need.

Our HP methodology follows all phases for the implementation of an enterprise-level business intelligence (BI) solution. It's been tried and tested across many software and database technology combinations. You can use it for projects of any size. And it has built-in templates for all phases of project execution, from requirements documenting, to scoping, project planning, design, build practices, test, and launch. This approach lends structure and discipline to the implementation—two keys for building consistent, maintainable, supportable solutions.

Skilling the team

The skills required for the implementation and support of an enterprise-level SAP HANA solution depend on the use case selected. Remember, you can choose to implement SAP HANA in different ways:

- Within the Business Suite landscape as an accelerator
- Underneath BW as the database for that solution
- Underneath Business Suite as the database for that solution (currently in ramp-up)
- As a standalone operational data tool for analytics with front-end tools attached
- Or via an RDS from SAP

Each of these will require slightly different skills in order to build the solution and subsequently support or run it.

When following the approach of using SAP HANA within or underneath existing SAP applications, the only area of specific new skills (or potentially new skills) that the existing team may not be able to cross-skill into are those at the data center and Basis level. This is because of the need to understand the functionality and workings of the HANA database and the requirement to use SUSE Linux.

In reality, cross-skilling into SUSE from other Linux operating systems (OSs) should not present a huge challenge for your OS maintenance team. And HP and SAP can provide the training for the Basis team. These individuals will need to understand the workings of the appliance in order to support and maintain the elements of the system not covered in support packages from SAP or HP, or not covered by the ITO or application management solutions purchased from third parties. Other Business Suite skills and BW skills will transfer well, whether the database underneath them is SAP HANA or something else.

When working on the application side of HANA for SAP Rapid Deployment Solutions, or when implementing bespoke solutions in which the modeling is carried out in the SAP HANA Studio, you'll need database modeling skills consistent with RDBMS skills, as well as resources skilled in SQL. That's because much of the logic is written in SQL; and while the wizards are useful for some coding, more complex elements must be coded manually. Also, when performance issues arise, being able to debug and manipulate SQL really provides you with significant additional value.

Summing up

The differences between implementing a local solution and an enterprise-level solution are significant, and the effort involved is on a different scale entirely. However, your upfront investment in making the solution an enterprise-level one, if that is what you need, will ultimately result in savings in time and effort. It will also reduce costs in the long run.

Attempts to scale a local solution without the structure to do so, or the cost of an uncontrolled, unintegrated support structure, can quickly get out of hand. Conversely, once you design and define the standards and best practices for an enterprise solution, you'll reap benefits as the system grows and developers need to work across the solution.

The key thing to remember: if your solution is planned as an enterprise solution, it's critical that you invest early in its design. Define the reference architecture, development standards, modeling and security approaches, and implementation methodologies up front, then follow your plan all the way to delivery and operation.

For more information

Learn more at
hp.com/go/bi

Share with colleagues

Rate this document

4AA4-4454ENN, February 2013, Rev. 2